iSeries. mySeries.

Common Europe Luxembourg
Introducing IBM eServer i5 & i5/OS

LPAR i5 "Up-to-Date" :
Logical Partitions on the IBM Power5
by Eddy PASTEGER

iSeries. mySeries.

Logical Partitions on the IBM Power5

Introduction

Introduction

Objectives

- Understand the LPAR functionality as implemented on the IBM Power5 servers
- Understand planning considerations and system requirements

Warning

- This presentation is built on functionalities available on i5 models 5xx running i5/OS V5R3

Logical Partitioning on the IBM Power5

Agenda

- Chapter 1. Architecture
- Chapter 2. Resources Management Concepts
- Chapter 3. LPAR-Capable i5 Models
- Chapter 4. Planning
- Chapter 5. Hardware Management Console
- Chapter 6. Virtual Partition Manager

iSeries. mySeries.

Logical Partitions on the IBM Power5

Architecture

Architecture

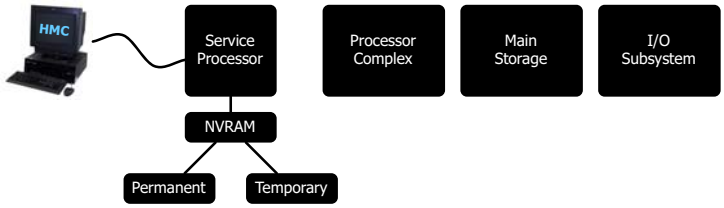

The Service Processor

- Standard built-in, independent component of every eServer i5 or p5
  - → Own processor, memory and boot code
  - → NVRAM stores the "firmware" and the LPAR configuration

Architecture

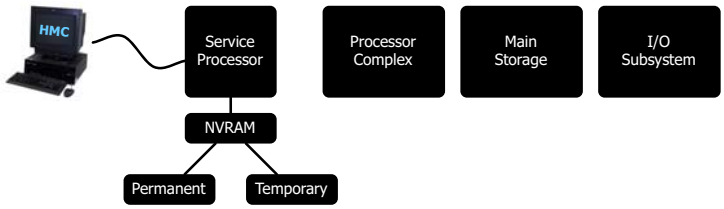

The Service Processor (aka SP)

- Once main power is connected ...
  - → The SP activates its own "firmware" operating system
  - → Control panel shows C1xx progress codes
    - Gathers vital product data (VPD) from the electronic components in the server
    - Runs diagnostic routines to verify its hardware and firmware
  - → Control panel shows 01 and the green power LED is flashing
    - The server is in "SP Standby" mode
    - The server is ready to interface with HMC

Architecture

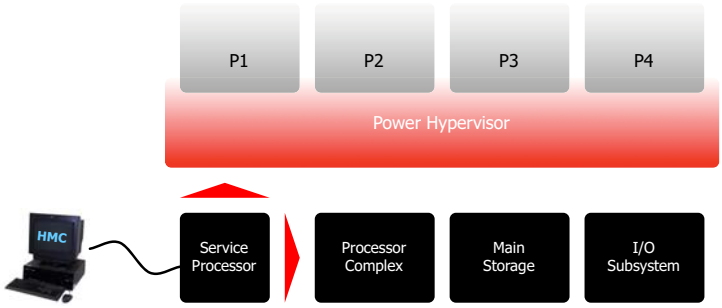

The Service Processor (aka SP)

- Once you power on the system ...
  - → The SP turns on the system power
    - All I/O towers via SPCN
  - → Control panel shows C7xx progress codes
    - Runs diagnostic routines to verify system processors and system main storage
    - Verifies that it can communicate with environmental sensors and monitors in the system
  - → Control panel shows 01 and the green power LED is on
    - The server is in "Hypervisor Standby" mode

Architecture

The Power Hypervisor

- Previously part of the System Licensed Internal Code (SLIC)
  - → Known as "Partition Licensed Internal Code (PLIC)"
  - → PLIC is enhanced to be used with the eServer i5 hardware
- Power Hypervisor is independent from any operating system
  - → There is no longer the concept of a "Primary Partition"
  - → There is no longer the need to IPL the whole system to make partitions changes
  - → There is no longer a single point of failure with a "Primary Partition"
- Also known as ...
  - → Firmware
  - → Server Firmware
  - → Hypervisor
  - → Power Hypervisor
  - → pHyp
  - → Underware

Architecture

The Power Hypervisor Functions

- Virtual processor support
  - → Allocation of physical processor time to partition thru "virtual processors"
- Virtual memory management
  - → Allocation of physical memory to partition thru "virtual memory"
- Virtual hardware support
  - → Processors
  - → OptiConnect
  - → SCSI
  - → Ethernet
  - → Serial
- Security an isolation between partition
  - → Partition allowed to access only to their allocated resources
  - → Enforced by the Power Hypervisor
    - Reinitializes processors, resets registers, flushes caches
    - Clears memory
    - Reset I/O devices

Architecture

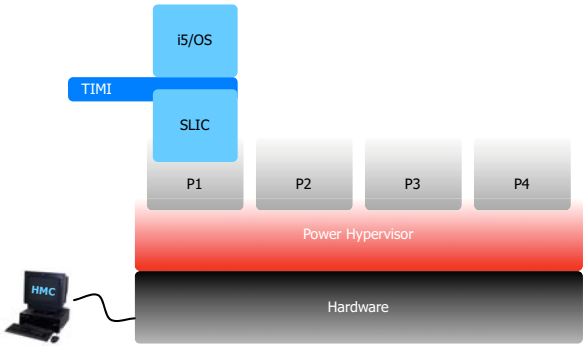

Starting an i5/OS partition

- Hypervisor will operates a "virtual service processor" to load and start SLIC
  - → SLIC is changed to interface with the new Power Hypervisor
  - → TIMI and all the layers above are still in place

Architecture

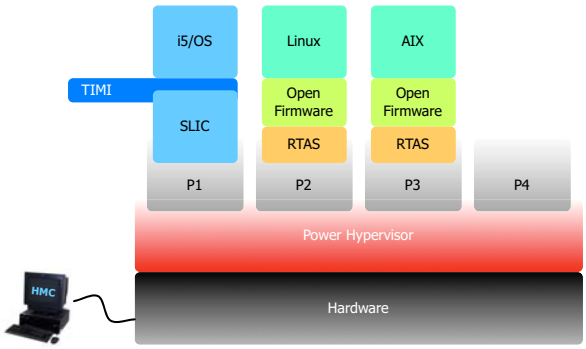

Starting a LINUX or AIX partition

- Hypervisor will operates ...
  - → Open Firmware contains boot-time drivers and the boot manager
  - → RTAS is a service that passes platform-dependent system calls to the Hypervisor
- Other is characteristic to the guest operating system
  - → E.g. Linux has a "monolithic" kernel, while AIX has a "microkernel-hybrid" structure

Architecture

Hypervisor Code Maintenance

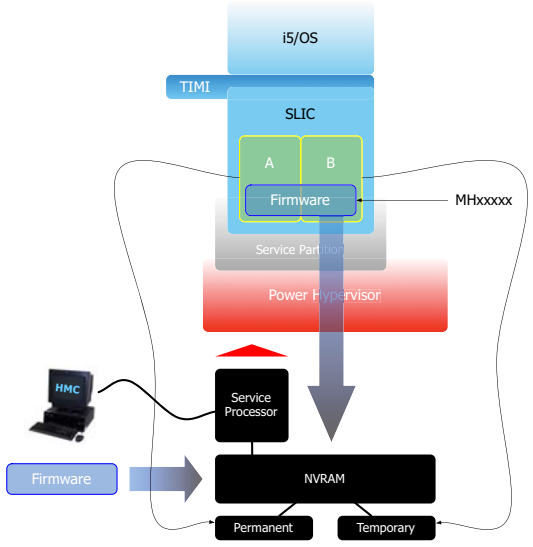

iSeries. mySeries.

Logical Partitions on the IBM Power5

Resources Management Concepts

Resources Management Concepts

What is "logical partitioning" ?

- The ability to make a single server run as if it were many independent systems
- Each logical system is called a "Partition"
- Each partition operates as an independent system
- Each partition will receive a portion of the system resources :
  - Processor
  - Memory
  - I/O slots
  - Virtual I/O slots

iSeries. mySeries.

Logical Partitions on the IBM Power5

Resources Management Concepts
Processor

Resources Management Concepts

Processor Concepts

- Dedicated vs. Shared Processors
- Virtual Processors
- Capped vs. Uncapped Processors
- Considerations about LPAR and SMT

Processor Concepts

Processor Concepts

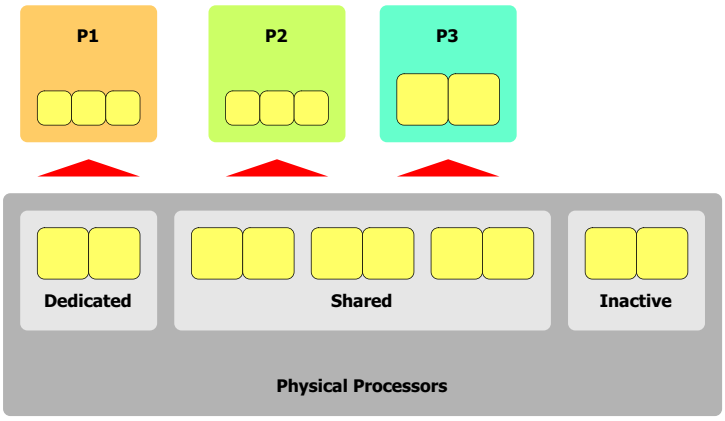

Processor Concepts

Dedicated Processors

- A "dedicated processor" refers to a whole processor that is dedicated to a single partition
- One or more processors can be dedicated to a partition
  - → The most simple way to distribute processors among logical partitions
- Rules & requirements :
  - → At least one dedicated processor
  - → Granularity of movements : one processor
  - → Processor moves are dynamic, without any IPL within the configured ranges
  - → Unallocated processors are not used

Processor Concepts

Shared Processors

- A "shared processor" allows to assign partial processors to a partition
- Physical processors lives in a "shared processor pool"
  - → By default, any unassigned processor is a member of the shared processor pool

Processor Concepts

Shared Processors

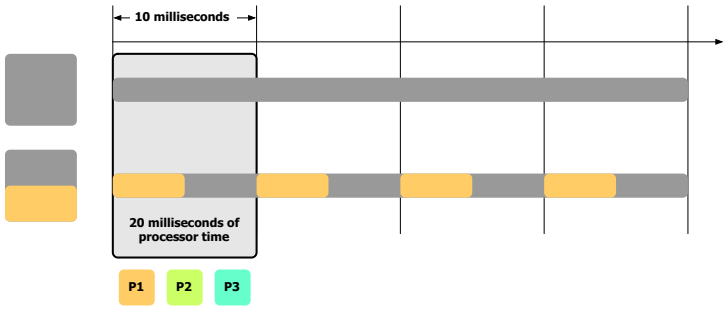

Processor Concepts

Shared Processors

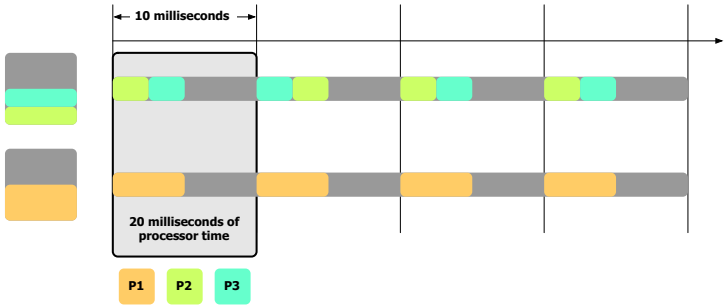

Processor Concepts

Shared Processors

- Rules & requirements :
  - A processing unit (PU) represents the cycles of a single processor
  - At least 0.10 PU (1 millisecond)
  - Granularity of movements : 0.01 PU
  - Processor moves are dynamic, without any IPL within the configured ranges
  - Unallocated PU are not used

Processor Concepts

Potential Shared Processors Penalty

- Increase the possibility that job's won't complete, and :
  - → Have to be re-dispatched, and potentially have to reload cache
  - → Increase the chance of a cache miss
- Reduce the chance for processor/memory affinity
- Increase the Power Hypervisor work :
  - → Tracking each partition use of its allocated milliseconds

Processor Concepts

Virtual Processors

- The number of "virtual processors" represents :
  - → The number of physical processors that the partition is able to use
  - → The number of simultaneous jobs/threads execution
  - → The number of processors that the operating system thinks he has to use

Processor Concepts

Virtual Processors

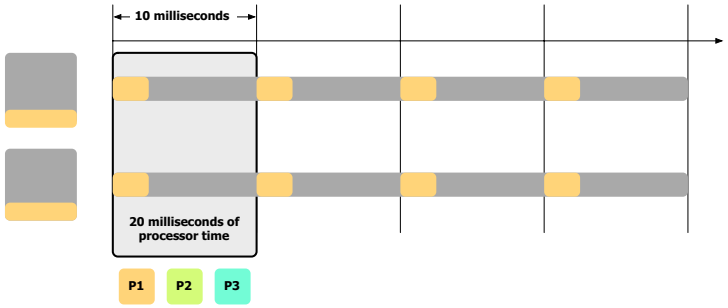

Processor Concepts

Virtual Processors

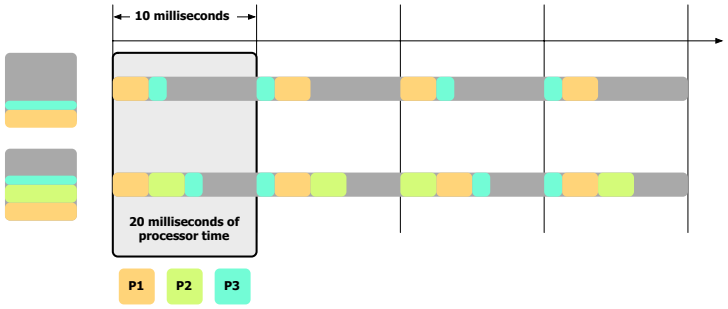

Processor Concepts

Virtual Processors

- Rules & requirements :
  - → It is possible to allocate more virtual processors than there are in the shared processor pool
    - The actual number of processor in the shared processor pool is a "floating" number
    - A single physical processor can report in the system as a n-ways
  - → You cannot allocate less than 0.10 PU per virtual processor
    - Example : requesting 0.20 PU to run on 4 VCPU means 0.20 / 4 = 0.05 < 0.10 !
  - → Changes are dynamic, without any IPL within the configured ranges

Processor Concepts

Uncapped Processors

- Uncapped allows to potentially utilizes unused processor capacity in the system
- Uncapped weight is used to determine how unused capacity is distributed :
  - → P1 weight is 200
  - → P2 weight is 100
  - → P1 will receive 2/3 of the unused shared processor capacity
  - → P2 will receive 1/3 of the unused shared processor capacity
  - → If P1 does not need all 2/3 of capacity, then P2 could get more !

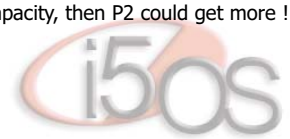

Processor Concepts

Dedicated, Shared, Capped or Uncapped ?

- How to decide on using :
  - → Dedicated or shared processors ?
  - → If using shared processors, it is to be capped or uncapped ?
  - → And ... how many virtual processors do I allocate ?
- The best performance may be achieved by using dedicated processors ...
  - → However, dedicated processors cannot utilizes excess capacity !
- Shared uncapped will allow use of excess capacity of the processor, but ...
  - → Setting virtual processor number too low will limit uncapped usage
  - → Setting virtual processor number too high will degrade performance
- What does IBM recommends ?
  - → It depends ...
- What do I recommend ?
  - → Keep things simple !

Processor Concepts

Simple Recommendations

- On small configurations, use shared uncapped processor :
  - → Allocate desired processing units according needs
  - → Keep a little portion of unused processing units
  - → Configure uncapped weight according partition importance
- Set virtual processor number to a "reasonable" value :
  - → Set a value "close" to the number of allocated processing units
    - 0.4 PU running on 1 VCPU
    - 1.8 PU running on 2 VCPU
  - → Take care of your workload pattern
    - 1.0 PU running on 1 or 2 VCPU ?
    - Single job, single thread applications will not take benefits when running on multiple VCPU
    - Multiple job, multiple thread applications will run more efficiently when running on multiple VCPU
    - Take care of SMT capabilities !
  - → Do not exceed the number of physical processors

Processor Concepts

About Simultaneous Multi-Threading (SMT)

- SMT is a technology that allows a single CPU to be seen logically as 2 logical CPU
  - → Physical processor
  - → Logical processors
- SMT allows 2 threads to run simultaneously on the same processor
  - → Using different execution units, or "pipes"
- SMT optimizes the usage of the 7 execution units of a PowerPC processor
- Hypervisor only deals with physical processors
  - → Each operating system can use the physical processor as its needs
  - → SMT can be used or not ... according partition's operating system configuration
    - For i5/OS, refer to system value QPRCMLTTSK

iSeries. mySeries.

Logical Partitions on the IBM Power5

Resources Management Concepts

Main Storage

Main Storage Concepts

Memory Allocation Facts

- Power4 Hypervisor allowed memory manipulation at the single megabyte (1MB) level
- Memory in the Power5 machines cannot be manipulated that finely :
  - → Memory must be assigned or moved in blocks
  - → Theses blocks are called "Logical Memory Blocks (LMB)" or "Memory Regions"
- About those "Memory Regions" :
  - → LMB are sized in multiples of 16 MB
    - Future implementations will allow use of 16, 32, 64, 128 or 256 MB blocks
    - LMB size is a system-wide parameter
  - → LMB is the smallest amount of memory that can be manipulated

Main Storage Concepts

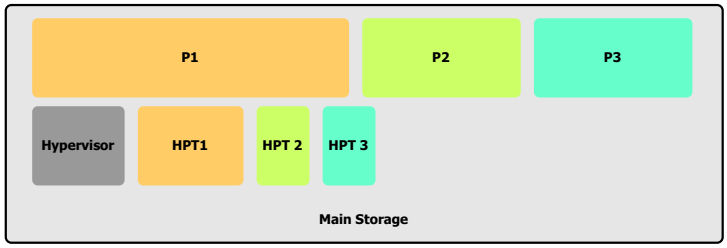

Memory Allocation Facts

- Hypervisor uses memory to run the system !
- Hypervisor memory used expands and contracts based on many variables :
  - → Number of partitions
  - → Amount of I/O attached to the system
  - → Amount of virtual I/O slots used in each partition
  - → HSL OptiConnect and Virtual OptiConnect links
- Starting and/or stopping a partition affects hypervisor memory requirements
  - → Hypervisor memory allocation made by 16 MB increments

Main Storage Concepts

Memory Allocation Facts

- A "Hardware Page Table (HPT)" is needed to manage memory
- Power4 Hypervisor reserved HPT within partition memory
- Power5 Hypervisor allocates HPT separately from partition memory
- Size of the HPT depends of the maximum memory amount the partition can receive
- Basic calculation : maximum memory divided by 64, rounded up to next power of 2
  - → A partition can allocates a maximum of 12 GB (12288 MB)
  - → HPT size will be 12288 / 64 = 192 ... next power of 2 is ... 256
  - → HPT size will be 256 MB
- Recommendation
  - → Take care of HPT when allocating memory to partitions !
  - → Plan to allocate 12288 - 256 = 12032 MB for the partition
  - → Memory actually used will be 12288 MB
- Main Storage size MUST be carefully planned !

Main Storage Concepts

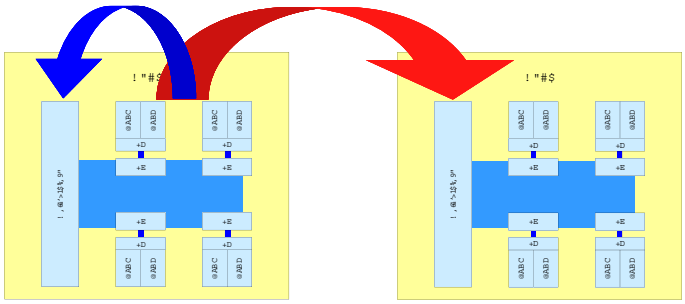

Memory Allocation Facts

- Remember
  - → A node is a packaged set of processors, caches (L1+L2+L3) and memory on a module
  - → Remember about ... "home node", "local node" and "remote node" ?
  - → Remember about ... "local memory access" and "remote memory access"
  - → Remember about ... "memory affinity" ?

Main Storage Concepts

Memory Allocation Facts

- Power4 Hypervisor allocated memory equally from all nodes
  - → This leads to more "remote" memory access
- Power5 hypervisor tries to do better ...
  - → Processors are assigned on the fewest nodes possible
  - → Configured memory will be assigned keeping as much memory as possible on the same node
  - → The HPT for the partition will be assigned on the same node
- How becomes memory affinity with dynamic processor/memory moves ?
  - → The system will "do its best" ...
  - → But there is no guarantee that the best configuration can be maintained
  - → A full system restart (hypervisor) may be needed to return system in its optimal configuration !

iSeries. mySeries.

Logical Partitions on the IBM Power5

Resources Management Concepts
I/O Subsystem

I/O Concepts

I/O Facts

- I/O allocations are done at slot level
  - → A single slot can be allocated
- I/O structure must be understood for proper configuration !

I/O Concepts

Hierarchy of Microprocessors

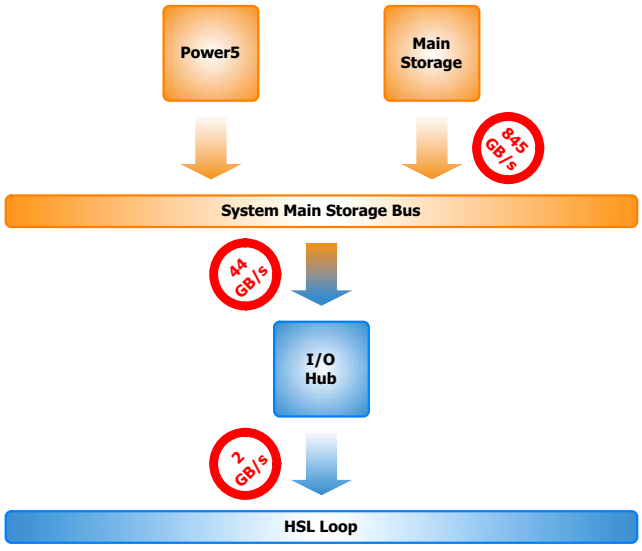

I/O Concepts

Hierarchy of Microprocessors

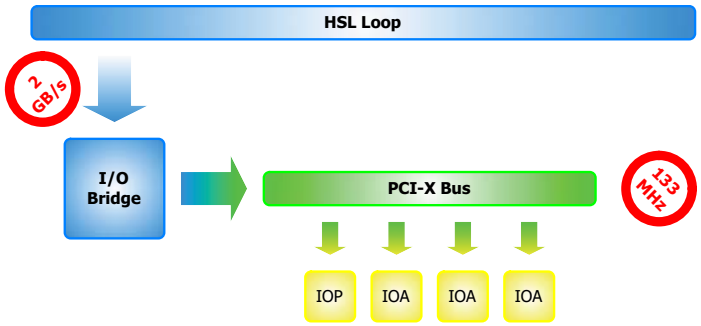

I/O Concepts

I/O Facts

- I/O allocations are done at slot level
  - → A single slot can be allocated
  - → BUS/IOP/IOA requirements must be met !
- Allocate a slot as "desired"
  - → If the resource is available at allocation, then grab it
  - → Otherwise continue to activate anyway
- Allocate a slot as "required"
  - → If the resource is available at allocation, then grab it
  - → Otherwise do not continue the activation process

I/O Concepts

IOP-Less IOAs

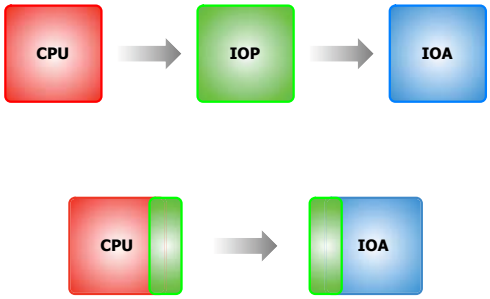

I/O Concepts

I/O Facts

- Once allocated, some I/O resources needs to be identified
  - → Load-Source Unit resource
    - Tells the hypervisor which IOA drives the LSU
    - This is mandatory
  - → Alternate-IPL resource
    - Tells the hypervisor which IOA drives the alternate IPL device
    - This is optional
  - → Console resource
    - Tells the hypervisor which IOA drives the console resource
    - This can be the HMC, thru virtual I/O
    - This is mandatory
  - → Operations Console resource
    - Tells the hypervisor which IOA will be used for ECS
    - This is optional
  - → Alternate Console resource
    - Tells the hypervisor which IOA drives the alternate console resource
    - This is optional

iSeries. mySeries.

Logical Partitions on the IBM Power5

Resources Management Concepts
Virtual I/O

Virtual I/O Concepts

Virtual I/O Facts

- Each partition has virtual I/O slots
  - The number of slots is configurable
    - The more virtual slots you configure, the more the hypervisor will need memory !
  - Similar in concept to physical I/O adapter slot
  - Each slot can be populated with a virtual adapter instance
    - Serial, Network (Ethernet) or SCSI (Server or Client)
- Adapters are defined into partition configuration profile
  - For each slot, the configuration profile contains adapter type and parameters
  - Adapters are "created" by the hypervisor when the partition is activated

Virtual I/O Concepts

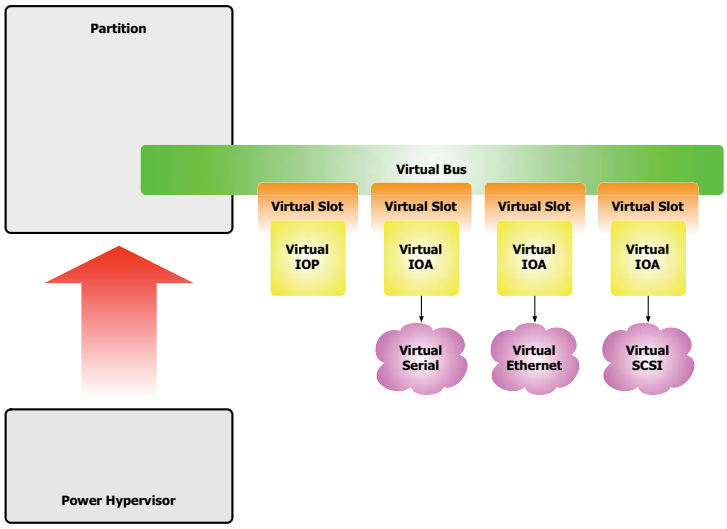

Virtual I/O Concepts

Virtual I/O Facts

- What is dynamic
  - → As for physical I/O adapters, virtual I/O adapters can be dynamically added, changed or removed
  - → Unlike physical I/O slots, virtual I/O slots cannot be moved from one partition to another
- What is not dynamic
  - → Changing the number of slots requires a partition deactivation and reactivation

Virtual I/O Concepts

Virtual Serial

- Virtual serial adapter provides point-to-point connection between
  - The partition and the HMC
  - The partition and another partition
- Purpose
  - Virtual console
  - Virtual terminal
- Parameters
  - Slot number
  - Adapter type : server or client
  - Connection information : who can connect ...
    - The HMC
    - Any remote partition
    - Selected remote partition

Virtual I/O Concepts

Virtual Ethernet

- Virtual ethernet adapter provides connections between the partition and virtual networks
- Purpose
  - Provides high-speed connections between partitions
    - 1 Gbps ethernet
  - Provides secure connections between partitions
    - A virtual network cannot be "tapped" or "sniffed"
- Parameters
  - Slot number
  - Network number : up to 4096 virtual networks !
- Who can participate to a Virtual LAN ?
  - Any i5/OS partition
  - Any Linux partition
  - Any AIX 5L v5.3 partition

Virtual I/O Concepts

Virtual Ethernet

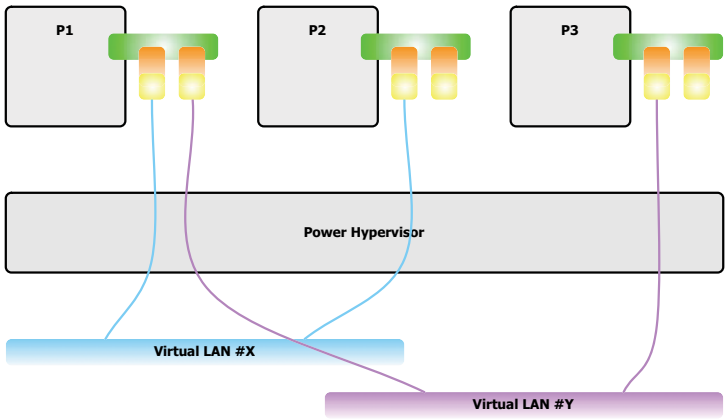

Virtual I/O Concepts

Virtual SCSI

- Virtual SCSI adapter provides SCSI storage services connection between partitions
- Purpose
  - → Storage virtualization
- Parameters
  - → Slot number
  - → Adapter type : server or client
  - → Connection information for the server : who can connect ...
    - Any remote partition
    - Selected remote partition
  - → Connection information for the client : where to connect ...
    - Which partition
    - Which slot
- Who can be SCSI server ?
  - → Any i5/OS partition
- Who can be SCSI client
  - → Any Linux partition
  - → Any AIX 5L v5.3 partition

Virtual I/O Concepts

Virtual SCSI

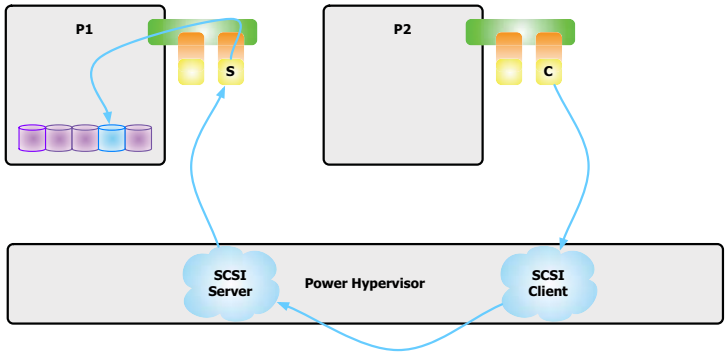

Virtual I/O Concepts

Virtual OptiConnect

- What is OptiConnect ?
  - → Stands for OPTImized CONNECTion
  - → Provides very high performance communications between systems
    - Works at HSL bus speed : 2 GB/s (20 Gbps) !
  - → Also known as "HSL OptiConnect"
  - → Requires hardware planning
    - Multiple CECs within the same HSL loop
  - → Requires additional software
    - 5722SS1, Option XX, OptiConnect
- What is "virtual OptiConnect"
  - → Virtualization of HSL OptiConnect facility between logical partitions
  - → Requires additional software
    - 5722SS1, Option XX, OptiConnect
- OptiConnect and Virtual OptiConnect usages :
  - → DDM
  - → Clusters
  - → Switchable IASPs

iSeries. mySeries.

Logical Partitions on the IBM Power5

Dynamic LPAR !

Dynamic LPAR

What is called "Dynamic Logical Partitioning"

- The ability to add, move or remove resources without system disruption

Dynamic LPAR

What can be changed dynamically !

- Processors
  - → Within established minimum and maximum
    - Configuration maintained in the partition profile
    - Changes to minimum or maximum requires partition deactivation an reactivation
  - → Granularity of moves
    - Dedicated processors : 1 processor
    - Shared processors : 0.01 processing unit (PU) per virtual processor !
    - Virtual processors : 1 virtual processor
- A processor move implies :
  - → Deallocation of processor or processor time
  - → Clear caches and registers
  - → Allocation of processor or processor time
- Nearly immediate !

Dynamic LPAR

What can be changed dynamically !

- Memory
  - → Within established minimum and maximum
    - Configuration maintained in the partition profile
    - Changes to minimum or maximum requires partition deactivation an reactivation
  - → Granularity of moves
    - The size of a Logical Memory Block (LMB) : 16 Megabytes
    - In the future, those Moves
- A memory move implies :
  - → Deallocation of memory
    - From the *BASE pool
    - Memory pages must be written to disk before deallocation
  - → Clear of memory
  - → Allocation of memory to the new partition
    - To the *BASE pool
- Time required !
  - → Up to 5 minutes !

Dynamic LPAR

What can be changed dynamically !

- I/O Slots
  - → Only "desired" slots can be moved !
  - → Remember i5/OS still needs IOPs
    - An IOP without IOA is functional but not very useful
    - An IOA without IOP is nothing
    - BUS/IOP/IOA requirements must be kept !
  - → Remember "IOP-less" IOAs
    - Ethernet adapters
    - Cryptographic processors
    - ... more in the future : the direction is to get away from the IOP !

Dynamic LPAR

What can be changed dynamically !

- An I/O slot move implies :
  - → Deactivation of the dependent device
    - VARY OFF !
  - → Deallocation of the slot
    - IOP/IOA gets first the "inoperative", then "not detected" status
  - → Reset of the IOP/IOA cache, memory, registers
  - → Allocation of the slot to the destination partition
    - IOP IPL : download of its microcode, activation of the IOP
    - IOA initialization : download of its microcode, activation of the IOA
    - IOP/IOA gets the "operational" status ... some "failed" status may temporarily appears !
  - → Activation of the dependent devices
    - VARY ON !
- Time required !
  - → Up to 10 minutes !

Dynamic LPAR

What can be changed dynamically !

- Virtual I/O Slots
  - → Can only be filled or emptied
    - Cannot be moved from one partition to another
  - → Only "desired" slots can be moved !
    - Configuration maintained in the partition profile
    - Changes to the desired/required status of a slot requires partition deactivation an reactivation
    - Changes to the number of virtual slots requires partition deactivation and reactivation
  - → Granularity of moves
    - One slot

iSeries. mySeries.

Logical Partitions on the IBM Power5

LPAR-Capable Server Models

IBM eServer i5/520

Model i5/520 - 1000 CPW

- Processor : 1-Way Power5, 1.50GHz, L1+L2 Cache
- Processing units available : 0.43
- Maximum partitions : 4

Model i5/520 - 2400 CPW

- Processor : 1-Way Power5, 1.50GHz, L1+L2 Cache
- Processing units available : 1.00
- Maximum partitions : 10

Model i5/520 - 3300 CPW

- Processor : 1-Way Power5, 1.65GHz, L1+L2+L3 Cache
- Processing units available : 1.00
- Maximum partitions : 10

Model i5/520 - 6000 CPW

- Processor : 2-Ways Power5, 1.65GHz, L1+L2+L3 Cache
- Processing units available : 2.00
- Maximum partitions : 20

IBM eServer i5/550

Model i5/550 - CoD from 3300 to 12000 CPW

- Minimal Configuration
  - → Processor : 2-Ways Power5, 1.65GHz, L1+L2+L3 Cache
    - ! one i5/OS license !
  - → Processing units available : 2.0
  - → Maximum partitions : 20
- Maximal Configuration
  - → Processor : 4-Ways Power5, 1.65GHz, L1+L2+L3 Cache
  - → Processing units available : 4.0
  - → Maximum partitions : 40

IBM eServer i5/570

Model i5/570 - CoD from 3300 to 44700 CPW

- Minimal Configuration
  - Processor : 1-Ways Power5, 1.65GHz, L1+L2+L3 Cache
  - Processing units available : 1.0
  - Maximum partitions : 10
- Maximal Configuration
  - Processor : 16-Ways Power5, 1.65GHz, L1+L2+L3 Cache
  - Processing units available : 16.0
  - Maximum partitions : 160

IBM eServer i5/595

Model i5/595 - CoD from 24500 to 165000 CPW

- Minimal Configuration
  - Processor : 8-Ways Power5, 1.65GHz, L1+L2+L3 Cache
  - Processing units available : 8.0
  - Maximum partitions : 80
- Maximal Configuration
  - Processor : 64-Ways Power5, 1.65GHz, L1+L2+L3 Cache
  - Processing units available : 64.0
  - Maximum partitions : 254

iSeries. mySeries.

Logical Partitions on the IBM Power5

Planning

Planning

Think ...

- Understand eServer i5 capabilities
  - → Review previous chapters ...
  - → System Builder !
- Understand Operating System requirements
  - → i5/OS
  - → Linux
  - → AIX

Planning

Requirements for i5/OS

- Are mandatory :
  - At least 0.10 processing unit
    - Commonly represents +/- 330 CPW
  - At least 256 MB of main storage
    - Just enough to load and start SLIC & i5/OS ... plan a few bit more to open a session !
  - One direct-attach (internal) disk unit (LSU)
    - Minimum size : 8.58 GB
  - One console
    - Can be a HMC thru virtual serial port
  - One alternate-IPL device
    - DVD or tape unit ... can be switchable or dedicated
  - One backup device
    - A tape unit ... can be switchable or dedicated
- Operating system :
  - i5/OS V5R3M0
    - DO NOT USE SLIC RSA !

Planning

Requirements for i5/OS

- Plan also for :
  - LAN ports
  - WAN ports
  - Alternate console
  - SCSI ports
  - FC ports
  - VIRTUAL ports !
- Plan also for :
  - IXS/IXA
    - More storage needed ?
    - Same ASP, other ASP or IASP ?
  - Other partitions : Linux or AIX thru virtual I/O
    - More storage needed ?
    - Same ASP, other ASP or IASP ?

Planning

About Linux and AIX ...

- CPW
  - → Commercial Processing Workload
  - → Representative for a commercial workload environment
    - Workload composed by CPU and I/O operations
- rPerf
  - → Relative Performance
  - → Derived from multiple "standard" benchmarks
    - TPC : Technical Publication Center
    - SPEC : System Performance Evaluation Center
    - Other IBM internal workloads
  - → Does not simulate I/O operations

Planning

About Linux and AIX ...

eServer model	Number of CPUs	GHz	rPerf	CPW
520	2	1.65	9.86	6000
550	2	1.65	9.86	6200
	4	1.65	19.66	12000
570	2	1.65	9.86	6350
	4	1.65	19.66	12000
	8	1.65	37.22	23500
	12	1.65	53.43	33400
	16	1.65	68.40	44700

Planning

Requirements for Linux/AIX 5L v5.3

- Are mandatory :
  - → At least 0.10 processing unit
  - → Some memory ... depending on the applications you'll have to run !
  - → One disk unit ...
    - Physical or ... virtual !
  - → One console
    - Can be a HMC thru virtual serial port
  - → One alternate boot device
    - CD or DVD ... physical or virtual !
  - → One backup device
    - A tape unit ... can be switchable, dedicated or virtual !
- Operating system :
  - → A PowerPC compatible, 64-bits kernel Linux edition :
    - Red Hat Enterprise Linux AS for POWER v4
    - Novell Suse Linux Enterprise Server for POWER v9
  - → AIX 5L v5.3

Planning

Requirements for Linux/AIX 5L v5.3

- Plan also for :
  - → LAN ports
  - → WAN ports
  - → SCSI ports
  - → FC ports
  - → USB ports
  - → IDE ports
  - → VIRTUAL ports !

Planning

Requirements for AIX 5L v5.2

- REMEMBER : this version of AIX does NOT SUPPORT "virtual things" !
  - → Virtual processors cannot be used ...
  - → Memory cannot be dynamically assigned ...
  - → No virtual networks, nor virtual storage can be used
  - → Virtual console can be used !
- Are mandatory :
  - → At least 1 dedicated processor
  - → Some memory ... depending on the applications you'll have to run !
  - → One disk unit
    - One console
    - Can be a HMC thru virtual serial port
  - → One alternate boot device
    - CD or DVD
  - → One backup device
    - A tape unit ... can be switchable or dedicated

Planning

Requirements for AIX 5L v5.2

- REMEMBER : this version of AIX does NOT SUPPORT virtual "things" !
  - → Virtual processors cannot be used ...
  - → Memory cannot be dynamically assigned ...
  - → No virtual networks, nor virtual storage can be used
  - → Virtual console can be used !
- Plan also for :
  - → LAN ports
  - → WAN ports
  - → SCSI ports
  - → FC ports
  - → USB ports
  - → IDE ports

Planning

Know restrictions

- At this time, Linux/AIX partitions does not handle any IOA with an IOP placed in an upstream position on the same bus !
  - → This means that you cannot share a bus between a Linux/AIX partition and an i5/OS partition
  - → This means that you cannot switch a device between a Linux/AIX partition and an i5/OS partition
    - You can switch a device between i5/OS partitions
    - You can switch a device between AIX/Linux partitions

Planning

Know restrictions

- Disk formats ...
  - → i5/OS uses 540 bytes/sector format
    - Sector data + CRC + TAG bits !!!
  - → Following units comes with i5/OS format :
    - #4319 : 35.16 GB
    - #4326 : 35.16 GB
    - #4327 : 70.56 GB
    - #4328 : 140.12 GB
  - → AIX uses
    - 522 bytes/sector format for RAID-5 protected units
    - 512 bytes/sector format for all others
  - → AIX won't recognize disks with 540 bytes/sector format
    - Those units will need a low-level formatting !

Planning

... Think ...

- Make a capacity planning
  - → Determine your needs
- Inventory existing system
  - → Determine which component may/will be reused
- Build a new system
- Validate the new system
  - → Use LPAR Validater Tool
  - → Call IBM or a Business Partner
- Consider to upgrade existing system to supported OS version/release
  - → Consider EVERY partition !
  - → Upgrades to i5/OS V5R3 ...
    - V5R3 does NOT support models 150, 6xx, Sxx
    - V5R3 is the last release to support models 170, 250, 7xx
    - V5R3 is the last release to support SPD hardware
    - V5R1 planned end of support is September 2005

Planning

... Think ...

- Order your new system
  - → ...
- Upgrade existing system(s) to supported OS version/release
  - → Consider EVERY partition !
- Write procedures for installation/upgrade
  - → Disk moves or unload/reload ?!
  - → Side-by-side !?

Planning

... Then act !

- Backup existing system
  - → Twice !
- Perform physical installation
  - → Install every hardware piece at its planned position
  - → Install and configure HMC
    - Setup wizard !
  - → Power on the new system to "Hypervisor Stand-by" mode
  - → Create LPAR definitions
- Start or reload
  - → If disk were moved ... just start !
  - → Else ... just reload !
- Validate your new system ... and ... enjoy it !

iSeries. mySeries.

Logical Partitions on the IBM Power5

Hardware Management Console

Hardware Management Console

What is it ?

- Hardware dedicated to console functions
- Used to create and maintain a multiple-partitioned environment
  - → Startup & shutdown partitions
  - → Performing resource movements
- Displaying a virtual console
- Displaying a virtual operator panel
- Detecting, reporting and sorting change in hardware conditions
- Gathering and reporting system error events
- Activating CoD resources
- Supports i5/OS, Linux and AIX

Is it mandatory ?

- YES, for LPARed systems
- YES, for CoD systems
- Not required to operate partitions

Hardware Management Console

Appliance based on xSeries technology

- Minimum configuration includes
  - → Intel Pentium-4 processor
  - → RAM : 1 GB
  - → HDD : 40 GB
  - → LAN : 2x 1000/100/10 Mbps ethernet
  - → DVD-RAM : 1x
  - → FDD : 1x
- "Cleaned" Linux with HMC application (GUI)

Connects to i5 and p5 servers

- Using dedicated HMC ports
- Up to 64 partitions on 16 servers

Packaging

- Desktop version : 7310-C03
- Rack version : 7310-CR3

iSeries. mySeries.

Logical Partitions on the IBM Power5

Hardware Management Console
Connecting HMC

Hardware Management Console

Ethernet Ports

- Private LAN
- Open LAN
- Auto-MDIX ports
  - → Medium Dependent Interface Crossover

Hardware Management Console

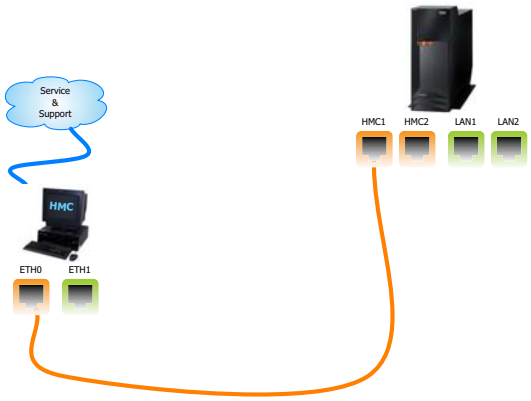

Simple HMC connection

Hardware Management Console

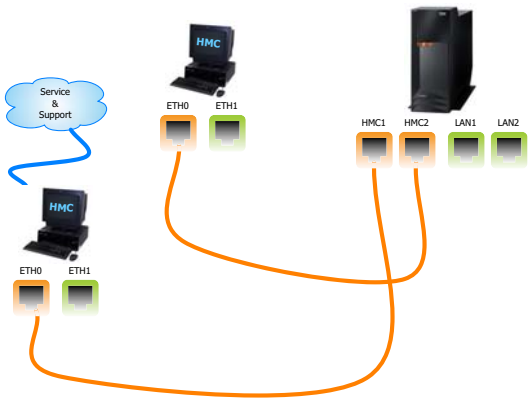

Cloning HMC configuration

- System profiles and partition profiles are stored in SP's NVRAM
- When SP detects that a redundant HMC is connected :
  - → SP copies NVRAM configuration to the new HMC

Hardware Management Console

Redundant HMC configuration considerations

- Both HMCs can be used concurrently
  - → Both are fully active and accessible at all times
- Both enabling you to perform management tasks at any time
  - → SP provides a lock-mechanism to allow operations in a parallel environment
  - → When one HMC issues an order, the other HMC is temporarily "locked-out"
- There is no "primary" or "backup" designation

But ...

- Other definitions are locally stored and must be set up separately on each HMC
  - → User authorizations
  - → Network settings
  - → Service Settings
- HMC provides Service Focal Point and Service Agents functions
  - → Enable service on only one HMC to avoid duplicate service calls
- Perform software maintenance separately on each HMC
  - → Ensure HMC availability when updating code level
  - → Meanwhile, update other HMC code to same level as soon as possible
- Consider save/restore critical data to duplicate configuration information

Hardware Management Console

Private indirect networking

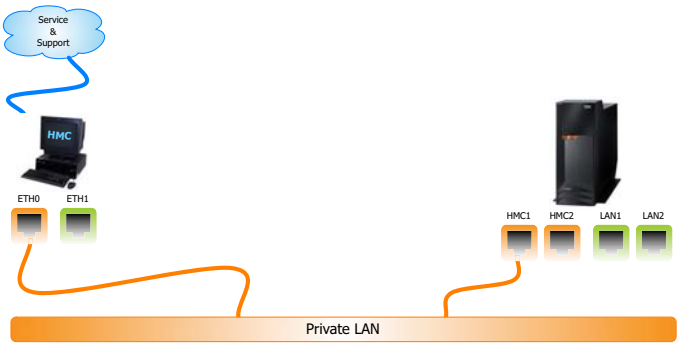

Hardware Management Console

Private indirect networking

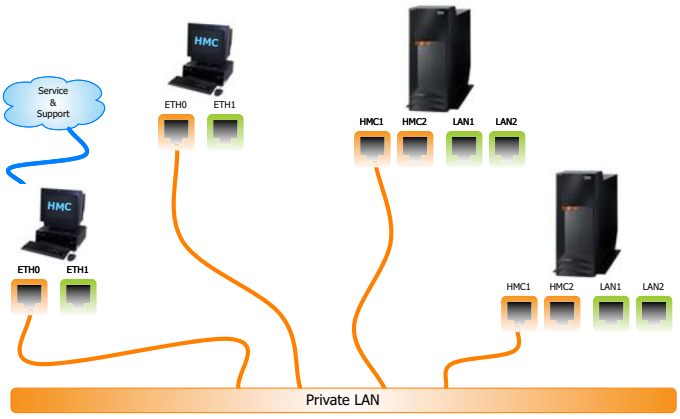

Hardware Management Console

Private and open networking

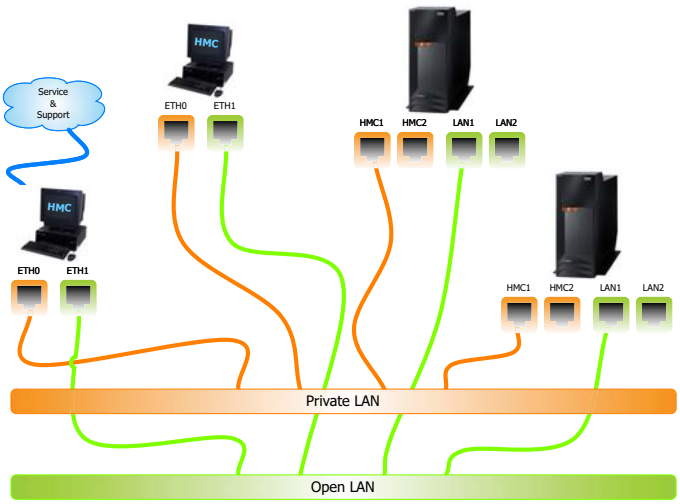

Hardware Management Console

Working with HMC in open networks

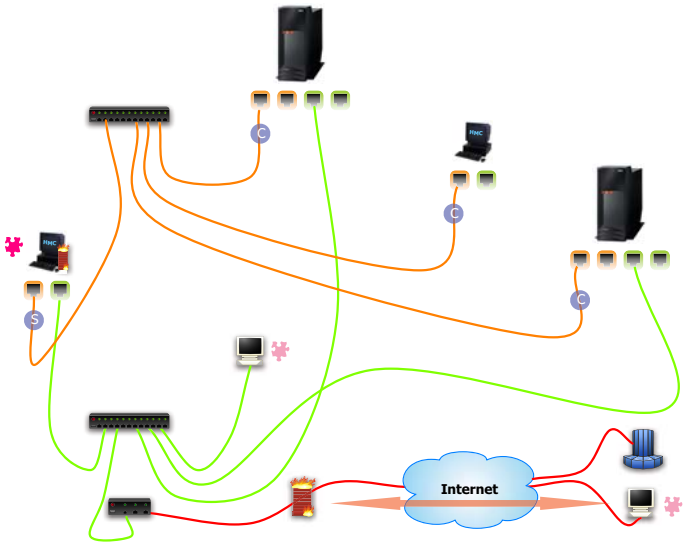

iSeries. mySeries.

Logical Partitions on the IBM Power5

Hardware Management Console
HMC Functionalities

Hardware Management Console

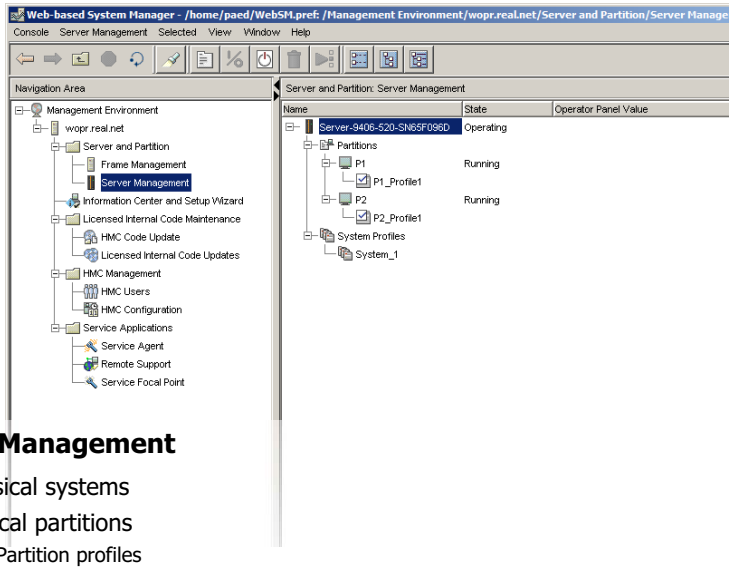

Server Management

- Physical systems
- Logical partitions
  - → Partition profiles
- System profiles

Hardware Management Console

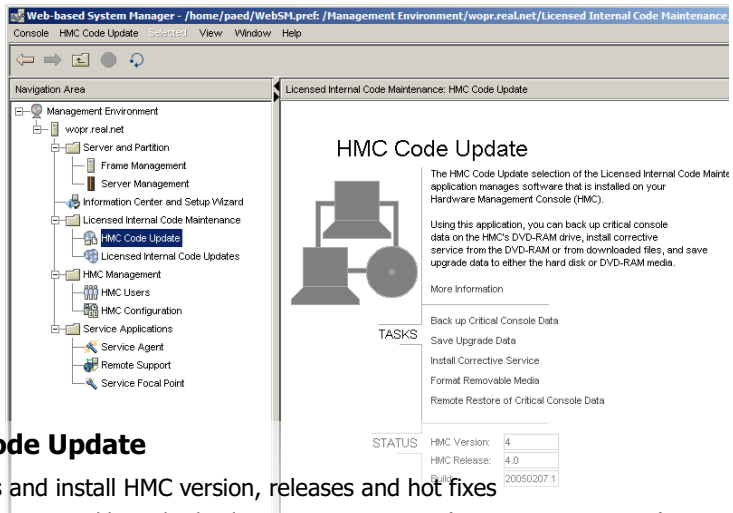

HMC Code Update

- Gets and install HMC version, releases and hot fixes
  - → Via removable media, local FTP server or internet (IBM service FTP server)
- Backups
  - → Critical console data vs. upgrade data

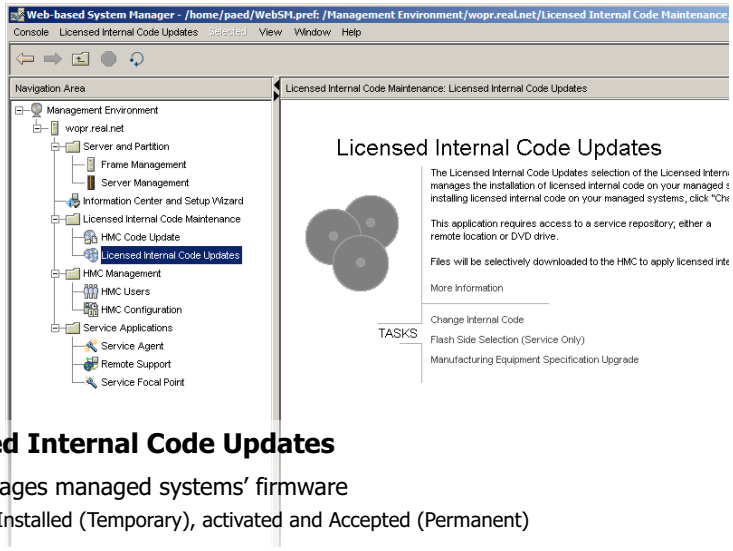

Licensed Internal Code Updates

- Manages managed systems' firmware
  - → Installed (Temporary), activated and Accepted (Permanent)

Hardware Management Console

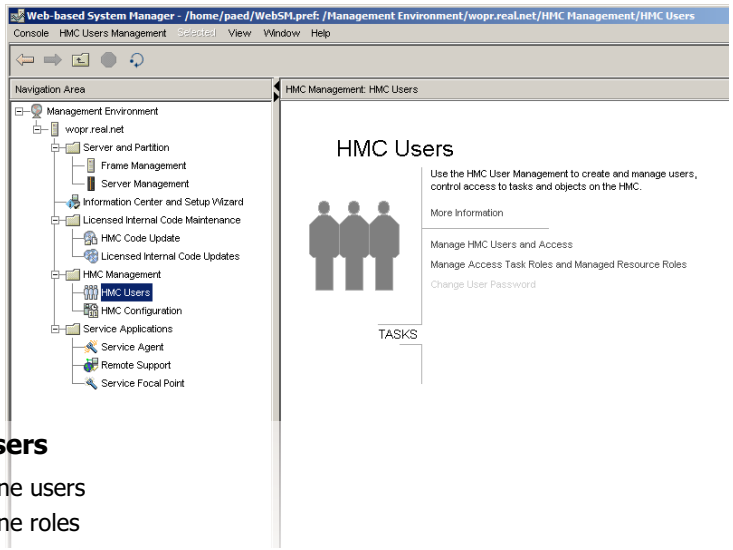

HMC Users

- Define users
- Define roles

Hardware Management Console

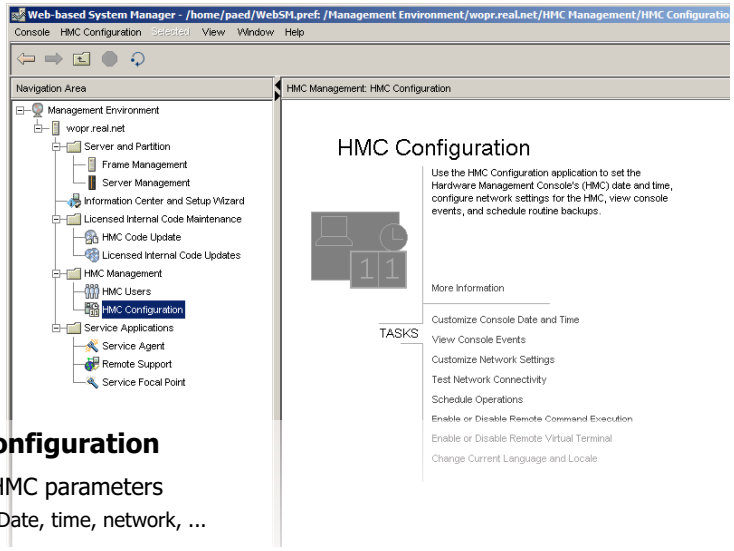

HMC Configuration

- All HMC parameters
  - → Date, time, network, ...

Hardware Management Console

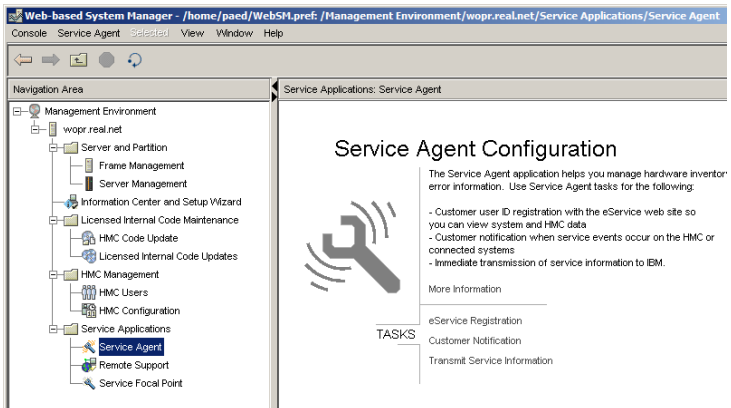

Service Agent Configuration

- Notify problems
  - → Kind of problems : "all" or "home-call"
  - → To the user : via SMTP
  - → To the service : via PSTN, VPN or another HMC

Hardware Management Console

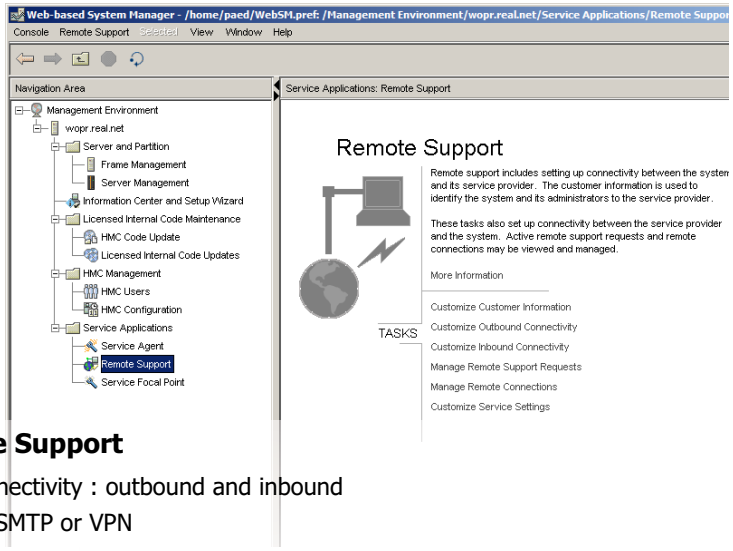

Remote Support

- Connectivity : outbound and inbound
- Via SMTP or VPN

Hardware Management Console

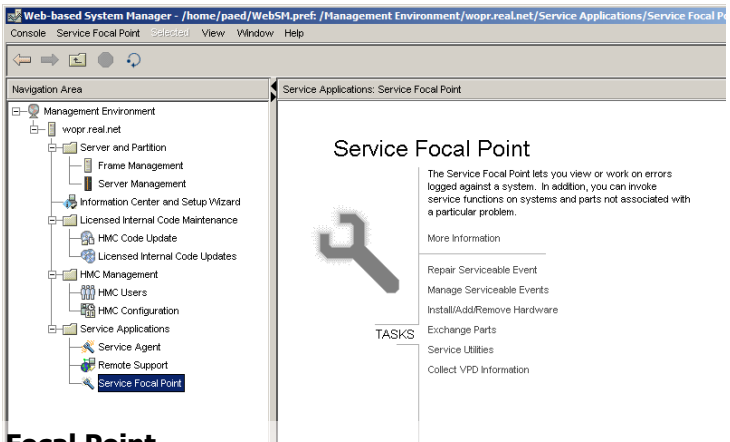

Service Focal Point

- Manage reported events and errors
- Concurrent maintenance
- Service utilities
  - → Access to control panel service functions (e.g. Function 21 - Request DST)

One more thing ...

iSeries. mySeries.

Logical Partitions on the IBM Power5

Virtual Partition Manager

Virtual Partition Manager

What is it ?

- VPM introduces a way to create and manage Linux partitions without the use of a HMC
- VPM allows small and medium configuration to add simple Linux workloads

Available on May 27th, 2005 (tomorrow) !

Virtual Partition Manager

Planning for VPM ...

- An IPL is requires to apply enablement PTF
  - → MF34753, MF34956, MF34962
  - → Those PTF should be permanently applied before SAVSYS
- Ensure that you have the last firmware level
  - → Minimum level is SF230_113, via PTF MH00265
- Maximum 4 Linux partitions
- I/O for all linux partition must be managed by a single i5/OS instance
  - → Only a single i5/OS instance can exist
- Linux must use virtual I/O to access disk, tape, DVD, ethernet
  - → Direct-attach I/O cannot be installed in any Linux partition
- Maximum 4 virtual ethernet connection per Linux partition
- Linux partition creation and maintenance can only be performed thru DST/SST
  - → iSeries Navigator does not support VPM
  - → LPAR Toolkit does not support VPM
  - → Likewise, there is no API nor SSH enablement capabilities

Virtual Partition Manager

Planning for VPM ...

- Automatic processor balancing between i5/OS and Linux partition is supported
  - → Thru uncapped shared processor
- Dynamic movements of resources (processor, memory and I/O) is not supported
  - → Partition must be restarted for the changes to take effects
- Capacity on Demand (CoD) is not enabled thru DST/SST
  - → Only permanently activated processors can be allocated
- Partition configuration cannot be saved
  - → Hardcopy prints must be kept to re-create partitions
- Migration to HMC is not available
  - → If a HMC is deployed in a future stage, Linux partitions must be re-created
  - → Data stored thru virtual I/O on i5/OS is preserved
- Configurations screens are available only in english
  - → No translation

Virtual Partition Manager

Virtual Partition Manager vs. Hardware Management Console

	VPM	HMC
Operating systems supported	i5/OS and Linux	i5/OS, Linux and AIX
Maximum number of partitions	5 (1x i5/OS, 4x Linux)	254
Uncapped partition support	Yes	Yes
Dynamic resource movement	No	Yes
I/O support for Linux	Virtual	Virtual and Direct
Maximum number of ethernet connections	4	4096
Maximum virtual disk per partition	64 TB	64 TB

Questions & Answers

iSeries. mySeries.

Logical Partitions on the IBML Power5

Bibliography

Bibliography

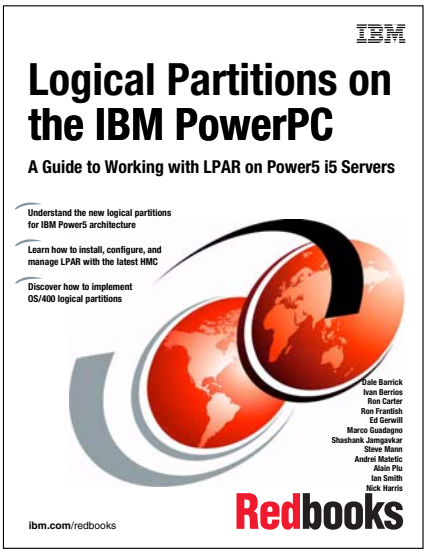

Bibliography

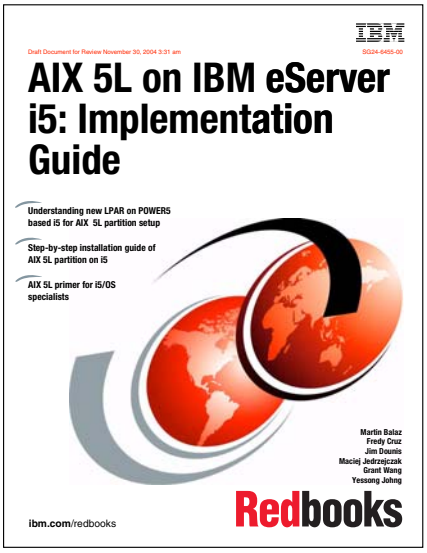

Bibliography

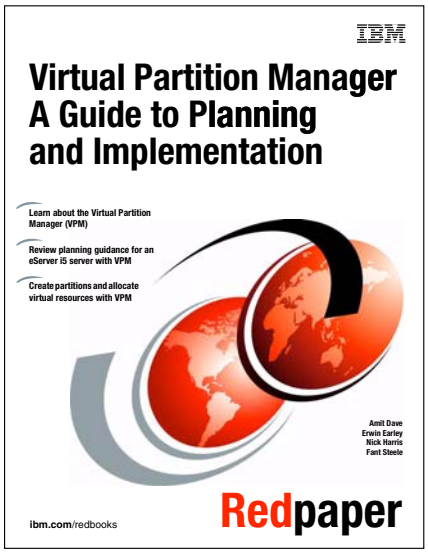

Bibliography

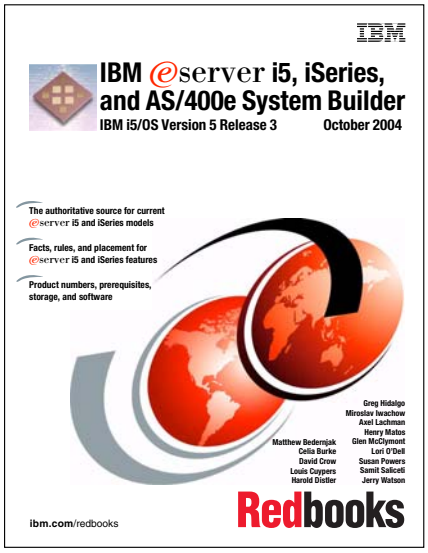

Thank You !

Contact Information

REAL Solutions S.A.
Rue d'Eich, 33
L-1461 Luxembourg

Eddy Pasteger
System Engineer Manager
iSeries Certified Solutions Expert

Voice : (+352) 43 65 22 1
Fax : (+352) 42 26 38
E-mail : eddy.pasteger@real.lu