

# Contents

Introduction.....	4
Agenda .....	5
Logical Partitioning Concepts.....	6
Defining LPAR .....	7
LPAR Terminology .....	8
Defining LPAR (cont'd) .....	9
Advantages of Partitionable Servers .....	10
Physical Partitioning .....	11
Logical Partitioning .....	12
LPAR Requirements .....	13
Hardware and Firmware .....	14
Hardware and Firmware Mechanisms .....	16
Hardware Management Console.....	17
Applications in an LPAR Environment .....	18
Logical Partitioning Details .....	19
Logical Partitioning Details .....	20
Planning for LPAR .....	21
Dividing the Server .....	22
CPU Resources .....	23
Memory Resources .....	24
LPAR Memory Overhead .....	25
Memory Resources .....	26
Partition Page Table Size .....	27
TCE Table Size .....	28
I/O Resources.....	29
I/O Resources.....	30
I/O Resources.....	31
I/O Resources.....	32
IP Addresses for LPAR Mode .....	33
Hardware Management Console.....	34
Hardware Management Console.....	35
Hardware Management Console (continued) .....	36
Hardware Management Console (continued) .....	37
Hardware Management Console.....	38
Remote HMC Support .....	39
LPAR Setup with HMC .....	40
Hardware Management Console.....	41
Hardware Management Console.....	42
HMC Login .....	43
HMC GUI .....	44
HMC GUI .....	45
HMC System Manager Security .....	46
HMC User Management .....	47
HMC Problem Determination.....	48
HMC Software Maintenance .....	49
HMC Software Maintenance .....	50
HMC System Configuration .....	51

HMC System Configuration .....	52
HMC GUI .....	53
HMC GUI .....	54
HMC Inventory Scout.....	55
HMC Inventory Scout.....	56
HMC Service Focal Point .....	57
HMC Service Agent.....	58
HMC Service Agent.....	59
Working with Partitions .....	60
System Profiles .....	61
Creating a System Profile .....	62
Creating a System Profile (continued).....	63
Creating a System Profile (continued).....	64
Creating a System Profile (continued).....	65
Creating a System Profile (continued).....	66
Creating a System Profile (continued).....	67
Starting a System Profile .....	68
Starting a Specific Partition .....	69
Partition Activation Failures .....	70
Partition Activation Failures .....	71
Partition Activation Failures .....	72
Partition Activation Failures (continued).....	73
Opening a Console.....	74
Opening a Console (continued) .....	75
Virtual Console .....	76
Stopping and Restarting Partitions.....	77
CEC Properties.....	78
CEC Properties (continued) .....	79
CEC Properties (continued) .....	80
Resetting a Partition.....	81
Resetting a Partition (continued) .....	82
Deleting a Partition Profile .....	83
Deleting a Partition Profile (continued).....	84
Powering Off.....	85
Powering Off (continued) .....	86
Powering Off (continued) .....	87
Powering Off (continued) .....	88
Dynamic Logical Partitioning on a pSeries server .....	89
Defining Dynamic LPAR .....	90
Moving LPAR Resources .....	91
Moving LPAR Resources (continued).....	92
Monitoring DLPAR Operations.....	93
Hardware Management Console.....	94
Summary .....	95
Hotlinks/References.....	96
Trademarks .....	97

## Author Biographical Information

### **Sebnem Jaji, IBM**

Sebnem Jaji is a senior software engineer in the IBM eServer Solutions Enablement organization. Currently, she is working with independent software vendors to develop and optimize software solutions for the IBM eServer platforms. She has over 15 years experience in operating systems and application design and development.

### **Kim Greene, President, Kim Greene Consulting, Inc.**

Kim Greene is president of Kim Greene Consulting, Inc. specializing in Domino™ for iSeries consulting, customized education, and resolving performance issues. Kim has more than 17 years' experience with IBM midrange server platforms and also does technical writing for a number of different projects, recently including an IBM whitepaper "Experiences with Lotus Domino and AIX Dynamic Logical Partitioning." Kim is a frequent writer for technology magazines and speaks at many conferences worldwide. Prior to starting her own consulting business, she was employed at IBM in Rochester, Minnesota, where, among other responsibilities, she worked on several areas of operating system performance in the OS/400 development laboratory.

# IBM eServer pSeries: Logical Partitioning (LPAR) Management

## Introduction

Welcome to this on-line course “IBM® eServer™ pSeries™: Logical Partitioning (LPAR) Management.” This course takes you from A to Z in working with LPAR on a pSeries server. You will start by learning the many benefits of LPAR and then learn how best to work with this exciting partitioned environment on a pSeries server.

If you have not already taken the on-line course, “IBM eServer pSeries Logical Partitioning: Installation and Configuration”, we strongly encourage you to do so prior to taking this course.

*[NOTE: This course is not intended to teach the fundamentals of LPAR, the references at the end of the course provide this introductory information along with details that help support the materials presented here.]*

## **Agenda**

**Logical partitioning concepts**

**Planning for LPAR**

**Hardware management console (HMC)**

**Working with partitions**

**Dynamic logical partitioning on a pSeries server**

**LPAR information resources**

## **Agenda**

This on-line training course provides a comprehensive look at utilizing LPAR on a pSeries server. We will start out by discussing some logical partitioning concepts to ensure you have a solid basis for the many reasons you may want to implement LPAR. Next, we move into planning for LPAR. In this section, we will cover memory, CPU, and I/O resource assignment.

The hardware management console (HMC) is required to work with an LPAR environment on a pSeries server, so we will show you how to work with this console. The majority of this on-line course shows you how to work with and manage LPAR partitions. You will learn how to create, start/stop, work with, and remove partitions.

One of the newer features available on a pSeries server is the ability to dynamically move resources between LPARs. We will show you the many benefits of dynamic LPAR. We conclude this course by showing you a number of pSeries LPAR information resources.

By the end of this course, you should be very well versed in how to manage and work with your pSeries LPAR environment.

## **Agenda**



**Logical partitioning concepts**

**Planning for LPAR**

**Hardware management console (HMC)**

**Working with partitions**

**Dynamic logical partitioning on a pSeries server**

**Summary**

## **Logical Partitioning Concepts**

Here is a brief outline of what we will be reviewing in this course. As you can see, we will start with some logical partitioning concepts.

## Defining LPAR

### LPAR means software isolation

- No program access between partition memory and I/O spaces
- Software exceptions/crashes are contained within a partition
- Partitions are independently rebootable



### LPAR is not hardware fault containment

- Some core resources are still implicitly shared (PPC, fan, etc...)
- A few failures in shared hardware can stop all partitions
- POWER4 introduces many improvements for fault containment
  - Fewer global checkstops, more errors handled by inline recovery - like S/390!

### Security Certification

- Partitions are designed for isolation and security

## Defining LPAR

Logical partitioning allows you to carve separate system environments from one physical server, providing total software isolation between partitions. Programs cannot access memory or I/O spaces on other partitions. Additionally, problems on one LPAR do not affect other partitions. Software exceptions or partition crashes only affect the partition where the exception or crash occurs. Because of this partition isolation, each can be rebooted without causing any impact to any other active partitions.

It is important to understand that LPAR is not creating a redundant environment. It is not creating a complete backup for a particular hardware system. There is a small exposure to hardware failures which affect more than one partition, although advances in Regatta system design have significantly reduce this exposure. Core system resources, such as the processor power component (PPC) or the fan, will impact the whole server. If any of the core components has a failure, all partitions will be affected.

The new POWER4™ pSeries processors contain many improvements for fault containment—such as fewer global checkstops and greater error handling by inline recovery. For example, a bad word in memory (double-bit error), which caused a global checkstop in previous servers, now propagates to the requesting processor with a special flag on the data. The processor takes a less critical server check interrupt, and AIX® decides whether the bad memory value is critical (perhaps causing a partition crash) or in a user space (which terminates the failing process). These enhancements, which are also present in the IBM eServer zSeries® platform (formerly S/390®), provide greater partition stability and reliability to ensure application isolation and security.

With LPARs providing software isolation between partitions, they also provide security certification. Resources within each LPAR such as applications, enterprise data, user accounts, and others are protected. Unless an application or user has access to resources in a particular LPAR, those resources cannot be accessed.

Now that we covered briefly what LPAR is, let's look at specific LPAR terminology.

## LPAR Terminology

### LPAR

- Flexible deployment of compute resources
- More efficient use of compute resources

### NUMA

- Servers which can be grown incrementally over time
- Cost-effective approach for large, single-image computing

### Putting LPAR and NUMA together

- An integrated design that exploits the advantages of both
- Supported by unifying technologies
  - System management, high-speed interconnect, packaging, multi-path I/O subsystems
- A server that is truly **flexible**.

## LPAR Terminology

Logical partitioning provides flexibility in the deployment of computing resources. By allowing the allocation of resources such as memory, CPU, I/O, and bus bandwidth to specific LPARs, you are able make the most efficiently use the resources of your pSeries server.

The architecture of non-uniform memory access (NUMA) means it will take longer to access some regions of memory than others. This is because some regions of memory are on physically different busses from other regions. NUMA was designed to improve the scalability of symmetric multi-processing (SMP) environments. All memory accesses are posted to the same shared memory bus with SMP. This architecture can cause memory access to be a bottleneck, because of the contention on the shared memory bus. NUMA alleviates the bottlenecks of SMP by limiting the number of CPUs on any one memory bus.

The combination of LPAR and NUMA provides an integrated design that exploits the advantages of each architecture. This provides unified support of system management, high-speed interconnect, and I/O subsystems. The end result is a server that offers the best performance and one that is truly flexible.



## Defining LPAR (cont'd)

### **Symmetrical MultiProcessor (SMP)**

- Traditional computer system with multiple processors
- Contains processors, memory, and I/O
- All the resources required to run an operating system

### **Logical Partitioning (LPAR)**

- Partition the resources in an SMP system
- Each partition can support an operating system
- Mechanisms provide strong protection between partitions

### **Non-Uniform Memory Access (NUMA)**

- Connect multiple SMPs together into a larger, unified SMP
- All processors can physically access all memory
- Access time from processors to memory varies by location

## Defining LPAR (cont'd)

Now that we have looked at some basic LPAR terminology, let's compare SMP, LPAR, and NUMA.

We mentioned symmetrical multi-processing (SMP) on the previous page. SMP processing involves two or more similar processors connected via a high-bandwidth link and managed by one operating system, where each processor has equal access to I/O devices. All processors can perform kernel, I/O, and computational operations. The operation system manages a pool of identical processors, each of which may be used to control any I/O devices or to reference any storage unit. Since all processors are equal, the server can be reconfigured in the event of a processor failure, and following a reboot, continue to run. Performance can be optimized and more processors can be used.

LPAR creates independent servers within a single physical server. LPAR allows you to allocate processors, memory, disk storage pools, I/O processors, and LAN/WAN communication to each partition. Each partition is isolated from the other partitions and workload performance on one partition does not affect the workload performance on any other partitions.

As we defined previously, NUMA allows multiple SMPs to run together in a larger environment. All processors can access all memory. The biggest difference with NUMA is that some regions of memory are on different busses from other regions of memory, making memory access non-uniform. This makes access time from processors to memory variable, depending on which region of memory is being accessed by the application.

## **Advantages of Partitionable Servers**

### **Solution for test / migration issues**

- Avoids the need for dedicated test servers
- Allows testing on the same server used for production
- Introduction of new operating systems / applications

### **Provides dynamic capacity demands**

- Applications with rapidly changing workloads
- Workloads that peak (or simply run) at different times
- Rapid deployment of new servers/applications
- Enable both horizontal and vertical growth in the same server
- Long-term capacity planning (headroom)

### **Sharing of computer resources**

- True dynamic sharing, or simply one at a time
  - Summary: Partitioning can help you get the most efficient use out of your server hardware

## **Advantages of Partitionable Servers**

There are many benefits to implementing LPAR, from providing a solution to testing and migration issues to providing dynamic capacity demands. Implementing one or more LPARs avoids the need for stand-alone, dedicated test servers. Rather than having an additional physical box to integrate into your operating environment, you can create an LPAR and be ensured this test environment is isolated from any production activity. An additional benefit you get from this setup over implementing a dedicated, stand-alone test server is that you are doing your testing on the same hardware that is used for your production environment. This makes testing new applications and fixes easier, which at the same time ensuring you are not directly impacting critical business services.

LPAR allows you to introduce new operating systems to your server as well. As an example, you could have AIX 5.1 installed on one LPAR and AIX 5.2 installed on a second or third LPAR. Not only do you get the added benefit of being to work with multiple versions of the operating system, but you can also introduce new applications to your enterprise.

In summary, LPAR can help you get the most efficient use out of your server hardware by allowing for true dynamic sharing of resources. You no longer have the environment where the purchase of a server causes you to run one version of the operating system and have all applications running with the exact same configuration.

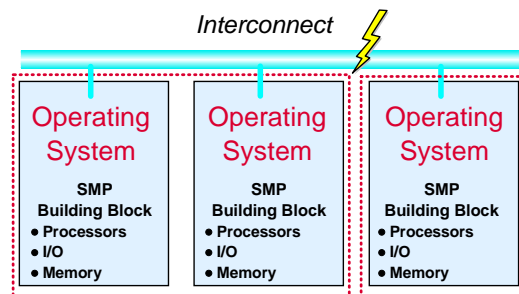
# Physical Partitioning

## Multiple memory coherence domains

- Each with an OS image
- Separation controlled by interfaces between physical units

## Strong software isolation, strong hardware fault isolation

- Granularity of allocation at the physical building block level
  - Resources allocated only by contents of complete physical group
  - No sharing of (internal) resources between partitions



# Physical Partitioning

There is a lot of confusion in the market about different types of partitioning. At IBM where both types were invented, partitioning is defined as Physical and Logical.

Physical partitioning involves connecting building blocks together. Each “group” or partition of building blocks share memory coherency like an SMP. This means the image of memory is kept consistent across all processor caches and memory subsystems within the group. Anything outside the group is fenced off.

With physically partitioned servers, you have multiple memory coherence domains. Each of those domains contains one version of operating system code. Physical partitioning is done without any special support in the processor or memory subsystems. A physical partition has a specific version of operating system (OS) installed, and all applications share that same operating system code when executing.


## Logical Partitioning

### One memory coherence domain with multiple OS images

- Separation controlled mainly by address mapping mechanisms

### Strong software isolation, fair-to-strong hardware fault isolation

- Granularity of allocation at the logical resource level (or below)
  - Resources allocated in almost any combination or amount
  - Some resources can even be shared



## Logical Partitioning

Unlike physical partitioning, logical partitions have one memory coherence domain with multiple operating system images. Separation is controlled by address mapping mechanisms. Logical partitioning offers strong software isolation, allowing you to install a different version of the operating system in each partition.

Logical partitioning provides finer granularities than offered with physical partitions. Each logical partition can have a different combination of resources allocated to it. Some resources can even be shared by the LPARs. As depicted in the diagram on this slide, I/O adapters and devices can be assigned to specific logical partitions, along with processor resources and memory. There are virtually limitless combinations of resources that can be allocated among logical partitions.

## **LPAR Requirements**

### **Requires LPAR-enabled version of AIX**

- Version 5.1 and above
- Includes AIX changes to use Hypervisor calls
  - Virtual Memory Manager calls Hypervisor to maintain Page Tables
  - Debugger calls Hypervisor to access specific memory and registers
  - Special "Virtual Console" device driver calls Hypervisor

### **LPAR configuration changes**

- Configured and monitored through a hardware management console (HMC)
- In AIX 5.1 configuration changes require partition reboot
  - Only the partition(s) affected by the configuration change
  - No impact to other running partitions
- LPAR dynamic reconfiguration in AIX 5.2


## **LPAR Requirements**

Enabling LPAR on a pSeries server requires AIX version 5.1 or higher. LPAR requires new interfaces in AIX that are supported starting with the 5.1 release. Previous version of AIX will not run in a partition. LPAR requires Hypervisor calls to support virtual memory management to maintain page tables, debugger calls to access specific memory and registers, and a virtual console device driver specifically required for LPAR.

The hardware management console (HMC) is required to configure and monitor an LPAR environment on a pSeries server. Depending on which version of operating system you have installed, you have different configuration changes available for working with LPAR. AIX 5.1 configuration changes require a reboot of the partition being changed. All other partitions are not impacted by changes to the partition undergoing configuration changes.

Starting in AIX 5.2, dynamic LPAR has been introduced. This allows resources to be moved without having to reboot partitions. Now you can add more memory, disk, or CPU to an LPAR in a 24x7 environment, offering a very flexible operating environment.

## Hardware and Firmware



**Based on enhancements to PowerPC architecture**

### **Processors: Real Mode Offset and "Hypervisor Mode"**

- Special Hypervisor call mechanism, a variation of System Call
- Protection of control registers managed by firmware and Hypervisor

### **I/O subsystem: Partitionable DMA address controls (TCEs)**

#### **Firmware**

- New support for over 20 Hypervisor system calls
- Partition's subset of resources passed to AIX in Open Firmware Device Tree
- Restricted platform service functions in Run-Time Abstraction Services (RTAS)

## Hardware and Firmware

LPAR was first introduced in 2001 with the release of the Regatta hardware line. These PowerPC systems had special hardware features added to support LPAR that are not available in Intel®-based systems. Each operating system image requires a range of memory that can be accessed in real addressing mode. In a non-LPAR environment, no virtual address translation is performed, and addresses start at address 0. Operating systems typically use this address range for startup kernel code, fixed kernel structures, and interrupt vectors. Since multiple partitions can not be allowed to share the same memory range at physical address 0, each partition must have its own real mode addressing range.

Special assistance is required from firmware (microcode) to support LPAR. When each partition is started, firmware assigns that partition a unique real mode address offset and range value, and then sets these offset and range values into registers in each processor in the partition. These values map to a physical memory address range that has been exclusively assigned to that partition. When partition program access instructions and data in real addressing mode, the hardware automatically adds the real mode offset value to each address before accessing physical memory. Each logical partition programming model appears to have access to physical address 0, even though addresses are being transparently redirected to another address range.

Hardware logic prevents modification of these registers by operating system code running in the partitions. Any attempt to access a real address outside the assigned range results in an addressing exception interrupt, which is handled by the operating system exception handler in the partition.


PCI I/O adapter direct memory access (DMA) operations move data between I/O adapters and system memory, and they use a similar address relocation mechanism to page tables. PCI host bridge hardware translates addresses generated by I/O devices into physical memory addresses. The I/O bridge makes this translation with a translation control entry (TCE) table, which is also stored in physical memory. As with page tables,

this TCE table resides in a physical address region of system memory that is inaccessible by partitions, and only accessible by the Hypervisor. Unlike page tables, however, TCE tables are not tied to any one partition. By calling a Hypervisor service, partition programs can create, modify, or delete TCE table entries for the specific address offset to the DMA target address provided by the partition program. When the I/O bridge translates to an I/O adapter DMA address into physical memory, the resulting address falls within the physical memory space assigned to that partition.

The Hypervisor program is trusted to create partition environments, and is the only program that can directly access special processor registers and translation table entries. Partition programs have no way to access the Hypervisor's instructions or data, other than through controlled Hypervisor service calls that are part of the processor architecture. These protections allow the Hypervisor to perform its duties in a simple and rigorous manner, resulting in the confinement of each operating system installed in each partition.

pSeries servers provide a set of architected firmware Run-Time Abstraction Services (RTAS) calls. In an LPAR environment, these calls perform additional validation checking and resource virtualization for the partitioned environment. For example, there is only one physical non-volatile RAM chip and one physical battery-powered Time-of-Day chip. RTAS makes it appear to each partition as though it has its own non-volatile RAM area and its own uniquely settable Time-of-Day clock. Because RTAS calls run inside a partition with the operating system, even they are not allowed to access anything outside the partition without a call to the Hypervisor.

The graphics on the following slides illustrate these mechanisms.



## Hardware and Firmware Mechanisms

This chart shows how the hardware and firmware mechanisms work to ensure that operating systems running in different partitions cannot interfere with each other. This example shows resources owned by two partitions, in red and blue.

Virtual addresses in the processors are translated by page tables into addresses in real memory. Only the Hypervisor is allowed to change those tables. There are several Hypervisor calls that the operating system uses to request new pages and delete old ones.

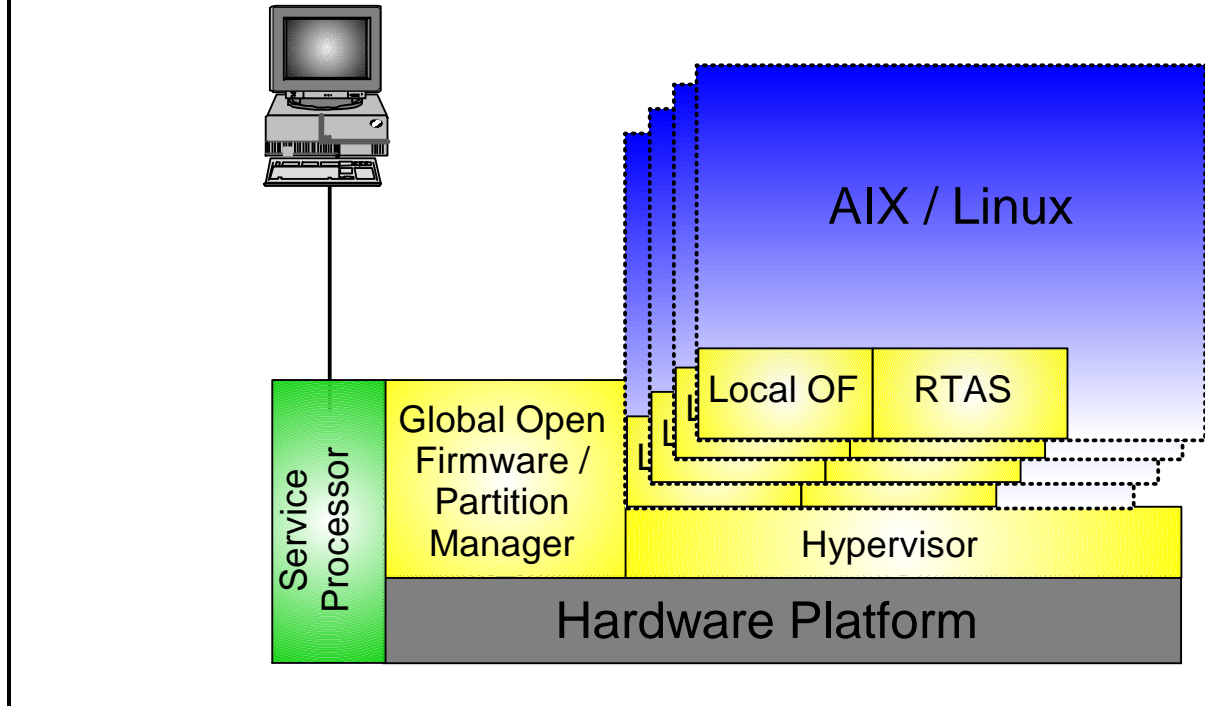
Each partition also needs some amount of real address memory, starting at address zero. Since the partitions cannot really have access to address 0, there is a special register in the processor that gets automatically added to all real address requests. This way each operating system thinks it has an address 0, but in fact that address range can be redirected anywhere in physical memory.

For I/O devices that do DMA, there is an equivalent translation method called Translation Control Entries (TCEs) that are also controlled by the Hypervisor. This is controlled at an I/O adapter slot level. When a device tries to transfer data into or out of memory, this mechanism ensures that it can only access memory that is valid for that partition.

Lastly, the processors have to read and write registers in the I/O devices to set up I/O operations. Access to those address spaces is also controlled through the virtual-to-real translation in the page tables.

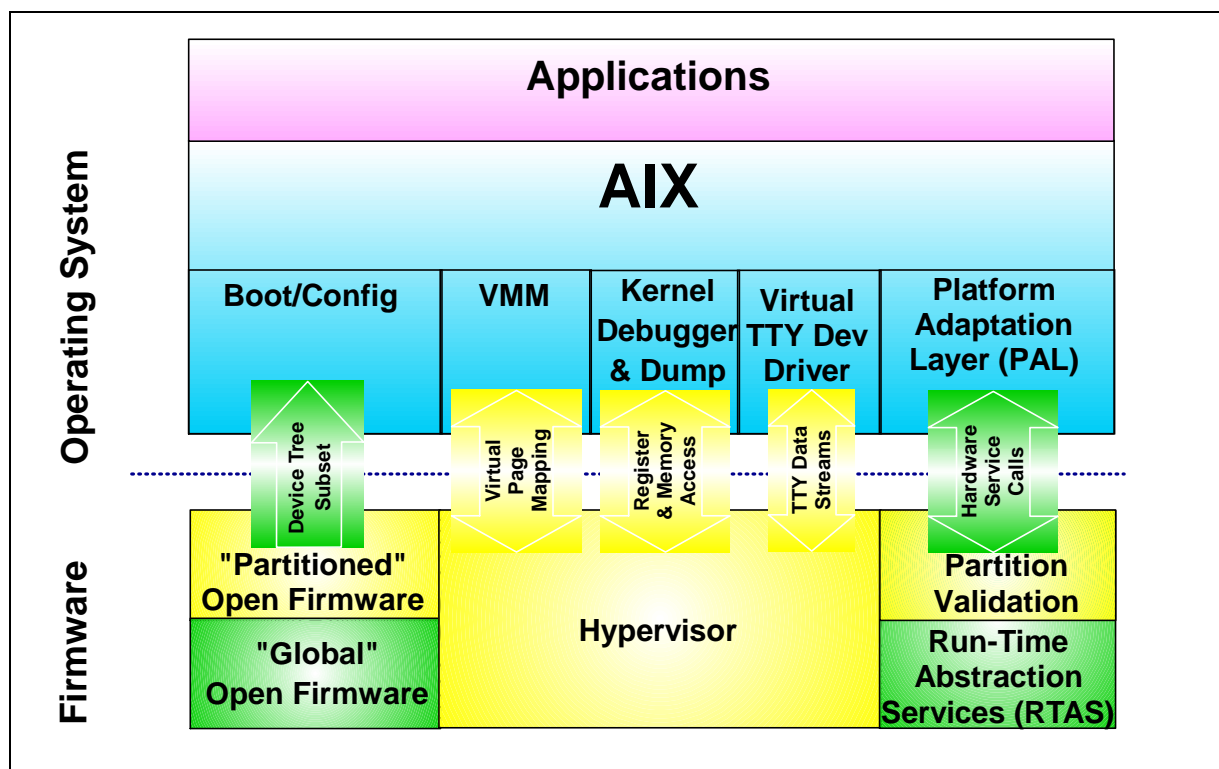


## Hardware Management Console (HMC)



## Hardware Management Console

LPAR requires an interface in AIX called the Hardware Management Console, which is a separate console used to configure and monitor an LPAR environment. This console hooks into the service processor and allows you to work with and manage the LPAR configuration. We will cover the HMC in much greater detail later in this course.




## Applications in an LPAR Environment

This chart provides more of the software view of how LPAR is seen by the AIX operating system. At boot time, AIX receives a list of information from boot firmware that describes the system resources. This list of information is called an Open Firmware Device Tree. In an LPAR environment, AIX only receives information about the subset of system resources that are assigned to that partition.

The Hypervisor provides three main types of services calls to AIX. One controls all of the virtual memory management, which maps virtual and DMA addresses into physical memory. Another set provides special controlled access to certain registers for debug purposes. The third provides a set of services for virtual TTY consoles that are presented and consolidated at the hardware system console. The virtual TTY consoles are used for two main purposes: 1) To give the shell a place to input and output. Each shell opens a new TTY from which it reads command lines, all output from the shell goes to the same TTY; 2) To provide each virtual server with virtual display hardware and a virtual keyboard. All virtual servers have a TTY allocated for them when they are created.

AIX also uses callbacks to firmware (RTAS) to perform certain system functions. The interfaces to these are not changed by LPAR, but in some cases, they must protect or virtualize certain things in an LPAR environment. A key example is the system Time-Of-Day (TOD) clock, which AIX can get and set through RTAS. Since there is only really one battery-powered clock in the system, RTAS creates a virtual TOD clock for each partition, based on an offset from the physical TOD clock value.

## Logical Partitioning Details



### Partitions

- Hardware can support up to 32 concurrent partitions
  - Limited by number of installed resources (processors, memory, I/O)
- Partitions are associated with an installed image of AIX
- Any resource can be allocated to any partition
  - At any given time, a resource can only be assigned to one partition, or none at all

### Configuration Granularities

- Processors: from 1 to total number installed, in increments of 1
  - No time-sliced processor sharing between partitions
- Memory: 1GB minimum, increments of 256MB
- I/O: Allocation at the PCI adapter slot level
  - No suballocation of devices beyond the PCI adapters
  - Some integrated devices and associated slots must be allocated as 1 unit
  - No I/O sharing between partitions (except standard multi-host I/O sharing)

## Logical Partitioning Details

A single system can support up to 32 concurrent partitions. The number of actual partitions that can be supported on a particular system depends on the number of installed resources such as processors, memory, and I/O. Each partition contains its own installed image of AIX; the operating system code is not shared across partitions.

When configuring an LPAR, resources can be allocated to a single partition only. They cannot be shared across partitions. The level of granularity of assigning processors on a pSeries server is limited to a single processor. You cannot assign one-fourth of a processor to a particular partition, for example. A key reason for this is that much of the I/O subsystem is handled by real-time code running in AIX. Time-slicing would not allow predictable real-time responses to I/O interrupts.

The minimum amount of memory that can be assigned to a single partition is one gigabyte. Additional memory can be added in 256 megabytes increments. I/O is allocated at the PCI adapter slot level. This means there is no I/O sharing between partitions, except in standard multi-host I/O sharing environments. Additionally, no suballocation of devices is allowed beyond the PCI adapters. Some integrated devices and associated slots must be allocated as one unit.

## Logical Partitioning Details

### LPAR is invisible to AIX applications

- To applications, a partition appears to be a "small" system
  - May not have any standard I/O devices - Serial, Parallel, Diskette
- Applications do not make direct Hypervisor calls
  - Kernel extensions can, but still cannot affect anything outside the partition
- Some constraint of hardware platform functions in partitions
  - For example, system firmware updates

### LPAR does affect performance (slightly)

- Extra overhead in Hypervisor for virtual memory management
  - Minor for most workloads
- Impact depends upon the amount of page mapping activity
  - Sensitive to amount of allocated memory and type of workload
- Partitioning may actually help performance in some cases
  - For applications that do not scale well on large SMPs by enforcing strong separation between workloads

## Logical Partitioning Details

Applications that are deployed in a LPAR environment are unaware of the fact that they are running in a partition that has been carved out of a larger physical system. To applications, an LPAR appears to be a "small" system. Because applications do not make direct Hypervisor calls, they cannot affect anything outside the partition they are running in. LPAR does not affect any applications or create any binary incompatibilities; it is totally hidden under the lowest layers of the operating system.

Some functions in LPAR will be constrained. A key example is the ability to update the system firmware from AIX. A customer with a test and production partition may not want the test users updating the system firmware, which would affect all partitions. The management console (HMC) provides policy settings for the partitions which allow one (or none) of the partitions to have this type of control.

There is a small performance hit associated with LPAR. The impact varies from workload to workload, however, benchmark tests have shown the impact to be less than 3-4% for workloads that do a lot of paging, and around 1% for most other workloads. The impact varies because the overhead in the Hypervisor for virtual memory management is more significant in environments involving high paging rates.

LPAR can actually help improve performance. In particular for applications that do not scale well on large SMP environments, implementing those same applications often see improved performance in an LPAR environment.

## **Agenda**

**Logical partitioning concepts**



**Planning for LPAR**

**Hardware management console (HMC)**

**Working with partitions**

**Dynamic logical partitioning on a pSeries server**

**Summary**

## **Planning for LPAR**

Now that you have a good understanding of LPAR terminology and logical partitioning concepts, let's focus on the various planning aspects required for implementing LPAR.

## **Dividing the Server**

### **Decide how big you want each partition to be:**

- Number of CPUs
- Amount of memory
- Number of PCI slots

### **Decide if partition needs alternate configuration**

- e.g. Month-end batch workload
- Additional partition profiles describe alternate configurations

### **Keep track of which partition configurations will be active concurrently**

**Up to 8 LPARS can be created on the p650**

## **Dividing the Server**

There are a number of decisions that need to be made prior to creating a logical partition. First, you need to decide how many resources the LPAR will require. How many CPUs do you need added to the LPAR to perform the workload you plan to run in this LPAR? Additionally, how much memory will be required to accommodate the workload, along with how many PCI slots will be necessary to ensure proper I/O performance?

How will the partition be used? Will it be used 24x7 for a production application, or will it only be used for month-end payroll processing? Will additional partition profiles be required for the LPAR configuration, or is one partition profile adequate?

How many LPARs will be required for the environment you need to configure? Can the system you are planning to partition support the number of partitions you require?

We will address each of these questions in more detail on this next section.

## CPU Resources

### Decide number of CPUs *required* for partition

### Decide number of CPUs *desired* for partition

- Partition can be started if the system has at least the required number available at activation time
- If less than required number available, then partition will not start
- Will use up to the desired number (if sufficient CPUs available)

### CPUs allocated in units of 1

- Cannot be shared between multiple active partitions

## CPU Resources

When using the HMC to configure a partition, you will see the terminology “desired,” “minimum,” and “maximum.” “Desired” equates to the number of a particular resource that you would ideally like. The “minimum” equates to the “required” number of a resource for that partition. And the “maximum” is the maximum number allowed for a particular resource.

For example, when determining the number of CPU resources to configure for a partition, the “required” number of CPUs is one. Because CPUs cannot be shared between multiple active partitions, the minimum number of CPUs required for a partition is one. The number of CPUs “desired” may be four. This means the partition can be started if only one CPU is available, however, if the sufficient number of CPUs is available, then the partition will be started with four CPUs.

If the system you are starting the partition on only has four CPUs total, then the “maximum” would equal the number of CPUs “desired”. However, if the system was an 8-way, then a maximum of eight CPUs could be allocated to the partition.

The number of CPUs required to run the workload you plan to place in the partition will vary depending on the type of work that will be executing. If you will be running a CPU intensive application, then you need to ensure the appropriate amount of CPU is defined for the “minimum” or “required” configuration of this partition.

## Memory Resources

Decide amount of memory **required** for partition

Decide amount of memory **desired** for partition

- Partition can be started if the system has at least the required amount of memory available at activation time
- If less than required amount available, then partition will not start
- Will use up to the desired amount (if sufficient memory available)

**Partition minimum = 1GB, increments of 256MB**

- Hardware allocates memory in 256MB chunks

**Need to factor in LPAR memory overhead**

- e.g. A system with 64GB of memory cannot have two 32GB partitions



## Memory Resources

The same terminology applies to memory resources you plan to allocate to the partition. There will be an amount of memory “required” which is the minimum level the partition will need to perform its workload. The “desired” amount of memory is the amount that you would ideally like if it was available on the system when the LPAR is started.

The minimum amount of memory that can be allocated to an LPAR is 256 megabytes. Additional memory can be added in 256 megabyte allocations. The maximum amount of memory that can be assigned to an individual LPAR depends on how much memory is available on the physical system. If you plan to only create one LPAR on a system, then the total amount of memory available on the system can be assigned to that one LPAR.

If you plan to create two or more LPARs on a server, then you need to plan appropriately for the memory allocation. For example, if you have a server that has a total of 64 gigabytes of memory and you plan to create two LPARs, you cannot create two LPARs with 32 gigabytes of memory each. *[NOTE: The server cannot be divided into two 32-gigabyte partitions, even though the minimum is 256 megabytes, because the Hypervisor uses some memory that is not allocated to any of the partitions.]*



## LPAR Memory Overhead

### Hypervisor requires a 256MB chunk for private use

- Always allocated at physical address zero

### Each partition requires a partition page table

- An area of contiguous memory used to track the location of the partitions memory - in addition to the memory allocated to the partition
- Size of partition page table is 1/64th of the size of the partitions memory
  - Rounded up to nearest power of 2
- Aligned on same size boundary
  - e.g. 128MB page table is aligned on 128MB boundary
- Small page tables may be contained within the 256MB Hypervisor segment
- Allocate memory in powers of 2 for efficient page table useage
  - i.e. 1GB, 2GB, 4GB, 8GB, 16GB, 32GB etc.

## LPAR Memory Overhead

When configuring an LPAR, it is helpful to understand how much memory the firmware requires. The Hypervisor requires 256 megabytes memory. An additional 256 megabytes is allocated to translation control entries (TCEs) and to Hypervisor per partition page tables. Another 256 megabytes is used for the first page table for the first partition. TCE memory is used to translate the I/O addresses to system memory addresses. Additional small page tables for additional small partitions will fit in the page table block. Therefore, the memory allocated independently of AIX to create a single 256 megabytes partition is 768 megabytes of memory.

As you will recall from earlier in this course, each partition has its own exclusive page table. Processors use these tables to transparently convert a program's virtual address into the physical address where that page has been mapped into physical memory. The size of the partition page table is 1/64<sup>th</sup> of the size of the memory that has been allocated to the partition. Small page tables may be contained in the 256 megabytes Hypervisor memory segment; however, it is common for the page table to reside in a separate 256 megabytes memory allocation as stated above.

For best performance, it is advisable to allocate memory in powers of two for efficient page table usage. For example, use memory allocations of one gigabyte, two gigabytes, four gigabytes, and so on.

## Memory Resources

### Different memory considerations when using AIX 5.1, AIX 5.2, Linux

- AIX 5.1
  - number of LPARs larger than 16GB is limited to 2 in a system
  - 64 GB installed memory
    - this is due to memory alignment requirement,
- AIX 5.2 does not have this restriction
- To create LPARs running AIX 5L Version 5.2 or Linux larger than 16 GB
  - Checkbox *Small Real Mode Address Region* must be checked (on the HMC, LPAR Profile, Memory Options dialog)
  - Do not select this box if you are running AIX 5L Version 5.1

## Memory Resources

There are key differences in how memory is allocated between AIX version 5.1 and 5.2. In AIX 5.1, no more than two LPARs can be created with more than 16 gigabytes of memory. This restriction has been removed in AIX 5.2.

The maximum amount of memory that can be installed in AIX 5.1 is 64 gigabytes. The maximum amount of memory that can be installed in AIX 5.2 is 256 gigabytes. However, the hardware model typically has a maximum memory limit which is not really limited by the OS level.

When configuring an LPAR, you have two options for selecting LPAR real mode memory addressing region size: large and small. If you select small real mode memory addressing, system memory will be managed more efficiently, avoiding some of the memory-allocation constraints associated with large real mode memory addressing. To use small real mode memory addressing, you must have either AIX 5.2 or later or Linux installed in the partition.

If you select the large real mode addressing, each large real mode memory address range is associated with a scalable real mode address region of a particular size (256 megabytes, one gigabyte, and 16 gigabytes). The real mode address region size is determined by the maximum partition memory size that you specify. If large real-mode memory addressing regions are used, an LPAR logical memory range from one gigabyte to 16 gigabytes will use a real-mode addressing memory region size of 16 gigabytes. Anything over 16 gigabytes will also use a real-mode addressing memory region size of 16 gigabytes.

## Partition Page Table Size

Partition Memory Size	Partition Page Table	Alignment	256MB Segments Used
1GB	16MB	16MB	1
1.25GB - 2GB	32MB	32MB	1
2.25GB - 4GB	64MB	64MB	1
4.25GB - 8GB	128MB	128MB	1
8.25GB - 16GB	256MB	256MB	1
16.25GB - 32GB	512MB	512MB	2
32.25GB - 64GB	1GB	1GB	4
64.25GB - 128GB	2GB	2GB	8
128.25GB +	4GB	4GB	16

## Partition Page Table Size

The table on this slide shows the partition page table size, the alignment, and the number of 256 megabytes segments required for different partition memory sizes. Let's look at a couple of examples over the next few pages.

Example 1: A partition memory size of seven gigabytes would require a partition page table size of 128 megabytes, an alignment of 128 megabytes, and would only require one 256-megabyte segment. Compare that to a partition memory size of 90 megabytes. This would require two gigabytes for the partition page table and eight 256-megabyte segments, and the alignment would be on a one-gigabyte boundary.

## TCE Table Size

Number of Drawers	TCE Table size	Number of 256MB segments unavailable for partition memory
1 B&C Drawer	64MB	1
2 B&C Drawers	128MB	1
3 B&C Drawers	192MB	1
4 B&C Drawers	256MB	1
5 B&C Drawers	320MB	2
6 B&C Drawers	384MB	2

## TCE Table Size

The translation control entry table was initially discussed in the planning section of this course. As you will recall, the PCI I/O adapter direct memory access (DMA) operations move data between I/O adapters and system memory. The PCI host bridge hardware translates addresses generated by I/O devices into physical memory addresses. The I/O bridge makes this translation with a TCE table which is stored in physical memory.

The size of the TCE table varies by the number of I/O drawers configured in the partition. For example, a partition that is configured with four drawers requires a TCE table size of 256 megabytes and only one 256 megabytes segment of memory.

## **I/O Resources**

**A partition requires a minimum of:**

**One boot device on assigned I/O slot**

- To load AIX operating system

**One network adapter in assigned I/O slot**

- For diagnostics and monitoring - not exclusive use like SP ethernet

**Other I/O adapters for storage and network access**

**I/O slots can be *required* or *desired***

## **I/O Resources**

As with CPU and memory resources, there are requirements for the amount of “required” and “desired” I/O resources for the configuration of an LPAR. The minimum I/O resources required for a partition include one boot device and one network adapter. The boot device is required to load the AIX operating system and the network adapter is needed for diagnostics and monitoring.

Additional I/O adapters will most likely be required for additional storage capacity and network access. The amount of additional capacity will depend on the amount of workload that will be running in the LPAR. Going back to the initial set of questions we looked at for defining the requirements of an LPAR, it is important to address the type of work as a month-end batch process will have different storage and network requirements than an enterprise application that provides critical services for your company.

## **I/O Resources**

### **I/O devices are assigned at slot level to LPARs**

- an adapter in a given slot can only be assigned to one LPAR

### **In the p650 — internal disks, media bays, & external SCSI port are driven by one SCSI chip on I/O backplane**

- This chip is connected to one of the PCI-X-to-PCI-X bridges, which is equivalent to one PCI slot in LPAR terms.

### **Use an external attached DVD-RAM to have access to CD-ROMs & DVD-RAMS from different LPARs. You can use:**

- External attached DVD-RAM ( FC 7210 Model 025)
- Storage Device Enclosure ( FC 7212 Model 102)
- PCI SCSI adapter ( FC 6203)

## **I/O Resources**

When assigning I/O devices to a partition, you will need to know the PCI slots for network adapters and other devices. I/O adapters are physically installed in one of many drawers in a pSeries server. However, with logical partitioning, any I/O adapter in any I/O drawer can be assigned to any partition. Remember that once an adapter in a particular slot has been assigned to one LPAR, it cannot be allocated to another partition.

If you are configuring LPAR on a p650 server, it is important to note that the internal disks, media bays, and the external SCSI port are all driven by one SCSI chip on the I/O backplane. This is different from other pSeries server configurations.

When configuring an LPAR, you need to determine if each LPAR will have its own CD-ROM device, or if they will share a common device. If you would like to use an externally attached DVD-RAM device so multiple LPARs can share the same CD-ROM device, you can use one of the three models shown on this chart.

## I/O Resources

### Every LPAR needs its disk for the operating system

- p650 server has four disks in a 4-pack connected to the internal SCSI port, for additional LPARs external disk space is necessary, the following can be used
  - **IBM 2104 Expandable storage plus**
    - Can support two LPARs
    - [www.storage.ibm.com/hardsoft/products/expplus/expplus.htm](http://www.storage.ibm.com/hardsoft/products/expplus/expplus.htm)
  - **IBM 7133 Serial Disk Subsystem ( SSA)**
    - Can support up to four LPARs using an Advanced Raid plus adapter
    - Each 7133 provides a max of four loops
  - **(FC 6230) in each LPAR**
    - Disk space for booting can be in RAID or JBOD (just a bunch of disks) format
    - [www.storage.ibm.com/hardsoft/products/7133.htm](http://www.storage.ibm.com/hardsoft/products/7133.htm)
- With firmware level 7000 and above the 6230 can support booting from RAID configured disks
  - Fast write cache must not be enabled on the boot resource SSA adapter
    - [www.storage.ibm.com/hardsoft/ssa/faq.html#microcode](http://www.storage.ibm.com/hardsoft/ssa/faq.html#microcode)

## I/O Resources

When planning I/O resources for a partition, keep in mind that each LPAR needs disk storage space for the operating system, as each LPAR has its own version of the operating system installed. Again there are some special requirements for the p650 server.

The p650 server has four disks in a four-pack configuration connected to the internal SCSI port. If the amount of disk required for your LPAR exceeds this configuration, you will need to use external disk space. These external disk units can be any of the following configurations:

- The “IBM 2104 expandable storage plus”
- The “IBM 7133 serial disk subsystem” (SSA)
- Or, an FC 6230 in each LPAR

If you have chosen the 2104 expandable storage plus, up to two LPARs can be supported. The 7133 serial disk subsystem can support up to four LPARs using the advanced RAID plus adapter. And finally, FC 6230 could be used in each LPAR, providing disk space in either RAID or just a bunch of disks (JBOD) format.

If you have chosen the FC 6230, you can boot from RAID configured disks if the firmware level is level 7000 or above. Notes that fast write cache should not be enabled on the boot resource SSA adapter if choosing the boot from the RAID option.

## **I/O Resources**

### **IBM TotalStorage FastT Storage Server**

- Model 700 supports upto 64 LPARs
  - [www.storage.ibm.com/disk/fastt/](http://www.storage.ibm.com/disk/fastt/)

### **IBM TotalStorage Enterprise Storage Server ( ESS)**

- [www.storage.ibm.com/disk/ess/](http://www.storage.ibm.com/disk/ess/)

## **I/O Resources**

Additional disk storage solutions include the IBM TotalStorage® FastT Storage Server. Up to 64 LPARs can be supported with the model 700 TotalStorage FastT Storage Server.

IBM also offers its TotalStorage Enterprise Storage Server (ESS) family of products. This family of products supports fast data transfer rates through a wide variety of host attachment applications, making the ESS a perfect solution for managing a wide variety of heterogeneous operating environments.

For more details on either of these servers, refer to the Web sites found on the Reference page of this course.



## **IP Addresses for LPAR Mode**

### **Need at least one IP address per partition**

- For network interface used for diagnostics

### **Single physical machine (CEC) now has multiple hostnames**

- Each OS instance has its own hostname

### **CEC name shown in HSC application**

- No IP address for CEC in LPAR mode

## **IP Addresses for LPAR Mode**

When planning for your LPAR implementation, you also need to determine the number of TCP/IP addresses that will be required for the configuration. A minimum of one IP address is required for each partition. This IP interface is used for basic communication with the partition, including diagnostics.

Each partition must have a unique hostname that can be resolved. Hostnames cannot be reused between the physical server (CEC) and the logical partitions.

When LPAR'ing a server, the single physical machine, also called the CEC, can have multiple hostnames. Each operating system instance in an LPAR will have its own hostname. This means there is no TCP/IP address required for the CEC when in LPAR mode, because the CEC of the physical machine can contain multiple hostnames.

## **Agenda**

**Logical partitioning concepts**

**Planning for LPAR**



**Hardware management console (HMC)**

**Working with partitions**

**Dynamic logical partitioning on a pSeries server**

**Summary**

## **Hardware Management Console**

The next section of this course will show you how to work with the HMC. The HMC is required to configure and work with LPARs on a pSeries server.

## **Hardware Management Console**

### **Control point for LPAR servers**

- Define, set up, start/stop, monitor status

### **Operates outside of any AIX image**

- Need to create initial partitioning setup before installing AIX
- Avoid need for a "master" or "special" partition
- Provide security model independent of any AIX image
- Make management functions available, even when the server is not operational or powered off

### **Provides virtual devices for partitions**

- For AIX devices with only one physical instance per server
- Virtual serial port / TTY console, and Operator Panel, per partition

### **Remote, LAN-accessible management**

## **Hardware Management Console**

The Hardware Management Console is crucial to the pSeries LPAR environment. The HMC is required for LPAR'ing any pSeries server. It is also needed for the p655, p670, and p690, even if these servers are not LPAR'd. HMC is used for creating partition profiles that define the processor, memory, and I/O resources allocated to an individual partition; and it is also integral to starting, stopping, and shutting down a partition.

The HMC allows you to manage the system as a whole or as individual LPARs, depending on what your requirements are. The HMC is a separate device you install and configure separately from a pSeries server. It provides an interface into one or more physical pSeries servers, for managing those environments. This means you can perform management functions on a system that is not operational or one even one that is powered off.

Each LPAR that is configured will have a virtual serial port (or TTY console) and an operator panel. This is all available through the HMC. Because the HMC is a separately installed device, it allows remote management of your pSeries server environment.

## Hardware Management Console (continued)

### Partition Configuration and Control

- LPAR definition
- Partition activation and reset

### Diagnostics


### Operational Management

- Modem connected to HMC for call-home

### Console access

- Virtual tty console per partition

**Remote HMC control from another HMC workstation on the network or Web-based System Manager Remote Client for Windows or Linux**



## Hardware Management Console (continued)

There are a number of different functions the HMC can perform. The main usage of the HMC you will see in this course is for configuring and working with a partition. However, the HMC can also be used for problem determination and remote console access.

For remote access from one HMC to another, you can use the WebSM client component. This runs on either an AIX-based or PC-based client server. Alternatively, you can use the command line interface (CLI) to issue commands to the LPAR. We will focus on using the WebSM client in this course.

## Hardware Management Console (continued)

### HMC used for management only

- Server will continue to work if HMC fails


### Partition information stored in system NVRAM and on disk on the HMC

- New HMC can read configuration information from NVRAM

### HMC provides partition information backup facility to DVD-RAM or diskette

### HMC operates on transactional basis

### 6M or 15M cable to CEC



## Hardware Management Console (continued)

One of the benefits of the HMC architecture is that it is only used for managing the pSeries environment. It is not required for system operation. If the HMC connected to a server fails, the server and all of its LPARs will continue to work.

LPAR information is stored in the server's nonvolatile memory (NVRAM) and on disk on the HMC, providing a duplicate of system and LPAR configuration information. The newer HMCs can read the configuration information stored in the NVRAM, making it much easier to replace a failing HMC. Additionally, HMC allows partition information to be backed up to CD-ROM or diskette.

As you will discover in the next several slides, the HMC operates on a transaction basis. There are a number of different types of transactions that can be performed by the HMC. We will be examining most of the functions of the HMC in this course.

Depending on which cable you have purchased, the HMC must be located within six or 15 meters from the system CEC.

## Hardware Management Console

### HMC has two serial ports

- One should be used to attach to a modem for the Service Agent, the other one to the server
- For more servers to be managed from the same HMC
  - Need one of the following adapters to have more serial ports on the HMC
    - 8-port Async Adapter
    - 128-port Async Adapter

## Hardware Management Console

The HMC has two serial ports. One port needs to be attached directly to the server's CEC and the other one should be attached to the modem. The attachment to the modem is required for the Service Agent function within the HMC. We will be covering the different functions provided by the HMC in this section of this course.

If you plan to manage multiple pSeries servers from the same HMC, you will need a special adapter installed. There is an 8-port async adapter and a 128-port async adapter. The number of servers you need to connect into from a given HMC will dictate which adapter is appropriate for your operating environment.


# Remote HMC Support

**From another HMC server on the network**

**From any PC on the network running WebSM client application**

**Telnet to HMC and get limited access shell**

- Most HMC functionality available
  - Cannot create new partition profile - specifying I/O is too complex for command line



## Remote HMC Support

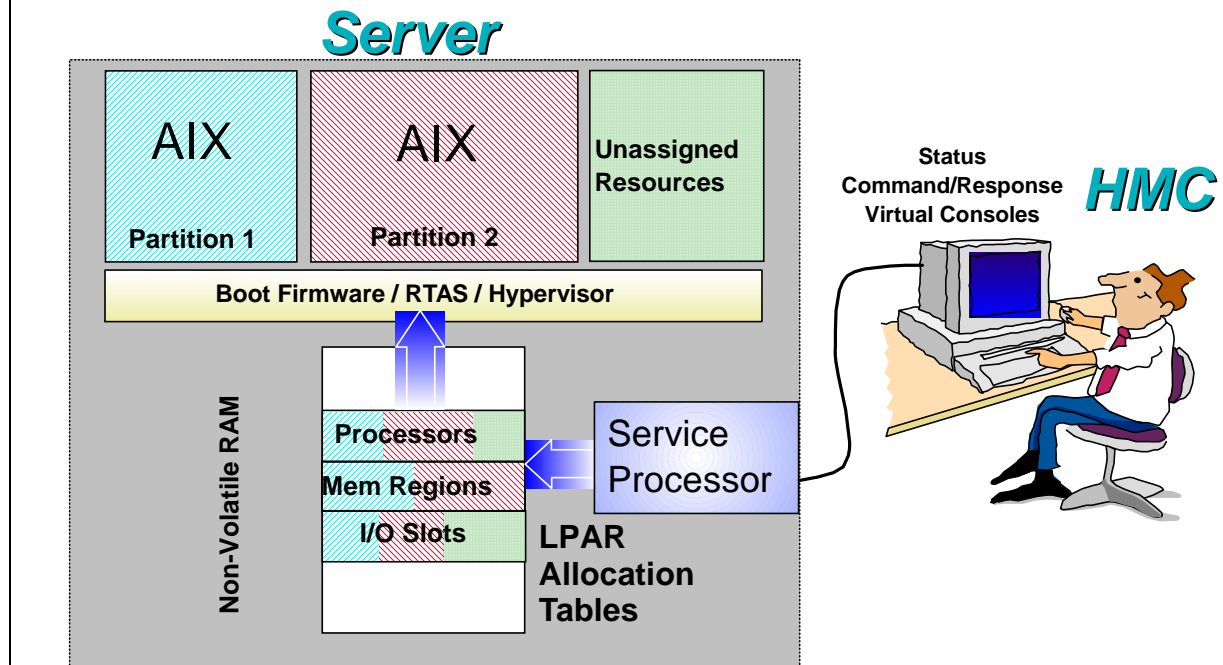
You have a variety of options available when configuring remote HMC support. The connectivity can be from another HMC on the same network; from a PC configured with the WebSM client; or by telnetting to the HMC. The most common remote HMC access is through the WebSM client application.

When using remote HMC support, you are not allowed to do the following:

- Configure System Manager Security for the certificate authority
- View overview and status information
- Determine the level of HMC code installed, or
- Reboot the HMC interface

Additionally, if you are using the option to telnet to HMC, you cannot create a new partition profile. The reason for this specific limitation is due to the complexity of specifying I/O resources for a partition. This function is only available in a HMC client that offers a graphical interface, such as the WebSM client, or accessing through another HMC on the same network.

## LPAR Setup with HMC



## LPAR Setup with HMC

The diagram on this slide shows you how the HMC is used to create a new partition. The HMC is connected to the service processor of the server being LPAR'd. Using the HMC, you define how many processors, how much memory, and which I/O resources will be allocated to the new LPAR.

Once the partition has been defined, you use the HMC to activate and install the partition. This step involves installing a specific version of AIX for the LPAR and booting the LPAR once it is installed.



# **Hardware Management Console**

**HMC workstation has multiple user logins**

**Six different types of HMC users**

- Viewer
- Operator
- Advanced Operator
- Service Representative
- System Programmer
- User Administrator

**Users can perform different tasks on the HMC**

## **Hardware Management Console**

HMC has built-in security in the form of user types. There are six types of HMC users:

- Viewer
- Operator
- Advanced Operator
- Service Representative
- System Programmer
- User Administrator

Depending on which type of user has logged onto the HMC, they can perform different sets of tasks. We will examine each of these HMC user types on the next slide.

# Hardware Management Console

## Types of users

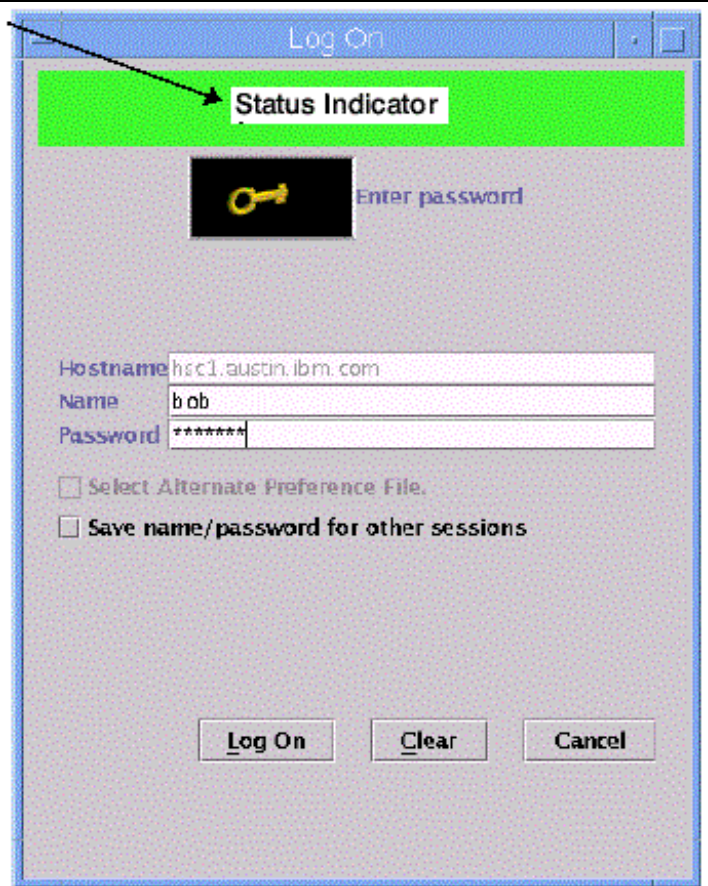
- Viewer
  - Can only view HMC information
- Operator
  - User who is responsible for daily routines
- Advanced Operator
  - User who can do some LPAR or system configuration, but does not have any user management authority
- User Administrator
  - Can perform user management tasks but cannot perform any other HMC functions
- Service Representative
  - IBM employee that has come on-site to install or repair the server
- System Administrator
  - Superuser or manager of the HMC server; has unrestricted authority to access and modify most of the HMC server

## Hardware Management Console

Each HMC user can be a member of one of the six different HMC user roles. Each of these roles allows the user to access different parts of the HMC. Each role is described as follows:

- Viewer — Can view HMC information, but cannot change any configuration information
- Operator — Is responsible for daily system operation
- Advanced Operator — Can perform some partition or system configuration and has access to some user-management functions
- User Administrator — Can perform user-management tasks, but cannot performance any other HMC functions
- Service Representative — Is a person who is at your location to install or repair the system
- System Administrator — Acts as the root user, or manager of the HMC system. The System Administrator has unrestricted authority to access and modify most of the HMC system.

# HMC Login



Log On

Status Indicator

Enter password

Hostname: hsc1.austin.ibm.com

Name: bob

Password: \*\*\*\*\*

☐ Select Alternate Preference File.

☐ Save name/password for other sessions

Log On Clear Cancel

## HMC Login

The HMC is shipped with a predefined user ID and password. Both the user ID and password are case-sensitive. This default user ID and password are as follows:

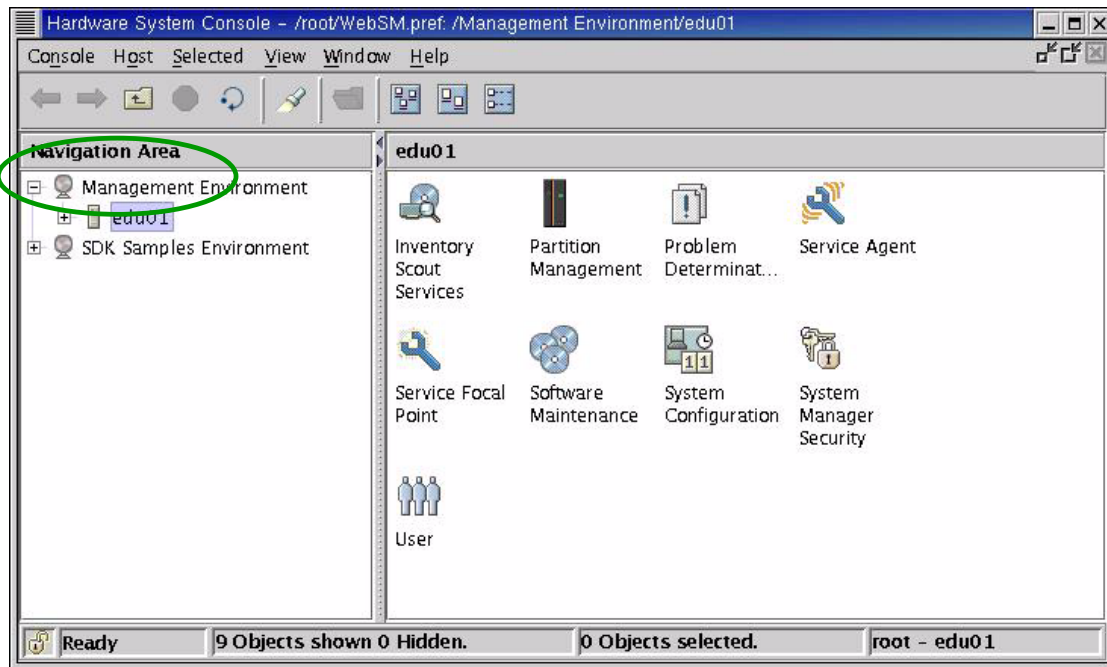
- ID: **hscroot**
- Password: **abc123**

When the console is powered on for the first time, use this user ID to log in. All subsequent logons to the HMC can be with a specific user's ID and password.

The login screen (shown on this slide) shows user **bob** logging into the HMC. User **bob** has been assigned a specific user role. Depending on which role he has been assigned, he will be able to perform only those operations allowed for the role he has been assigned.

After bob enters his user name and password, he presses the **Log On** button to log onto the HMC.

# HMC GUI



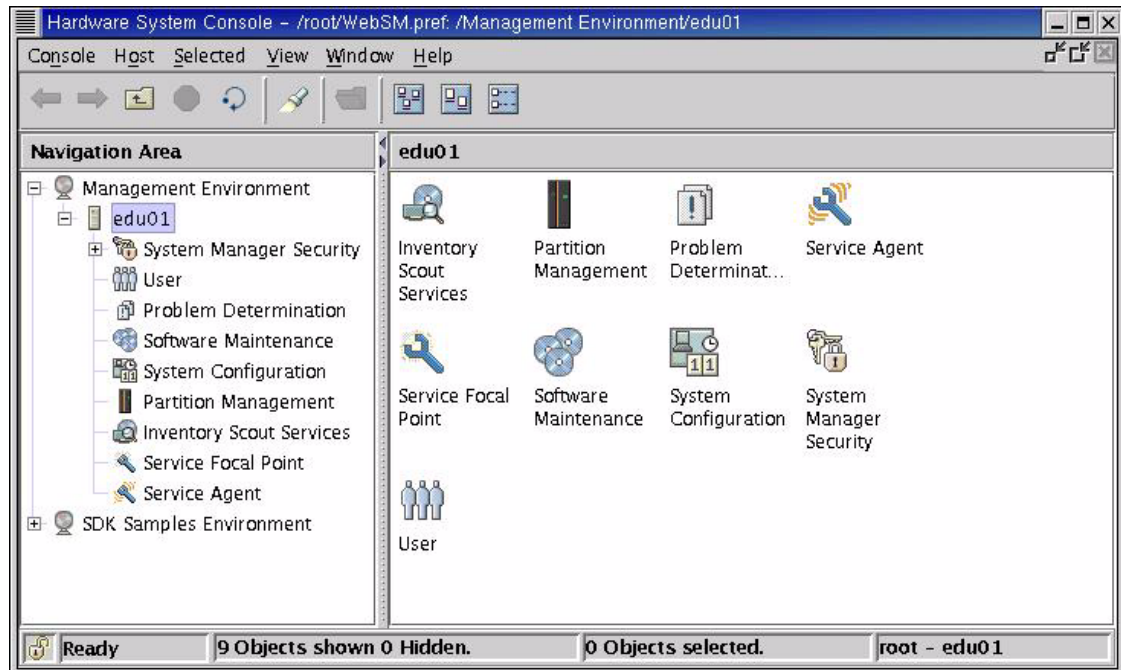
## HMC GUI

This screen shows you what the HMC GUI interface looks like. In the left pane, the Management Environment has been expanded, showing system edu01 in the navigation area. The types of functions that can be performed in the HMC against this system are:

- Inventory Scout Services
- Partition Management
- Problem Determination
- Service Agent
- Service Focal Point
- Software Maintenance
- System Configuration
- System Manager Security
- User

We will examine each of these functions in this section.

# HMC GUI

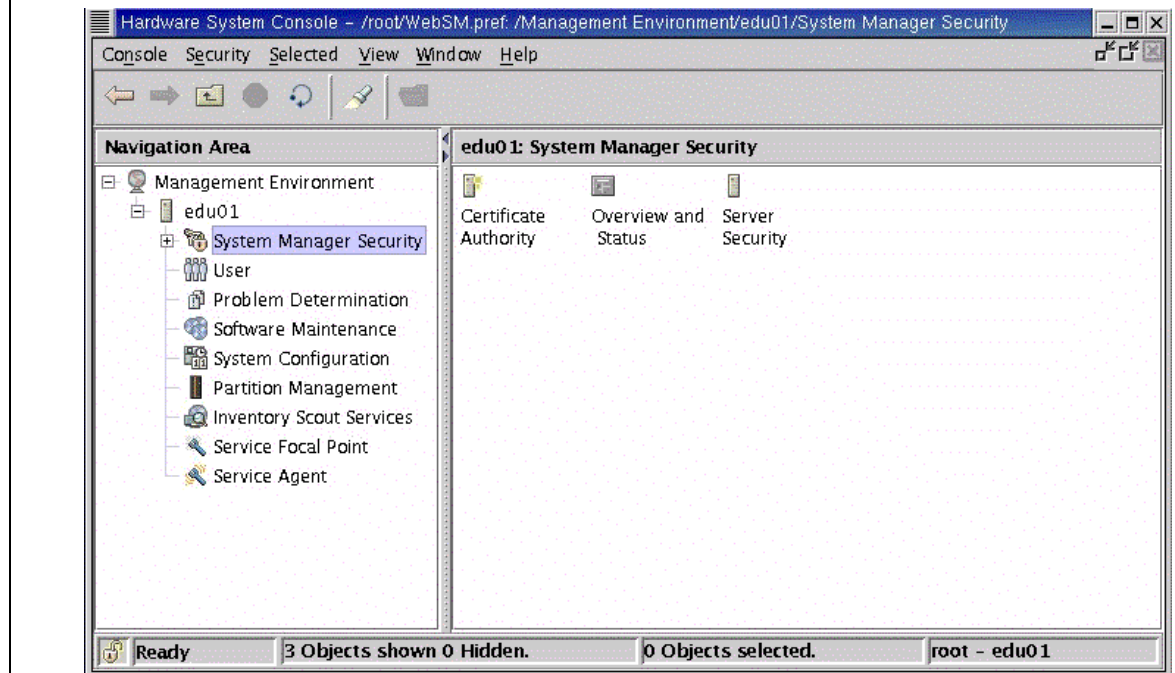


## HMC GUI

By expanding the "+" sign next to the server name, edu01, each of the functions that can be performed against the server are now listed under the server name in the navigation area in the left pane.

Let's continue to the next slide to delve into these functions one-by-one.

# HMC System Manager Security



## HMC System Manager Security

System Manager Security ensures the HMC can operate securely in client-server mode. Managed machines are servers and the managed users are clients. Servers and clients communicate over the Secure Sockets Layer (SSL) protocol, which provides server authentication, data encryption, and data integrity.

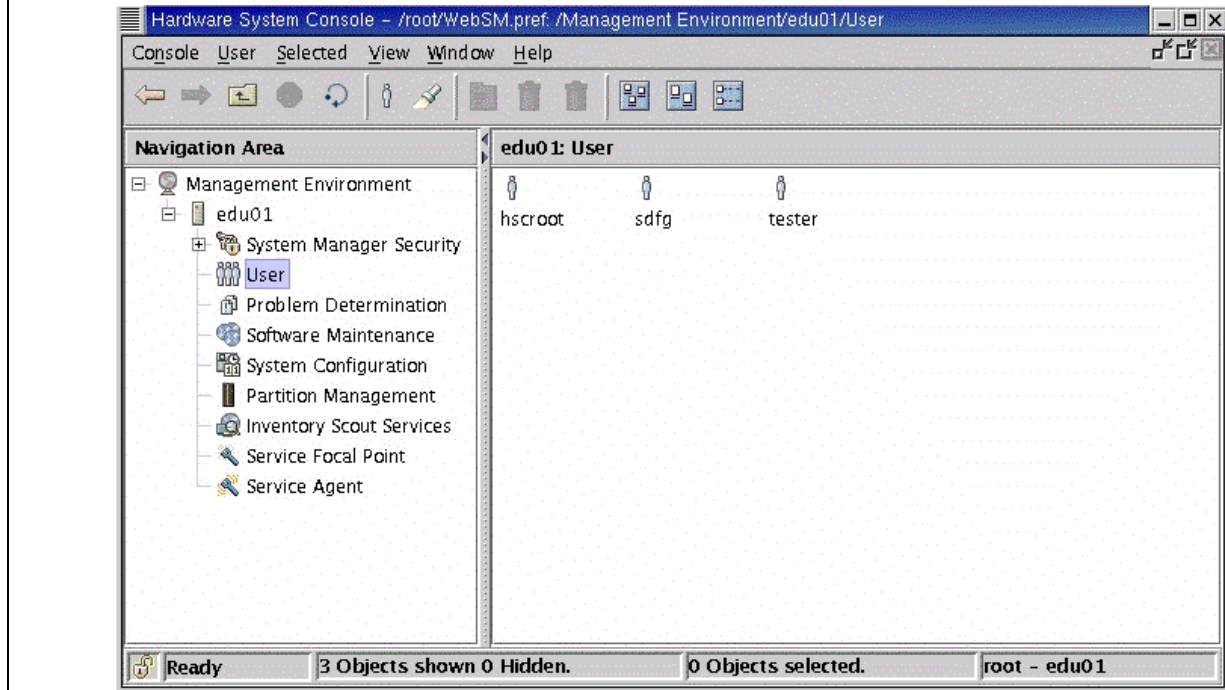
Each HMC System Manager server has its own private key and a certificate of its public key signed by a Certificate Authority (CA) that is trusted by the System Manager clients. The private key and the server certificate are stored in the server's private key ring file. Each client must have a public key ring file that contains the trusted CA's certificate.

You should define one HMC as a Certificate Authority. You then use this HMC to generate keys and certificates for your HMC servers and client systems. The servers are the HMCs you want to manage remotely. A unique key must be generated and installed on each server. You can generate the keys for all servers in one action on the CA and then copy them to diskette, install them at the servers, and configure the servers for secure operations.

The client systems are those from which you want to do remote management. Client systems can be HMCs, AIX, or PC clients. Each client system must have a copy of the CA's public key ring in its System Manager codebase directory. You can copy the CA public key ring file to the diskette on the CA and copy it from the diskette to each client.

To use the System Manager Security application, you must be a member of the System Administrator role. To ensure security during configuration, users of this application must be logged in to the HMC locally.

# HMC User Management



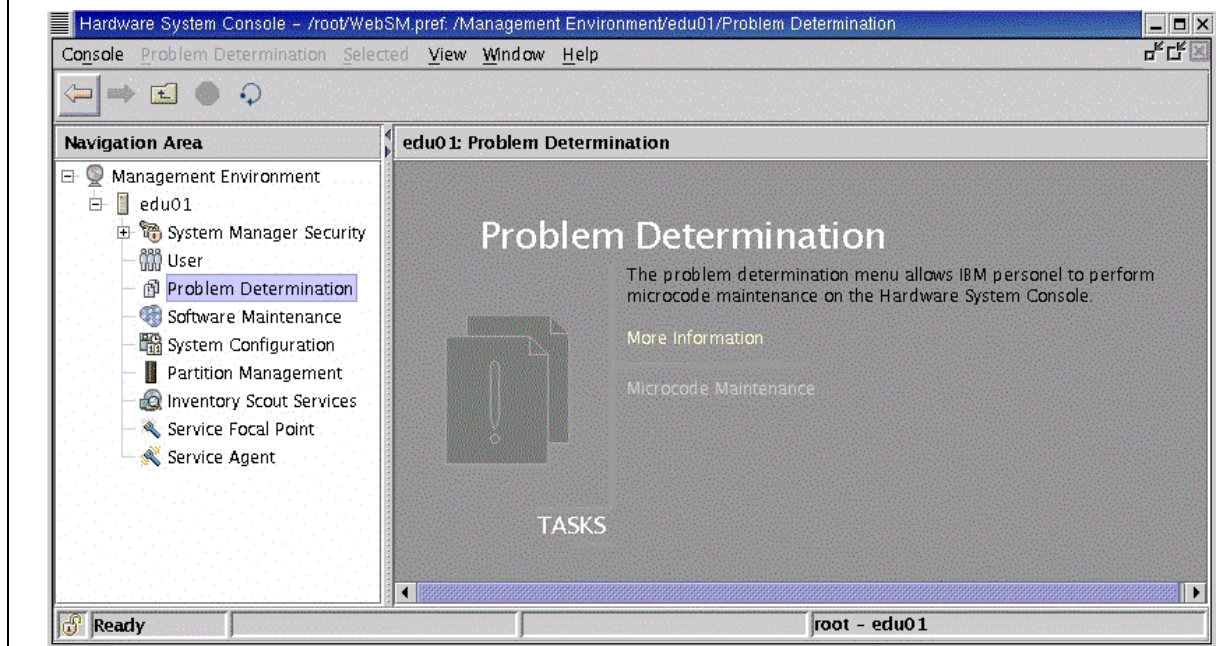
## HMC User Management

The HMC User Management role allows an HMC system administrator to manage users and assign roles. We previously discussed the six different roles that are available for users of the HMC. The User Management function within the HMC is used to create users and assign specific roles to these users.

To create users, you must be a member of one of the following roles:

- System Administrator
- User Administrator

# HMC Problem Determination

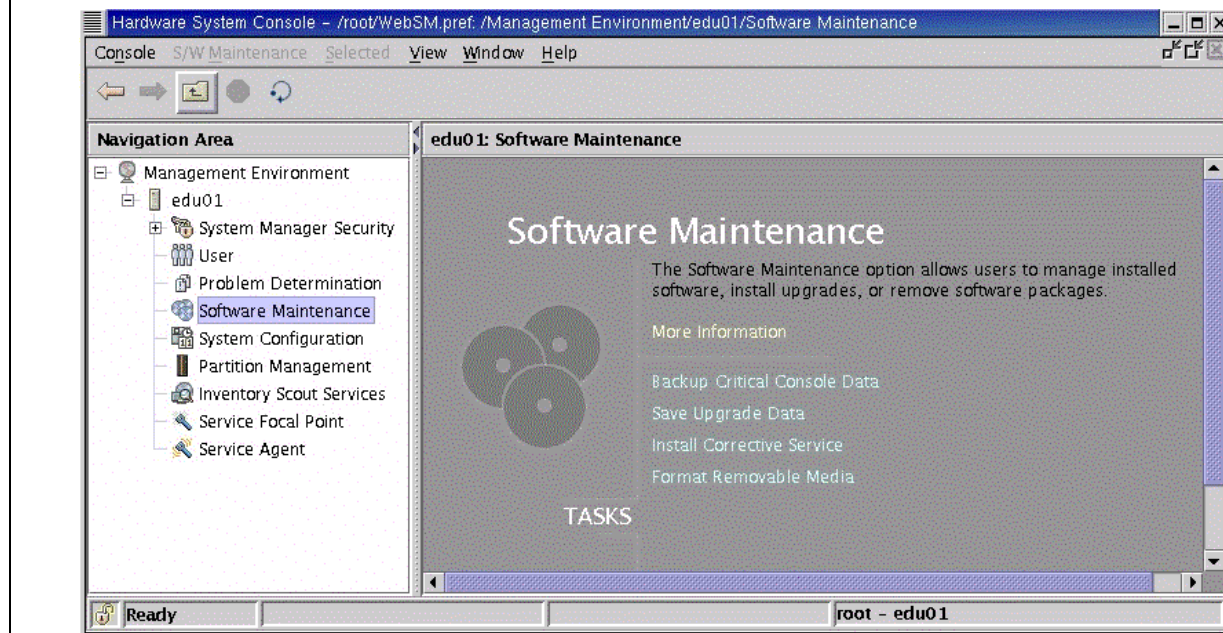


## HMC Problem Determination

The Problem Determination function is only used by an authorized service representative. It is used to view and diagnose HMC problems. We will not be providing details of the functions available through this option because it is used only by a service representative.



# HMC Software Maintenance



## HMC Software Maintenance

There are a number of applications in the Software Maintenance folder. These applications allow you to manage installed HMC software. The specific tasks you can perform include:

- Receive corrective service
- Install corrective service
- Back up critical console data
- Save upgrade data
- Format removable media
- Manage a frame of managed systems and resources

## **HMC Software Maintenance**

### **Upgrade HMC Release Levels**

- To keep current HMC configuration settings, follow instructions in the "Upgrade the HMC Software" section of the IBM Hardware Management Console for pSeries Installation and Operations Guide, SA38-0590

### **Apply Corrective Services**

- Information about corrective fixes to apply is available at:  
<http://techsupport.services.ibm.com/server/hmc>

### **Reinstall from recovery CD**

- Install/Recovery
- Upgrade

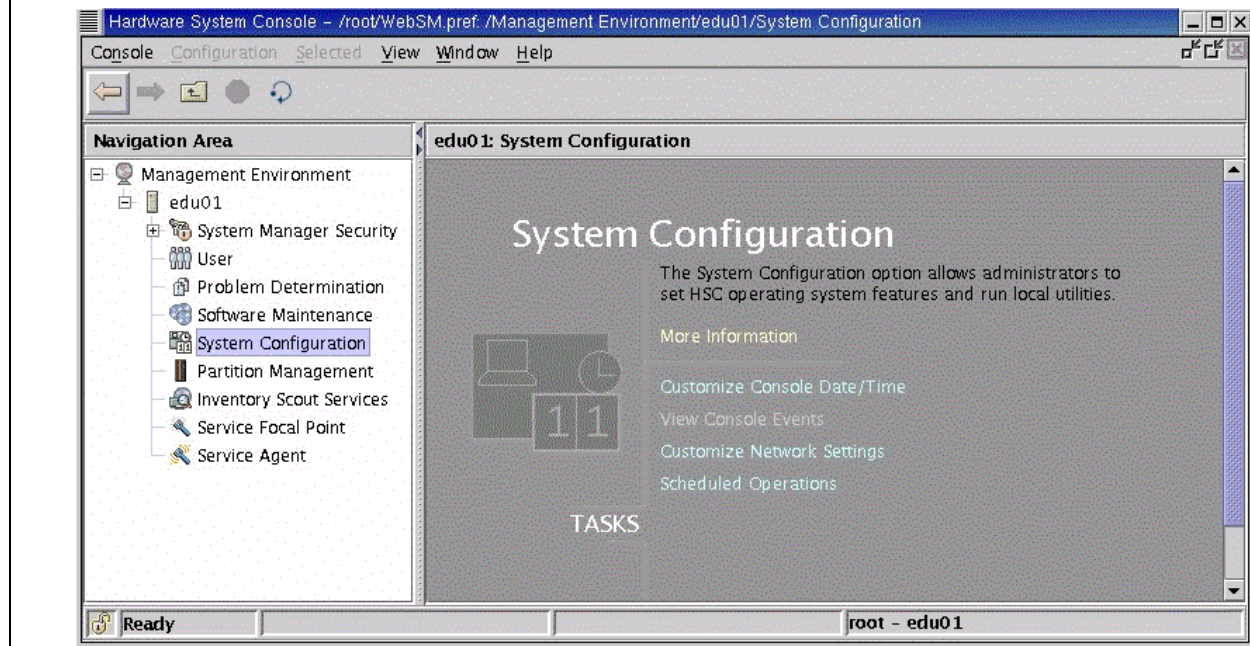
## **HMC Software Maintenance**

The HMC Software Maintenance applications can be used to keep current with HMC release levels. For detailed instructions on how to upgrade HMC release levels, see the "Upgrade the HMC Software" section of the IBM Hardware Management Console for pSeries Installation and Operations Guide, SA38-0590.

Other functions that can be performed within the Software Maintenance section of the HMC include applying fix packs and performing reinstalls. Fix packs can be downloaded from the pSeries Support Hardware Management Console Web site. (See the Hotlinks section of the course for the URL.)

HMC Software Maintenance can be used to perform reinstalls from the recovery CDs. This can either be an installation/recovery, or simply an upgrade.

# HMC System Configuration



## HMC System Configuration

The HMC System Configuration application allows you to perform the following functions:

- Set the console's date and time
- Enter and check HMC network information
- View console events
- Schedule routine backups
- Enable and disable remote commands
- Configure serial adapters
- Enable remote virtual terminal connections
- Change the HMC interface language

Any user role can view the console date and time. To update the console date and time, the user must be a member of one of the following roles: Advanced Operator, System Administrator, Service Representative, Operator, or Viewer.

Let's continue to the next slide to look at some of the functions performance by the System Configuration application in more detail.

## **HMC System Configuration**

### **View Console Events**

### **Customize Network Settings**

### **Scheduled Operations**

- Schedule the backup of critical console information

### **Verify data backed up during a scheduled back-up**

### **Enable or Disable Remote Command Execution**

- Recommendation: do not enable rexec because it transports a non-secure copy of the password across the network
- Free ssh client for Windows

## **HMC System Configuration**

To see a log of recent HMC activity, you can view console events. Each event has an associated time stamp. Below is an example of some of the events recorded in the log:

- When a server was powered on
- When a partition was activated
- When a partition was shut down
- The results of a scheduled operation

To view console entries, a user must be a member of one of the following roles: System Administrator, Advanced Operator, Service Representative, or Viewer.

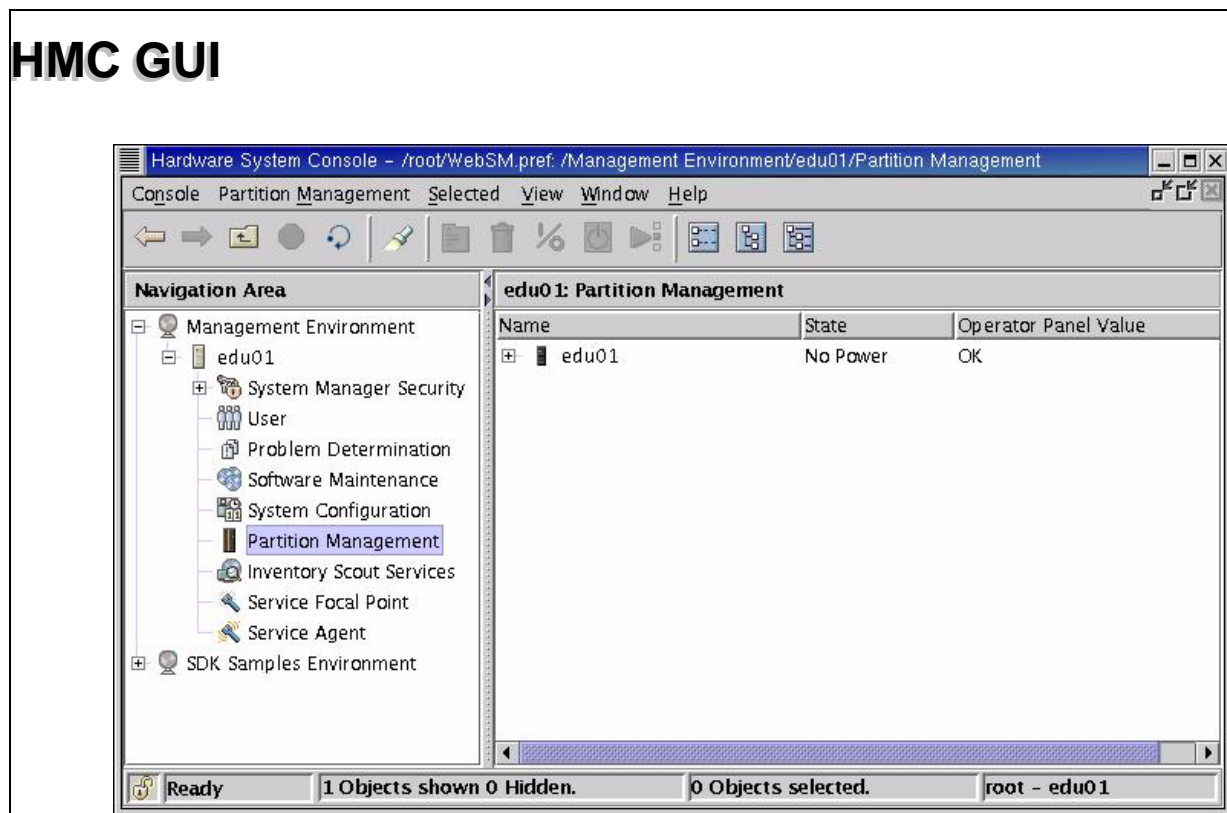
You can also customize the network settings through the System Configuration application. If you enable the HMC network connection, take advantage of some of the HMC capabilities. For example, you can access the HMC user interface remotely from a WebSM graphical user interface client; and you can configure the network to allow network connections between the HMC and the partitions in the managed server.

Using the System Configuration application, you can also verify data backed up during a scheduled backup operation. The data that is backed up is critical console information. When you schedule a backup operation, the information is saved on a CD-ROM on the HMC. To schedule a backup of console data, the user must be a member of either the Advanced Operator or System Administrator role.

Remote commands can be enabled or disabled through the System Configuration application as well. This option provides the ability to run commands remotely through the **ssh** command. To install **ssh** software on your PC, go to one of the URLs listed in the “Hotlinks” section of this course.

To enable or disable remote commands, the user must be a member of one of the following roles: Advanced Operator, System Administrator, or Service Representative.

# HMC GUI

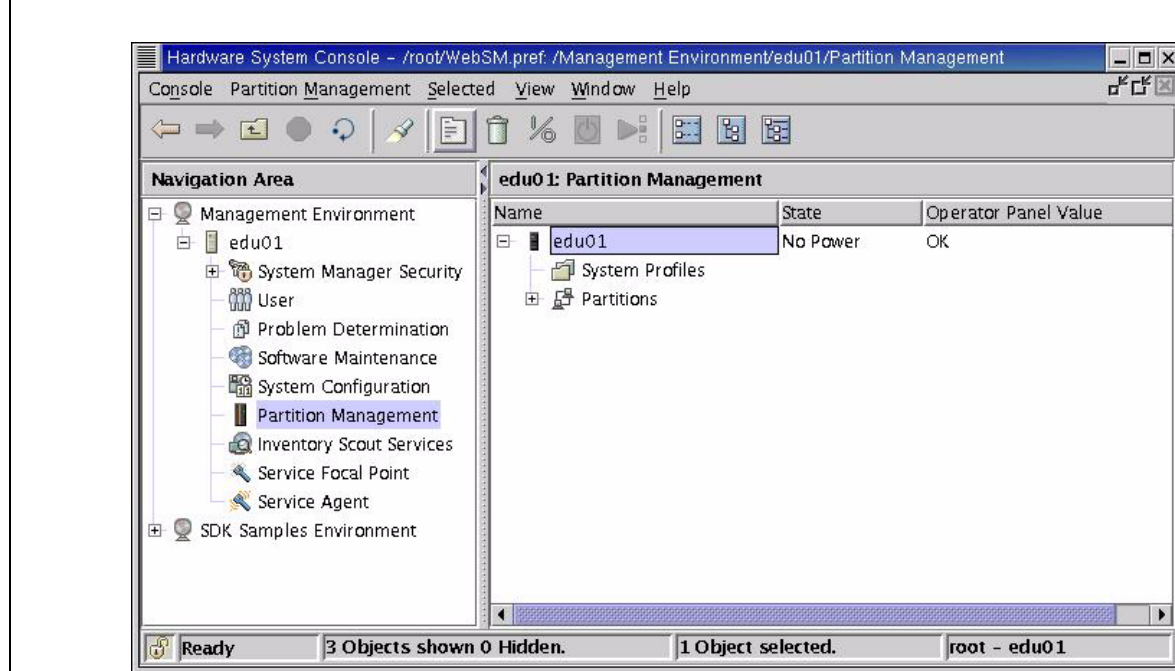


## HMC GUI

The HMC GUI allows you to perform a number of partition management functions. When you click on partition management in the navigation area, you see a list of the systems you can manage in the right pane.

Let's continue to the next slide to see how we can work with system **edu01** in the Partition Management application.

# HMC GUI



## HMC GUI

Within the Partition Management application, you can create system profiles or partitions for a given system. A partition profile defines the set of resources required to create a partition. More than one profile can be created for a partition, but only one partition profile can be active for a partition. You can use alternate profiles to start a partition with a different number of resources or to start partitions in different boot modes.

A system profile applies to the physical server, not to specific LPARs.

To create partition profiles, the user must be a member of either the System Administrator or the Advanced Operator role.

# HMC Inventory Scout



more info at : <http://techsupport.services.ibm.com/server/aix.invscoutMDS>

## HMC Inventory Scout

Inventory Scout Services is a tool that surveys the managed system for hardware and software information. This tool also provides a customized report indicating the latest microcode level. Inventory Scout Services helps users keep track of software updates and patches on managed systems.

Inventory Scout Services can be used to perform the following tasks:

- Inventory Scout Services Profile Configuration
- Conduct Microcode Survey
- Collect Vital Product Data (VPD) Information
- Restart the Inventory Scout Daemon

## **HMC Inventory Scout**

**Checks the firmware or microcode level of all the devices in the system and advises which levels need attention**

**More information at:**

**<https://techsupport.services.ibm.com/server/aix.invscoutMDS>**

## **HMC Inventory Scout**

HMC Inventory Scout can be used to determine if your pSeries server is at the latest microcode level. There are two ways to use this service: 1) Using a secure Internet connection or 2) Create a data file on the system and upload the data file to IBM.

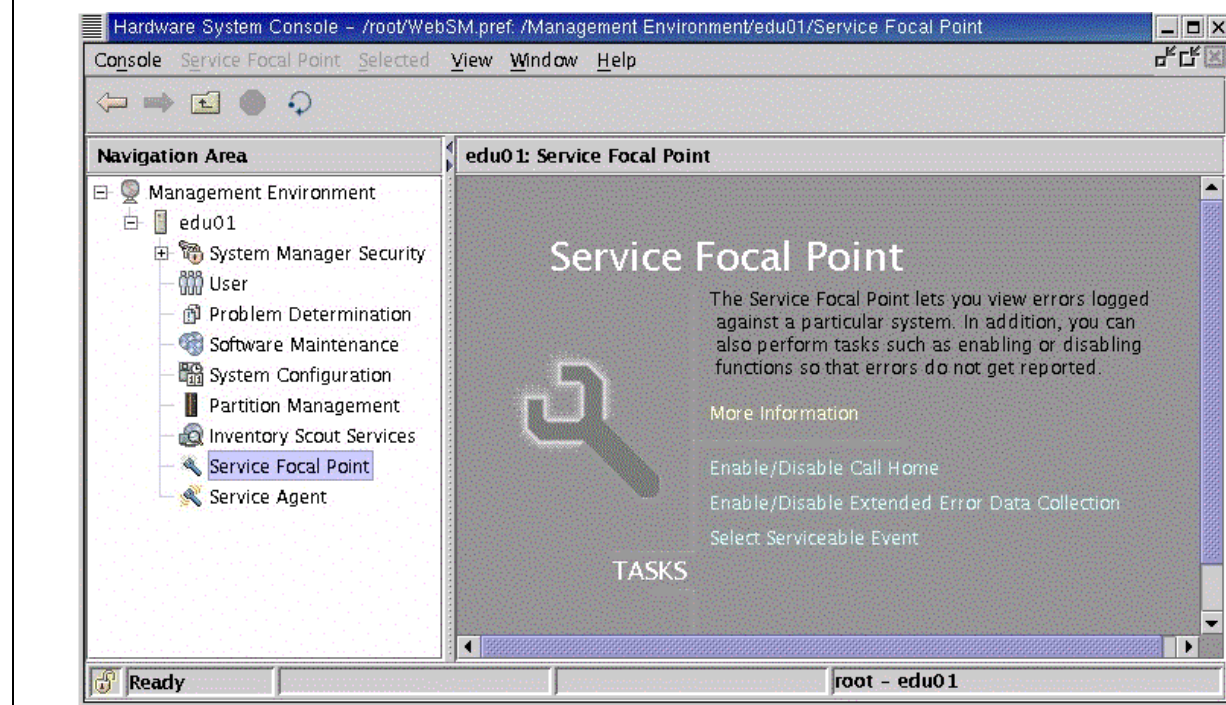
If you choose to use the secure Internet connection, you will be using a signed, trusted Java™ applet that runs on your Internet-connected workstation. This applet connects to the Inventory Scout application running on the server and captures hardware and microcode data. The workstation or PC running the browser must be able to connect to the pSeries server where the microcode check is to take place.

Alternatively, if the pSeries server is disconnected from the Internet, a data file can be created on the server and then uploaded to IBM from an Internet-connected system. To create the data file, you need to run an Inventory Scout command on the server.

For more information about using the Inventory Scout to check microcode levels, refer to the Microcode Discovery Service site listed in the “Hotlinks” section of this course.



# HMC Service Focal Point



## HMC Service Focal Point

Service Focal Point is used by service representatives and system administrators to view operating system error logs. It is used to diagnose and repair problems on partitioned servers. Service Focal Point can also be used to enable or disable functions so that errors do not get reported.

Traditional service strategies become more complicated in a partitioned environment. Each partition runs on its own, unaware that other partitions exist on the same physical server. If one partition reports an error for a shared resource, such as a managed system power supply, other active partitions report the same error. The Service Focal Point application enables service representatives to avoid long lists of repetitive call-home information by recognizing that these errors repeat, and by filtering them into one error code.

The following types of errors are reported to the Service Focal Point:

- Permanent hardware errors (detected by the managed system or operating system)
- LAN surveillance errors
- Hardware boot failure errors

# HMC Service Agent



## HMC Service Agent

The Service Agent application accepts hardware errors from the Service Focal Point. Service Agent reports serviceable events, assuming they meet certain criteria for criticality, for service without requiring customer intervention.

Service Agent enables automatic problem reporting. This means that service calls can be placed without customer intervention. Service Agent also makes possible the automatic sending of extended vital product data to IBM, as well as automatic customer notification. Service Agent also delivers network environment support with a minimum number of telephone lines required for modems.

Any user can access the Service Agent application.

## **HMC Service Agent**

**Automatic Problem Analysis**

**Problem-definable threshold levels for error reporting**

**Automatic Problem Reporting - service calls placed to IBM without intervention**

**View Hardware Event Logs for hardware errors**

**High Availability Cluster Multiprocessing (HACMP) support for full fallback**

- Includes High Availability Control Work Station (HACWS)

## **HMC Service Agent**

The Service Agent can be configured to automatically send error information to IBM. You set this up by using the Service Agent to define each server. After the servers are defined, they are registered with the IBM Service Agent Server (SAS). During the registration process, an electronic key is created, which becomes part of the resident Service Agent program. The IBM Service Agent Server verifies the current customer service status from the IBM entitlement database. If the customer is entitled for customer service, the service call is automatically placed.

The Service Agent can also be used to view hardware event logs for hardware related errors. As described on the previous slide, any hardware errors that are central to the physical server, rather than to just a specific partition, such as a power supply failure, those errors are consolidated for the system as a whole.

High Availability Cluster Multiprocessing (HACMP) support is also included in the Service Agent application. HACMP makes use of redundant hardware configured in a cluster to keep an application running, restarting it on a backup server if necessary. This minimizes expensive downtime for both planned and unplanned outages and provides flexibility to accommodate changing business needs. This functionality also includes High Availability Control Work Station (HACWS) support.

## **Agenda**

**Logical partitioning concepts**

**Planning for LPAR**

**Hardware management console (HMC)**



**Working with partitions**

**Dynamic logical partitioning on a pSeries server**

**Summary**

## **Working with Partitions**

Now that you have been introduced to the functionality of the HMC, let's move into working with LPARs. This next section will walk you through various aspects of managing a partition environment. For example, you will see how to create profiles, stopping and restarting partitions, deleting a partition, and powering a partition off.

## **System Profiles**

- **Collection of Partition Profiles**
- **Used to start CEC in System Profile mode**
- **Starts the Partition Profiles listed in the System Profile**
- **Can have many System Profiles**
- **Stored in NVRAM**
- **Can be backed up from HSC to diskette or DVD-RAM**

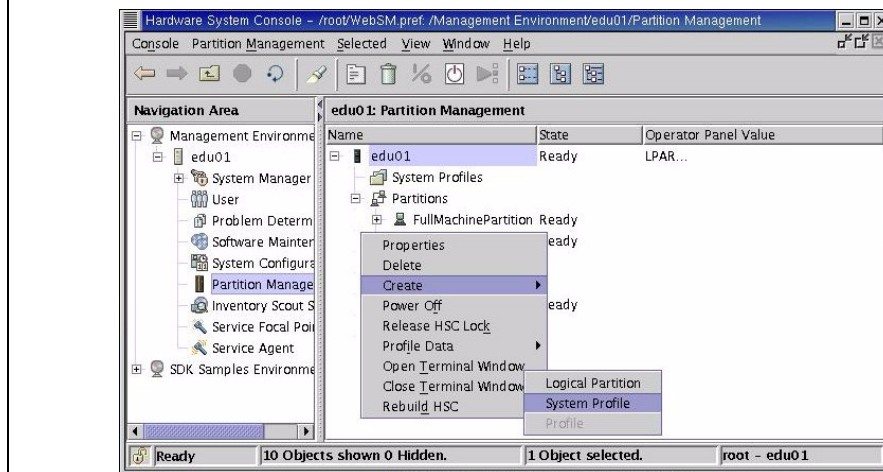
## **System Profiles**

In the HMC, you can create and activate often-used collections of predefined partition profiles. A collection of predefined partition profiles is called a system profile. The system profile is an ordered list of partitions and the profile that is to be activated for each partition. The first profile in the list is activated first, followed by the second profile in the list and so on.

You can have multiple system profiles for the same physical system. As an example, one system profile may contain six partitions which are used throughout the year for normal processing. At the end of the year when sales increase significantly, this system profile can be deactivated and a system profile that contains 10 partitions can be activated to handle the heavier workload.

All system profiles for a physical system are stored in the nonvolatile memory (NVRAM) on the server. Additionally, system profiles can be backed up to diskette or CD-ROM from the HMC.

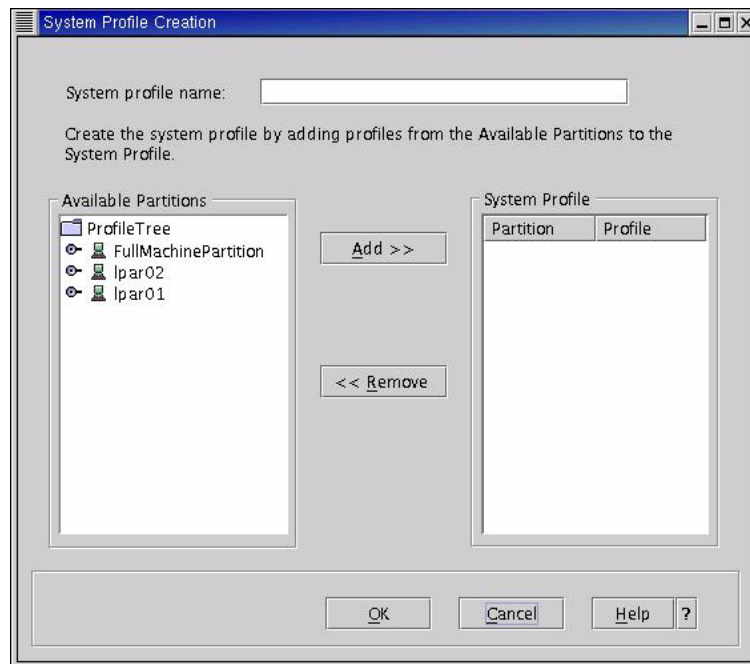
## Creating a System Profile



## Creating a System Profile

To create a system profile, click on **Partition Management** in the navigation area. The system name will show in the right pane. Right click the system name, in this example, the system name is edu01. Select **Create** and then **System Profile**. This will take you to the next screen.

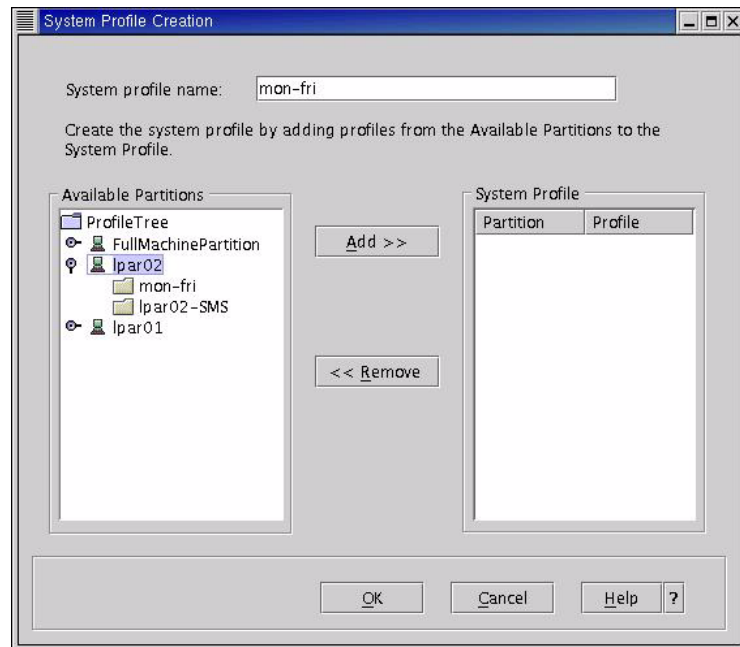
## Creating a System Profile (continued)



## Creating a System Profile (continued)

On the System Profile Creation screen, enter a name for the system profile. You see the available partitions listed under the ProfileTree. In this example, there are three partitions available to add to the system profile being created. They are: FullMachinePartition, lpar02, and lpar01.

## Creating a System Profile (continued)

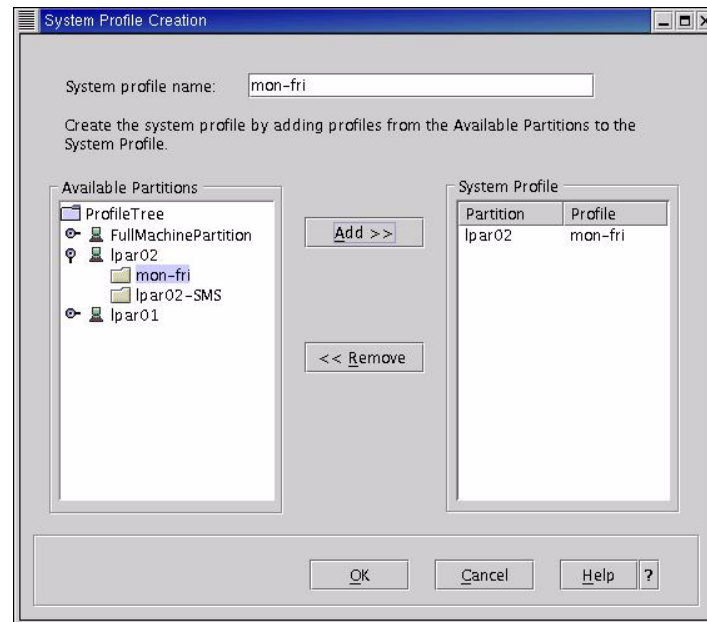


## Creating a System Profile (continued)

In this example, we are creating a system profile called mon-fri that will be used to start partitions that should be active Monday through Friday. By clicking on lpar02, we see mon-fri and lpar02-SMS listed.



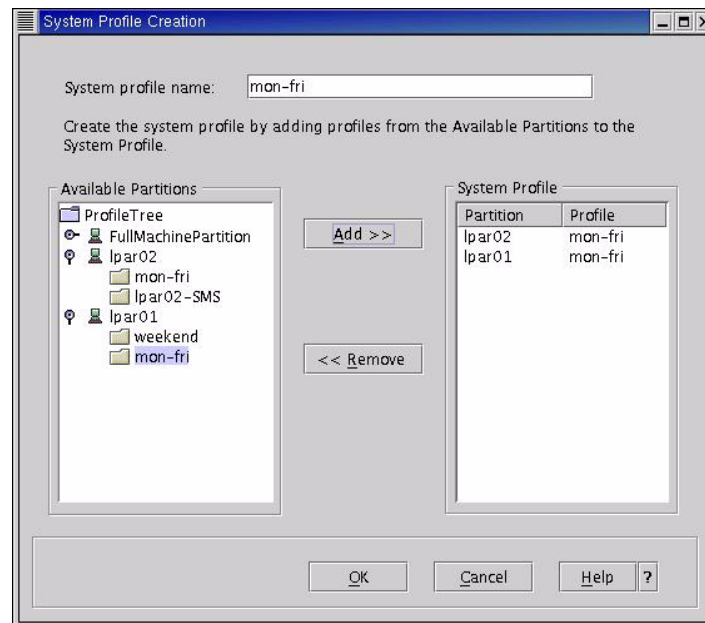
## Creating a System Profile (continued)



## Creating a System Profile (continued)

Because we want to add partition profiles for Monday through Friday to this system profile, highlight **mon-fri** under lpar02 and click the **Add** button. This adds the mon-fri partition profile for lpar02 to the mon-fri system profile.

## Creating a System Profile (continued)

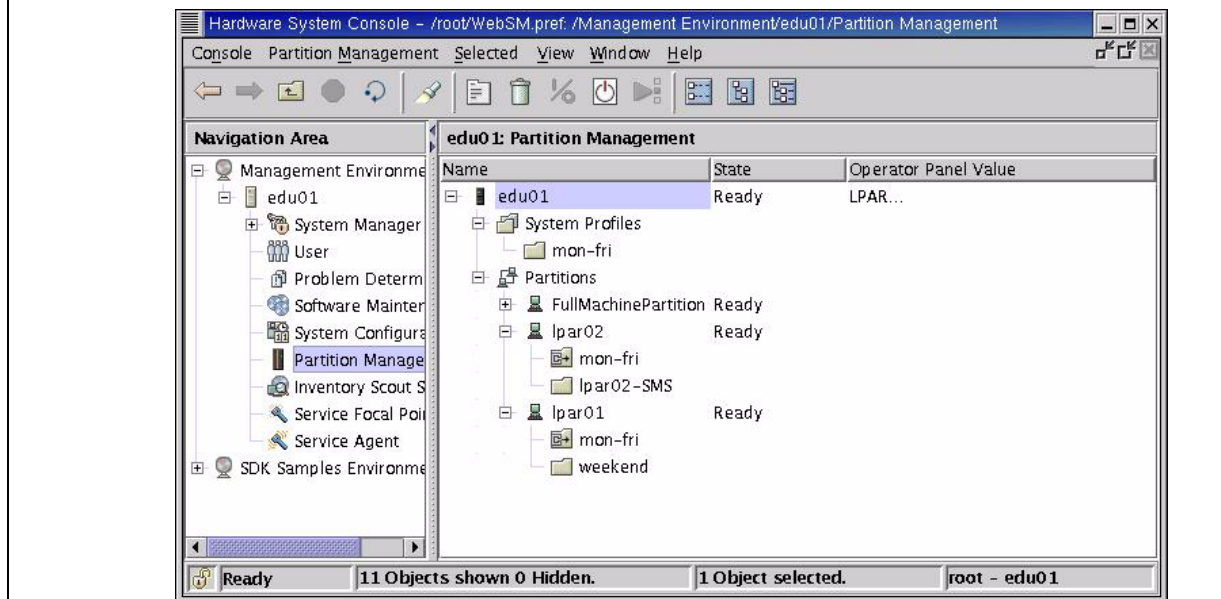


## Creating a System Profile (continued)

By clicking on lpar01, we see that there are two partition profiles defined: weekend and mon-fri. Highlight **mon-fri** and click the **Add** button to add this partition profile to the system profile.

To complete the creation of the system profile, click the **OK** button. This will take you to the next screen.

## Creating a System Profile (continued)

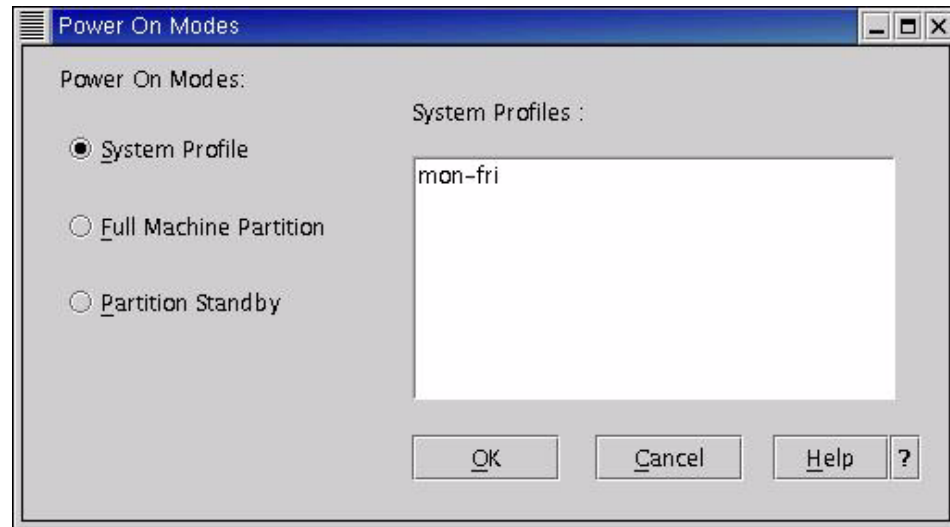


## Creating a System Profile (continued)

Once the partition profile has been created, you will see it listed under the System Profiles section for a particular system. Here we can see system profile **mon-fri** has been created for system **edu01**.

To start this system profile, highlight the system profile named **mon-fri** in the contents area in the right pane and click **Selected** from the menu. From the selected menu, click **Activate**. This will take you to the next screen.

# Starting a System Profile



## Starting a System Profile

You will be prompted for one of three power-on modes for the selected system profile.

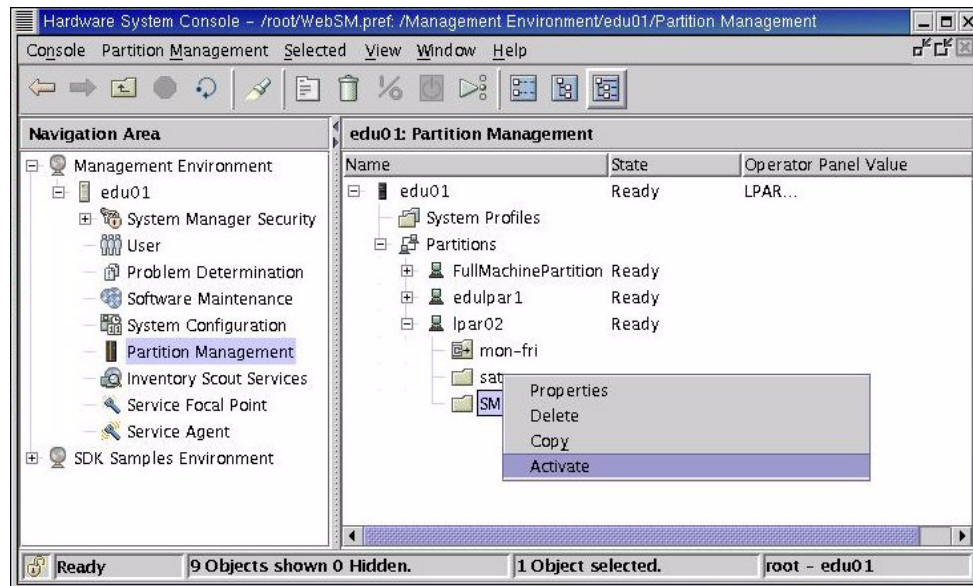
The **System Profile** mode causes the managed system activates partition profiles in the order listed in the given system profile.

The **Full Machine Partition** mode means that only one AIX operating system image is activated, which has access to all resources on a pSeries server. This is also applies to a full system partition, depending on which version of HMC you are using.

The **Partition Standby** mode allows you to create partitions or activate individual partitions. When the partition standby power-on is completed, the operator panel on the managed system displays LPAR, indicating the managed system is ready for you to use the HMC to partition its resources or to activate configured partitions.

We want to select system profile for the power-on mode. Click **OK** to power the system profile on.

## Starting a Specific Partition



## Starting a Specific Partition

To start a specific partition individually, highlight the partition name in the contents area or right pane. Right click on the partition name and select **Activate**.

## **Partition Activation Failures**

**Partition Profiles describe desired and required resources**

**Activation will fail when required resources not available**

**Can fail because of:**

- **Not enough CPUs to meet profile minimum requirement**
- **Not enough memory to meet profile minimum requirement**
- **I/O Slot in required list is not available**
- **Either active in another partition or missing from machine**

## **Partition Activation Failures**


When activating a partition profile, the profile contains both “desired” and “required” resources for CPU, memory, and I/O devices. If the required resources are not available when the partition tries to start, the activation will fail. The reasons for the failure depend on which resource or set of resources were not available to fulfill the required or minimum set of resources for the partition profile.

The partition may not be able to activate because there are not enough CPUs or memory available to meet the defined minimum or “required” for the partition profile. Another cause might be that no I/O slots are available.

A reason that these resources are not available could be that they are already active in another partition, or they have been removed from the server since the partition profile was originally defined.

In the next set of slides, we will show you what some of these errors will look like in case you encounter them when starting a partition profile.

# Partition Activation Failures




## Partition Activation Failures

Here is an example of a partition activation failure error. The message states "There are an insufficient number of processors." This tells you that the required number of CPUs defined for starting the partition profile were not available when the partition was started.

You will either need to add additional CPUs to the server by activating Capacity Upgrade on Demand, or you will need to deactivate another partition profile to free up some CPU resources.

# Partition Activation Failures




## Partition Activation Failures

The error message shown here indicates that the required minimum amount of memory defined in the partition profile was not available when you attempted to activate the partition. The specific error message is "There is not enough memory."

As with the lack of CPU resources being available to start the partition profile in the previous slide, you will either need to add additional memory to the server or you will need to deactivate another partition so you are able to start this partition profile.



## Partition Activation Failures (cont'd)

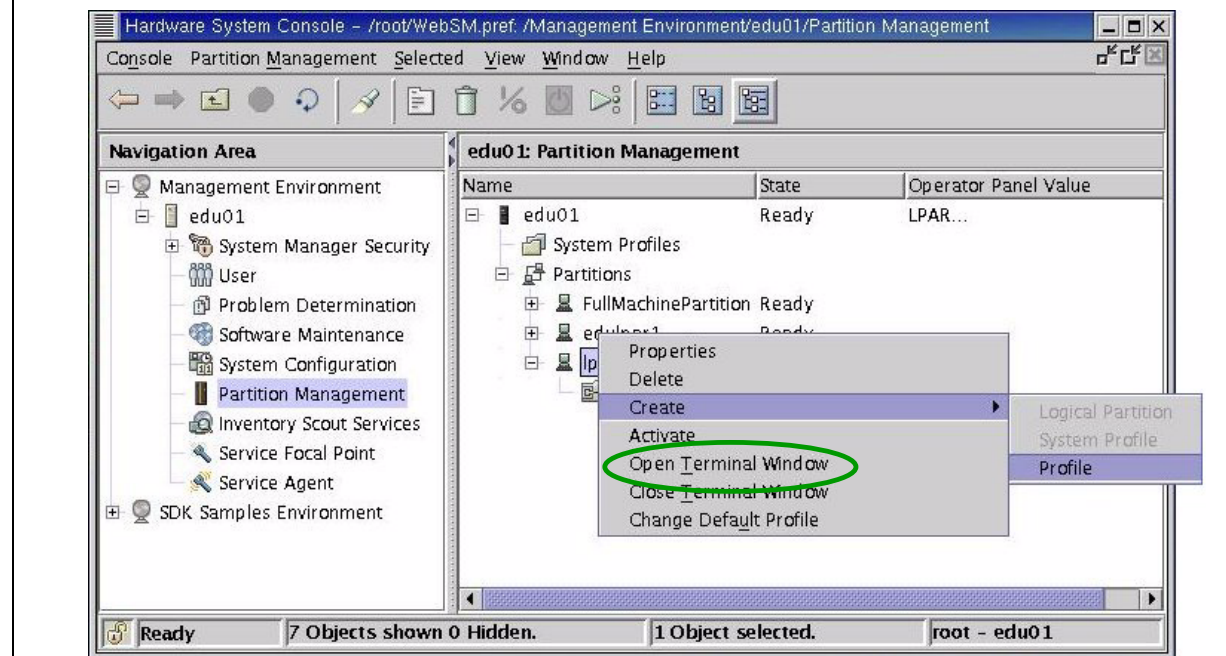


## Partition Activation Failures (continued)

The error message listed in this activation error dialog box looks to be more complicated than it is. This error message is telling you that the I/O slot 2 in the drawer (circled in red) is not available because it is currently being used by lpar02—which has already been activated and thus is already using this resource.

To be able to start this partition profile, you will need to end lpar02 prior to starting this partition.

## Opening a Console




## Opening a Console

AIX needs a console for installation and some service activities. The native serial ports on the managed system are only assignable together to one partition. The virtual terminal window provides virtual terminal console access to every partition without a physical device assigned.

A virtual terminal window is available for each partition or Full System Partition of the managed system. When working with a partition that has been activated, you can **highlight the partition, right click**, and then select **Open Terminal Window**. This will open the console session show on the next slide.

*[NOTE: If you open a virtual terminal window on a partition that has not been activated, the virtual terminal window will be blank until the partition is activated.]*

# Opening a Console



## Opening a Console (continued)

The virtual terminal window displays the server type and model number (circled in red here as “7017-S90”). It also displays the pSeries serial number (circled here in red as “1100091”). The partition name (circled in red as “edu01”) is also in the title bar to designate the partition to which the virtual terminal window is connected.

You are prompted for a console login ID at the bottom of the screen. The virtual terminal window supports the AIX **smitty** and other cursor-driven applications. The virtual terminal window emulates a VT3270 terminal. To set the terminal type on a virtual terminal window session, you can use the AIX **export TERM=vt320** command on the Korn shell prompt.

## **Virtual Console**

**One per partition**

**Java based**

**Management tasks only**

**Do not use for production applications**

## **Virtual Console**

You can open a virtual terminal any time, regardless of the state of a partition, but only one virtual console can be open per partition. The virtual terminal window is Java based, so what you are seeing is a virtual console served in the Korn shell environment from the AIX partition.

The virtual console should only be used for installation and services purposes. The virtual console does not support the following functions:

- Printing to a virtual console
- Transparent print services
- Modem connection for the virtual console port, or
- Real-time applications

## Stopping and Restarting Partitions

### Best way to stop a partition

- Login to AIX instance and use the shutdown command

### When all partitions are stopped, machine may power off

- Changeable in CEC Properties

### If AIX is not responding, can send OS reset from HSC

- Choice of hard reset or soft reset
- Soft reset is equivalent of pressing reset button on regular system
  - System might kernel dump (if Always Allow System Dump set TRUE)
- Hard reset equivalent to pressing the power button

## Stopping and Restarting Partitions

There are multiple ways to stop an AIX partition. The recommended method is to log into an AIX instance and use the **shutdown** command. To do this, you can use a virtual terminal window and issue the **shutdown** command, or you can Telnet into the partition and issue the **shutdown** command.

You need to end all partitions before you power off the server that is hosting the LPARs. You should verify the state of each partition or the Full System Partition has been changed from Running to Ready before you power the server off.

To power off the managed server, select the managed server in the Contents area (right pane) of the HMC, and select **Selected -> Power Off** from the menu bar. If you try to power off a server with active partitions, you will receive a warning message, allowing you to end the powering off function and end the partitions with the shutdown command.

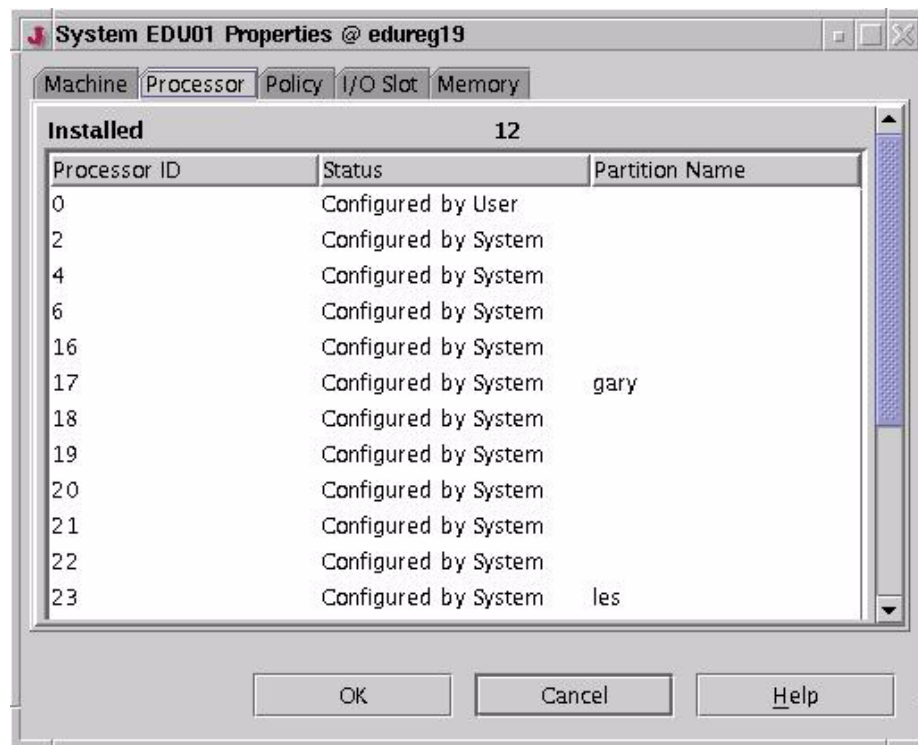
If the operating system in a partition is not responding, you can use the HMC to reset the operating system. **Note:** This is something you only want to do as a last resort, as doing so may corrupt data on the partition being reset.

To perform a reset, select the partition in the HMC GUID that will not end via the normal **shutdown** command. Select **Selected -> Operating System -> Reset** from the menu bar. You will be presented with two options: soft reset and hard reset.

What happens with a **soft reset** depends on the policy settings for the partition. Depending on how the policy settings have been configured, the operating system may perform a dump of system information or it may restart automatically. Selecting the soft reset option is like pressing the reset button on the managing system.

The **hard reset** acts as a virtual powering off of the partition, not the managed system. Issuing a hard reset forces termination and can corrupt information. Use this option only if the operating system is stalled and cannot send or receive commands.

# CEC Properties



## CEC Properties

When you open the properties for the server, you can view the information related to the machine, the processor, the policy, the I/O slot or the memory.

On this slide, we see the processor tab, which details information for different processor IDs on the server. You can see the status of each processor ID and also whether a particular processor or set of processors has been assigned to a specific partition.

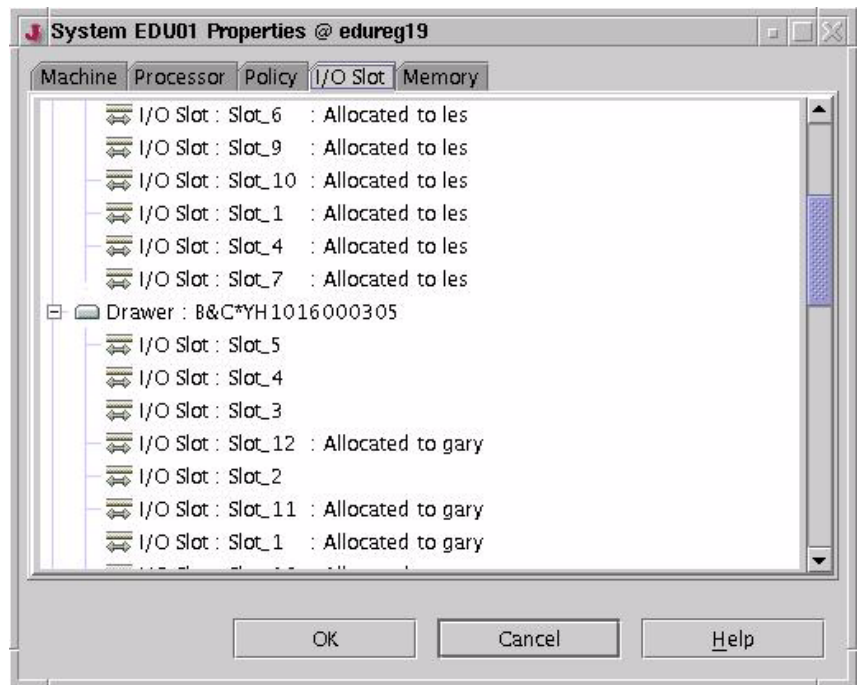
On the policy tab, you are presented with two options for this partition's policy settings:

- "Power off the system after all the logical partitions are powered off," or
- Common Service Processor Surveillance Policy

If you wanted to change the power off behavior, you would select the first option.

To get the properties for the whole server, click on the Central Electronics Complex (CEC) name and then select the properties.

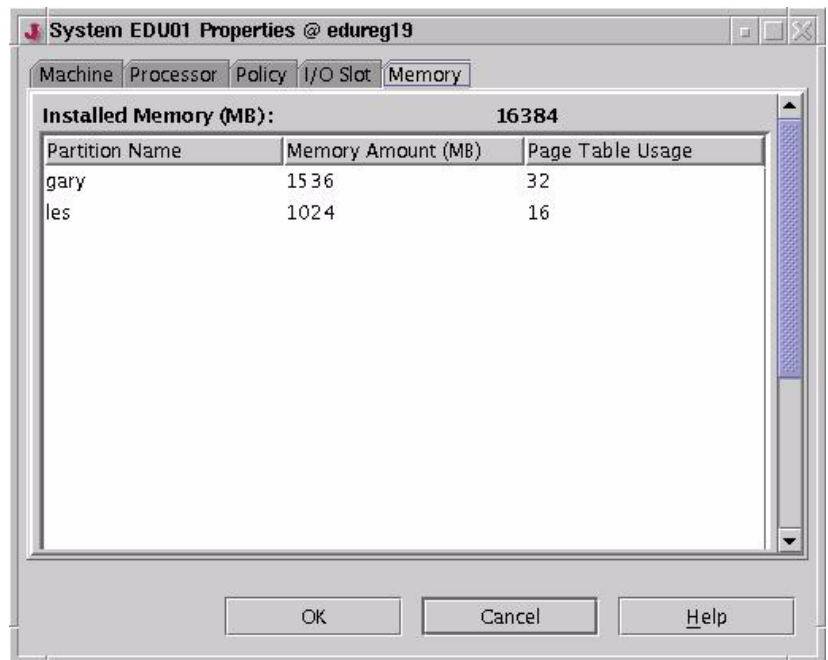
## CEC Properties (continued)



## CEC Properties (continued)

By clicking on the I/O Slot tab, you can see which I/O slots in the different drawers on the system have been assigned to specific partitions. Here we can see that some of the I/O slots have been allocated to a partition called 'les' and others have been allocated to a partition called 'gary.' For the I/O slots that are not listed as allocated to a particular partition, they have not yet been allocated to a specific partition and therefore are still allocated to the overall physical server.

## CEC Properties (continued)

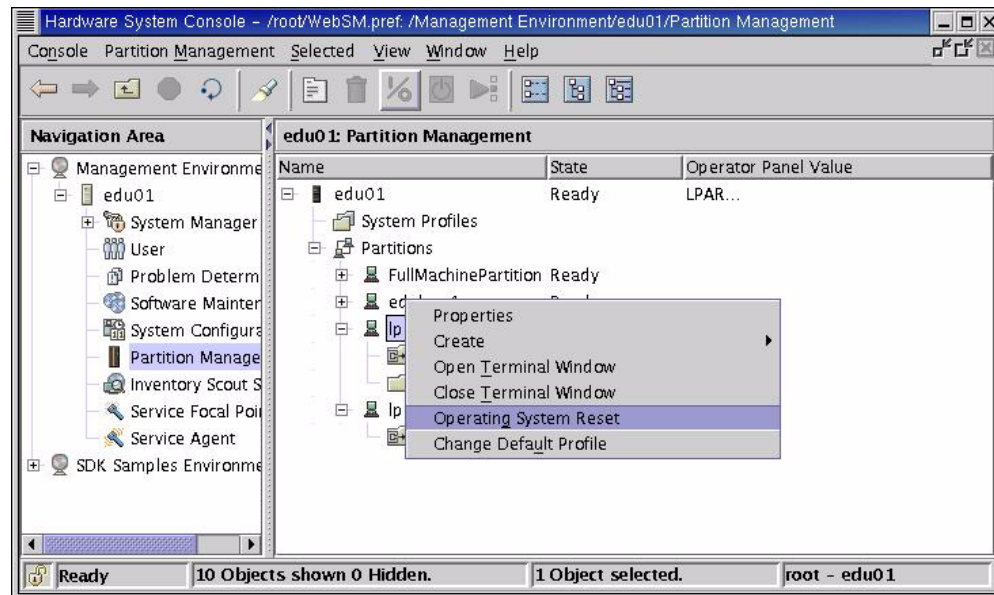


## CEC Properties (continued)

The memory tab shows how much memory is available on the server and which partitions the memory has been allocated to. Here we see there are two partitions that are consuming the memory. The page table size is 32 for partition 'gary' and 16 for partition 'les.'



# Resetting a Partition



## Resetting a Partition

You can use the HMC to reset a partition.

To do this, highlight the partition you would prefer.

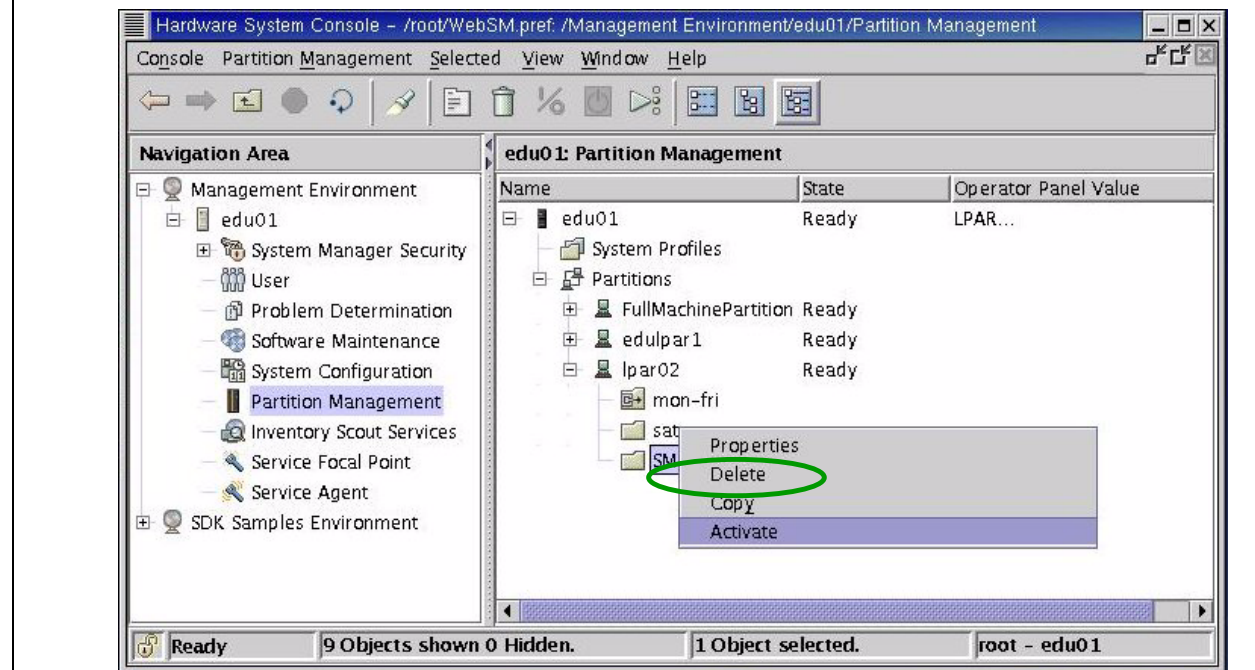
## Resetting a Partition (continued)



## Resetting a Partition (continued)

Here is the next screen you will see as you reset the partition from the HMC.

## Deleting a Partition Profile



## Deleting a Partition Profile

If you determine that you no longer need an LPAR that was previously created, you can remove a partition profile through the HMC. To do this, select the LPAR you wish to delete in the Contents area of the HMC GUI. Right click that partition profile and select **Delete**.

You will be asked to confirm the deletion as shown on the next screen.

## Deleting a Partition Profile (continued)

**Must be confirmed via dialog box**

**Cannot delete default profile**

- Need to make another profile the default

**Can only be done when profile is not active**

- Partition can be active with another profile

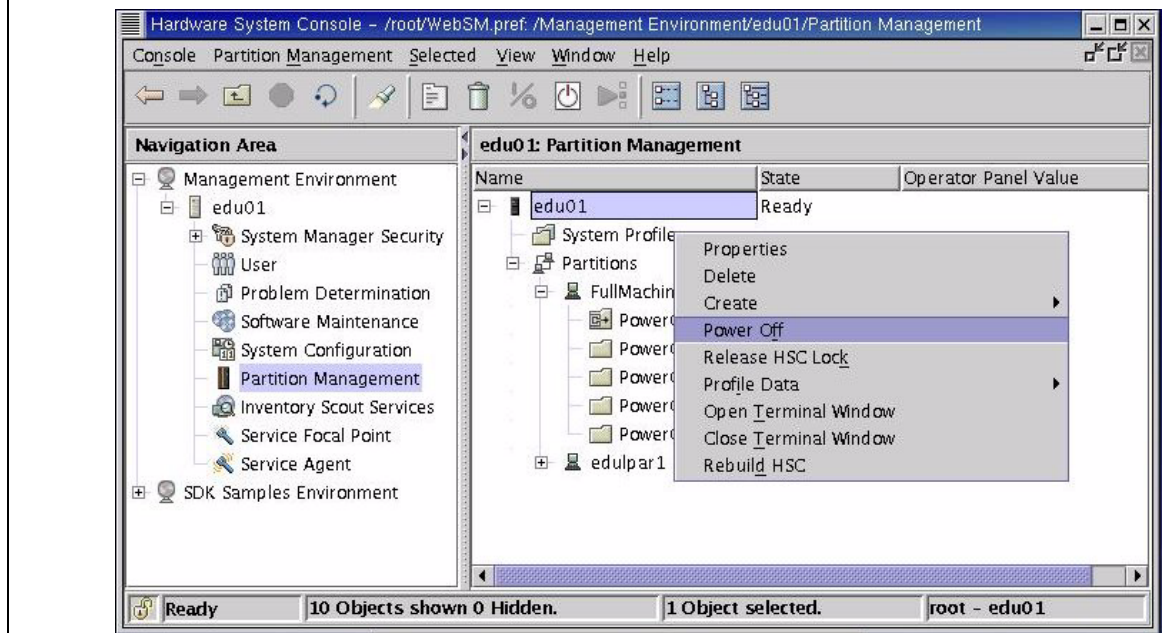


## Deleting a Partition Profile (continued)

Here you see the delete partition profile confirmation dialog box. To delete the selected LPAR, select **Yes**, otherwise select No. If the profile you have chosen to delete is the default profile for the server, you will need to assign another partition profile as the default profile before you can delete this profile.

Additionally, a partition cannot be removed when it is active. In order to successfully delete the partition profile, the partition needs to be shutdown prior to the removal.

# Powering Off



## Powering Off

One of the management functions you may need to do is to power off the server that is hosting the LPAR environment. Prior to performing this function, you will want to end each individual LPAR using the **shutdown** command explained earlier.

To power off a physical server, select the server name in the HMC GUI Content area. In this example, we are shutting down server edu01. Right click on the server name and then select the **Power Off** option.

## Powering Off (continued)




## Powering Off (continued)

When selecting to power the server off, you will be shown this power off confirmation dialog box. If you wish to continue powering off the server, select **Yes**; otherwise select **No**.

If all partitions are currently ended, the server will power off. However, if any partitions are currently active, you will receive the message shown on the next slide.

## Powering Off (continued)

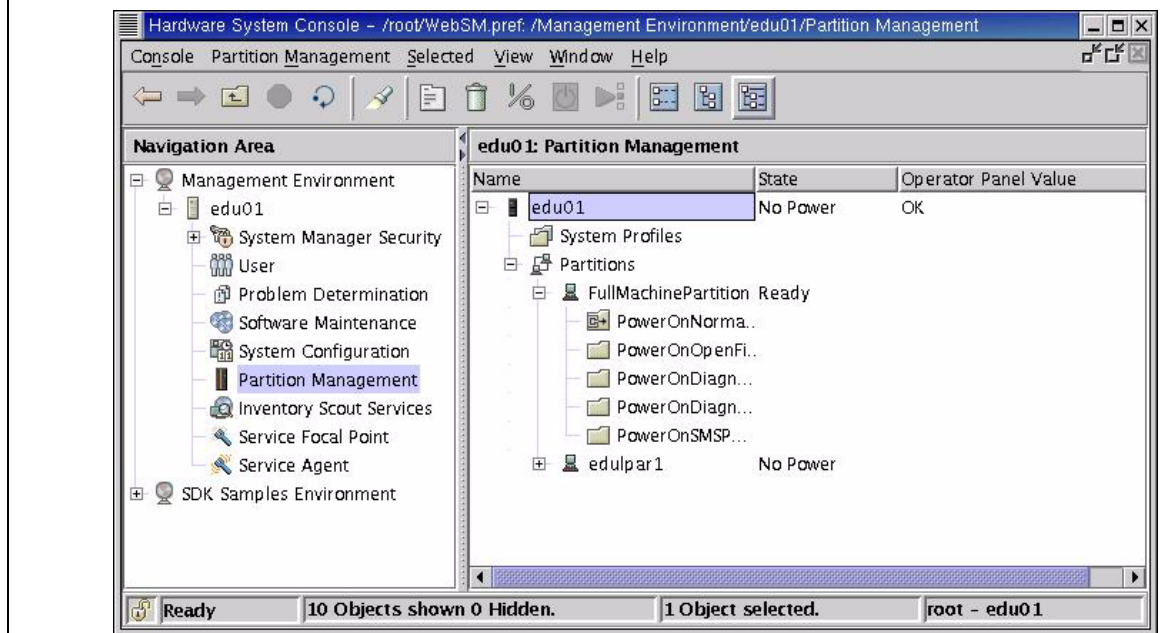


## Powering Off (continued)

On this screen, you can choose to continue with the power off function and have any active partitions automatically ended. If you choose this option, select Yes. If you were unaware that a partition was active when you started the power off option, select **No**. This will allow you to ensure work on the on the active LPAR can be ended in a controlled manner.

As you will recall, it is highly recommended to use the **shutdown** command to end a partition. Once all partitions are ended, you can resume with powering off the server.

## Powering Off (continued)



## Powering Off (continued)

Once the server has been successfully powered off, you will see the status changed to **No Power** as shown in this graphic. At this state, all partitions are deactivated and the physical server has been powered off.



## **Agenda**

**Logical partitioning concepts**

**Planning for LPAR**

**Hardware management console (HMC)**

**Working with partitions**



**Dynamic logical partitioning on a pSeries server**

**Summary**

## **Dynamic Logical Partitioning on a pSeries server**

The next section of this course defines dynamic LPAR and its many advantages over previous versions of LPAR. Up to this point, we have been focusing on LPAR in a rather static way. With the latest versions of the operating system, LPAR has been extended to allow for dynamic LPAR capabilities.

## Defining Dynamic LPAR

**While an LPAR is running, processors, memory, and I/O slots can be:**

- Released into a free pool
- Acquired from the free pool
- Directly moved into another LPAR

**Processor affinity partition mode allows:**

- Only I/O adapter resources to be dynamically reconfigured
- Allocates processors and memory to partitions based on MultiChipModule (MCM) boundaries

## Defining Dynamic LPAR

What is dynamic LPAR? Dynamic LPAR is a function that was added to the AIX operating system in version 5.2. This new functionality allows you to dynamically assign pSeries hardware resources to a partition. These resources include CPU processors, memory, and I/O slots.

With dynamic LPAR, as resources within a partition are no longer needed, they are moved into the free pool. This pool contains a set of resources that are available to be moved to other partitions requiring more resources. As an active partition becomes busier and needs additional resources such as CPU, those resources can be moved to the partition in need of the resource(s).

Resources can also be moved from one LPAR to another. This means resources do not have to be returned to the free pool for the resources to be reassigned. An example will help to explain this situation. If Partition A is very busy and is doing critical work such as receive orders from the Web, and Partition B is not very busy with its workload, resources can be dynamically moved from the less busy Partition B to the very busy Partition A.

If you have chosen to use processor affinity partition mode, there are some dynamic limitations you need to be aware of. First, with this option selected, only I/O adapter resources can be dynamically reconfigured. With this option enabled, if a partition needs additional CPU or memory resource, those will not be able to be dynamically assigned.

Additionally, with processor affinity partition mode enabled, processor and memory allocations to the partition needing the resources will happen on MultiChipModule (MCM) boundaries.

## Moving LPAR Resources

### Steps for moving LPAR resources:

1. HMC sends a network request request to AIX in Partition A, asking it to release a resource and put it in quiesced state. The resource is stopped and put under the control of the Hypervisor.
2. HMC send a command to the Hypervisor asking to relocate the resource from Partition A to Partition B.
3. HMC sends a network request to Partition B to acquire the resource and configure it.

**The move can be initiated thru the HMC GUI or the HMC command line interface**

## Moving LPAR Resources

With dynamic LPAR, the steps to move a resource are transparent; however, we will cover those steps so you understand what is happening under the covers. The steps outlined here involve resources being moved from Partition A to Partition B.

The first step involved in dynamically moving resources in an LPAR environment is to release the required resource from another active partition. In this case, the HMC sends a request to Partition A, asking it to release the required resource. Partition A stops the resource, releases it, and puts it under control of the Hypervisor.

Next, the HMC sends a command the Hypervisor asking to relocate the resource from Partition A to Partition B. The Hypervisor grants the request—provided that the resource has successfully been quiesced and deactivated from Partition A.

The final step is for HMC to send a network request to Partition B, letting that partition know it has access to the needed resource. At this point, Partition B configures this resource into its partition profile and activates the newly acquired resource.

If you are running a version of AIX prior to release 5.2, dynamic LPAR is not available. In that case, an administrator needs to manually move resources between partitions through the use of the HMC GUI or the HMC command line interface.

## **Moving LPAR Resources (continued)**

### **Reconfiguration of processors and memory is well integrated into AIX**

- System administrator needs to take no action to configure these resources.

### **Reconfiguration of PCI I/O slots requires interaction**

- System administrator (sysadmin) must follow PCI Hot plug procedures through the SMIT GUI

## **Moving LPAR Resources (continued)**

The flexibility offered with dynamic LPAR varies depending on which resources need to be moved from one to partition to the next. The ability to dynamically reconfigure memory and CPU or processor units is 100% automated. Moving these resources between partitions can be handled by the dynamic LPAR capabilities of AIX 5.2.

Reconfiguration of PCI I/O slots is a different story. If PCI I/O slots are required to be moved from one partition to the next, these assignment reallocations cannot happen automatically. In this case, a system administrator must intervene. The system administrator will need to use the SMIT GUI to move PCI resources from one partition to another. This administrator will need to use the PCI hot plug procedures to move PCI I/O slot resources.

## **Monitoring DLPAR Operations**

**Detailed reporting can be requested thru the HMC GUI.**

**If the DLPAR operation fails, check the AIX error log.**

## **Monitoring DLPAR Operations**

Dynamic allocation of resources certainly takes an amount of workload off the system administrator, because the resource allocations are done on the fly. However, there may be a need to monitor these dynamic allocation changes to ensure they are working properly.

An administrator can obtain a detailed report of the resources that have been dynamically moved through the HMC GUI. This report will show which resources were moved, what the source and target partitions were, and the time that the resources were dynamically moved.

There is also an AIX error log that can be checked to determine if there were any dynamic LPAR operations that failed. This log is especially helpful when trying to determine why performance of a specific LPAR did not improve as expected.

## Hardware Management Console

### **To prevent problems with DLPAR Operations, Inventory Scout, Service Agent, and Service Focal Point**

- Ethernet network between the HMC and the partitions is a mandatory administrative network

**This network is equivalent to the SP Ethernet network used in the RS6000 SP environment.**

## Hardware Management Console

The HMC provides a very rich set of functionality for working with an LPAR environment. Many of the HMC applications also provide additional information related to dynamic LPAR activities. In particular, the HMC can be used to provide dynamic LPAR information in the following areas:

- Inventory scout
- Service agent
- Service focal point

In order to perform this set of functions, you must have an Ethernet connection between the HMC and the partitions that are being dynamically adjusted. For those of you who are familiar with the RS/6000® SP environment, the network required for dynamic LPAR has equivalent settings.

## **Agenda**

**Logical partitioning concepts**

**Planning for LPAR**

**Hardware management console (HMC)**

**Working with partitions**

**Dynamic logical partitioning on a pSeries server**



**Summary**

## **Summary**

This has been a long course packed with information on the various ways you can use logical partitioning in your pSeries environment. The benefits of LPAR are enormous, providing secure, reliable, and 100% isolated servers inside of a server. This reduces your management tasks greatly—from spotting and resolving errors, to increasing efficiencies and manipulating different partition resources based on prescheduled routines. The POWER4 hardware and Hypervisor firmware mechanisms that drive the pSeries platform work in tandem to deliver flexibility equivalent in many ways to that delivered on IBM's largest platform, the zSeries family of servers.

CPU and memory resources, as well as I/O resources all play an active role in designing and continually tuning logical partitions, so you spent a fair amount of the last hour learning about how to assign and manage these.

Hardware Management Console also took up a large portion of this course material as it is a crucial (required) means of creating and managing LPARs. There were even enumerable screen captures to help you understand the interaction you will experience with the HMC as you learn to finesse your partitions.

There are many resources listed in the Hotlinks section of this course. We think you will find the Redbooks especially beneficial.

Thank you for your time, and good luck as you move into the highly flexible world of logical partitioning.

## Hotlinks/References

Here are some helpful links as you implement LPAR in your environment:

- *pSeries – LPAR Planning Redpiece (TIPS0119)*  
[ibm.com/redbooks](http://ibm.com/redbooks) → search for TIPS0119
- *Effective Management Using the IBM Hardware Management Console for pSeries (SG24-7038-00) Redbook*  
[ibm.com/redbooks](http://ibm.com/redbooks) → search for SG24-7038
- *Complete Partitioning Guide for IBM eServer pSeries Servers (SG24-7039) Redbook*  
[ibm.com/redbooks](http://ibm.com/redbooks) → search for SG24-7039
- *IBM Hardware Mgmt Console for pSeries Installation and Operations Guide, (SA38-0590)*  
[http://publib16.boulder.ibm.com/pseries/en\\_US/infocenter/base/hardware\\_docs/pdf/380590.pdf](http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware_docs/pdf/380590.pdf)
- *Logical Partitioning, Frequently Asked Questions Web site*  
[ibm.com/eserver/pseries/lpar/faq\\_2.html](http://ibm.com/eserver/pseries/lpar/faq_2.html)
- IBM TotalStorage FastT Storage Server  
[www.storage.ibm.com/disk/fastt/](http://www.storage.ibm.com/disk/fastt/)
- IBM TotalStorage Enterprise Storage Server (ESS)  
[www.storage.ibm.com/disk/ess/](http://www.storage.ibm.com/disk/ess/)
- Information about corrective fixes to apply is available at:  
<http://techsupport.services.ibm.com/server/hmc>
- To install **ssh** software on your PC, go to one of the following URLs:  
<http://www.chiark.greenend.org.uk/~sgtatham/putty/>  
[http://HMC\\_System\\_Name/remote\\_client\\_security.html](http://HMC_System_Name/remote_client_security.html)
- For Inventory Scout details/microcode levels:  
<https://techsupport.services.ibm.com/server/aix.invsoutMDS>
- *Guide for IBM eServer pSeries Servers*  
[ibm.com/redbooks](http://ibm.com/redbooks)
- *pSeries 650 Model 6M2 Technical Overview and Introduction*  
[ibm.com/redbooks](http://ibm.com/redbooks)
- *Managing AIX Server Farms*  
[ibm.com/redbooks](http://ibm.com/redbooks)
- *IBM eServer p650 User's Guide (SA38-0611-00)*  
[http://publib16.boulder.ibm.com/pseries/en\\_US/infocenter/base](http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base)
- *IBM Hardware Management Console for pSeries Operations Guide*  
[http://publib16.boulder.ibm.com/pseries/en\\_US/infocenter/base](http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base)



## **Trademarks**

IBM, eServer, pSeries, zSeries, AIX, S/390, TotalStorage, RS/6000, and POWER4 are registered trademarks or trademarks of the IBM Corporation in the United States or other countries or both.

Java and all Java-based trademarks and logos are registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.