

iSeries Dynamic Logical Partitioning

Simplicity in an on demand world

@ddy Pasteger / REAL Solutions

iSeries Dynamic Logical Partitioning

Objectives

- Understand the LPAR functionality in OS/400 V5R2
- Practical Scenarios that use new functionality in OS/400 V5R2
- Planning considerations and system requirements

Warning

- This presentation is build on functionalities available on iSeries model 8xx and OS/400 V5R2

@ddy Pasteger / REAL Solutions

Agenda

- Chapter 1. The iSeries Hierarchy of Microprocessors
- Chapter 2. The LPAR Concepts
- Chapter 3. The LPAR Benefits
- Chapter 4. Build an LPARed System by Example
- Chapter 5. Linux on iSeries LPAR

@ddy Pasteger / REAL Solutions

iSeries Dynamic Logical Partitioning

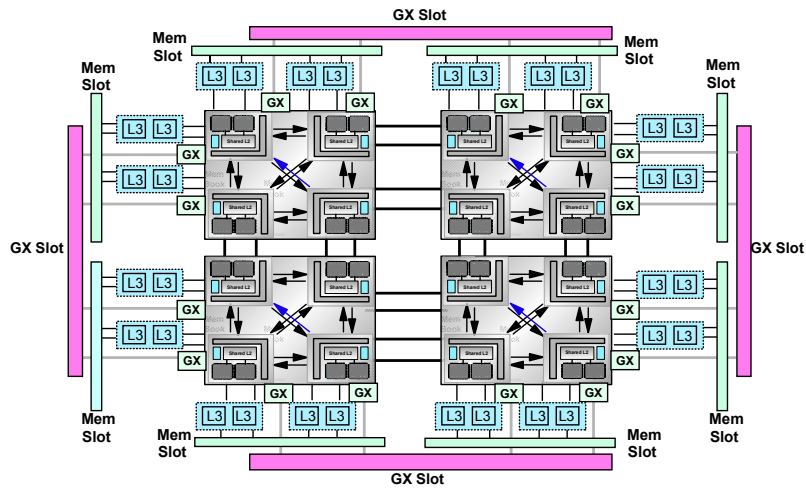
Chapter 1. iSeries Hierarchy of Microprocessors

"High performance on an iSeries server is achieved by using many individual high performance microprocessors, I/O devices, and interconnect technologies"

@ddy Pasteger / REAL Solutions

iSeries Hierarchy of Microprocessors

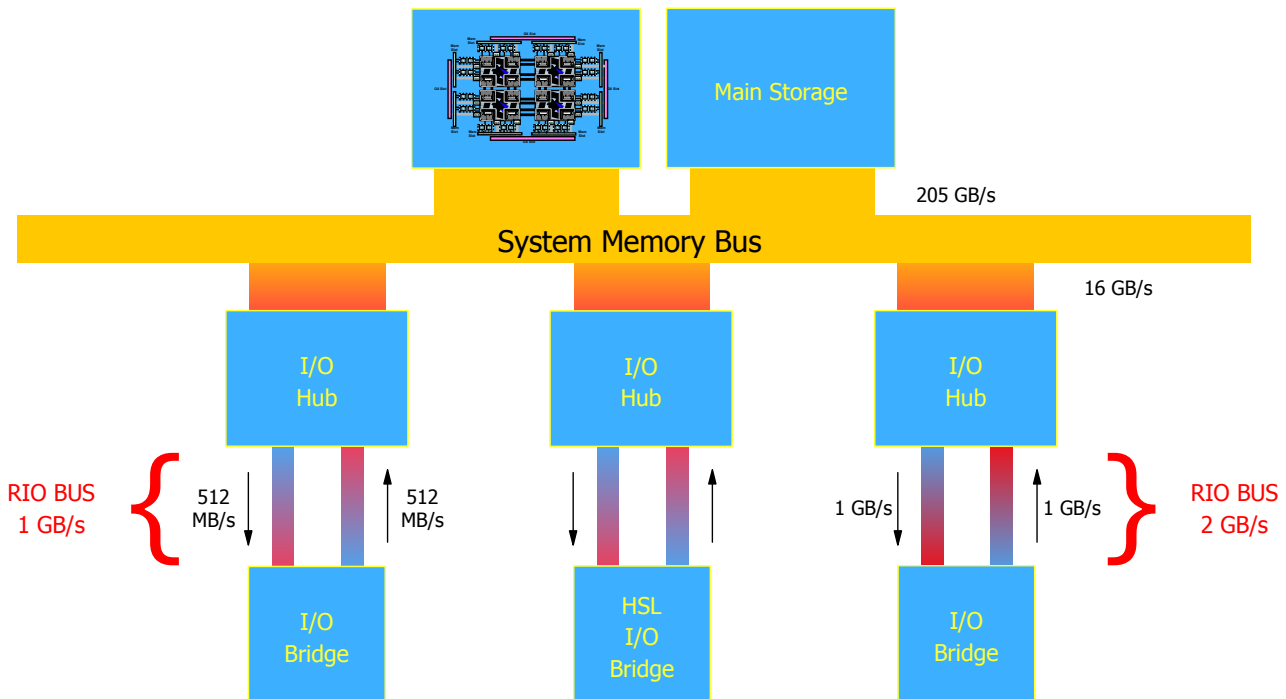
The processor complex



@ddy Pasteger / REAL Solutions

iSeries Hierarchy of Microprocessors

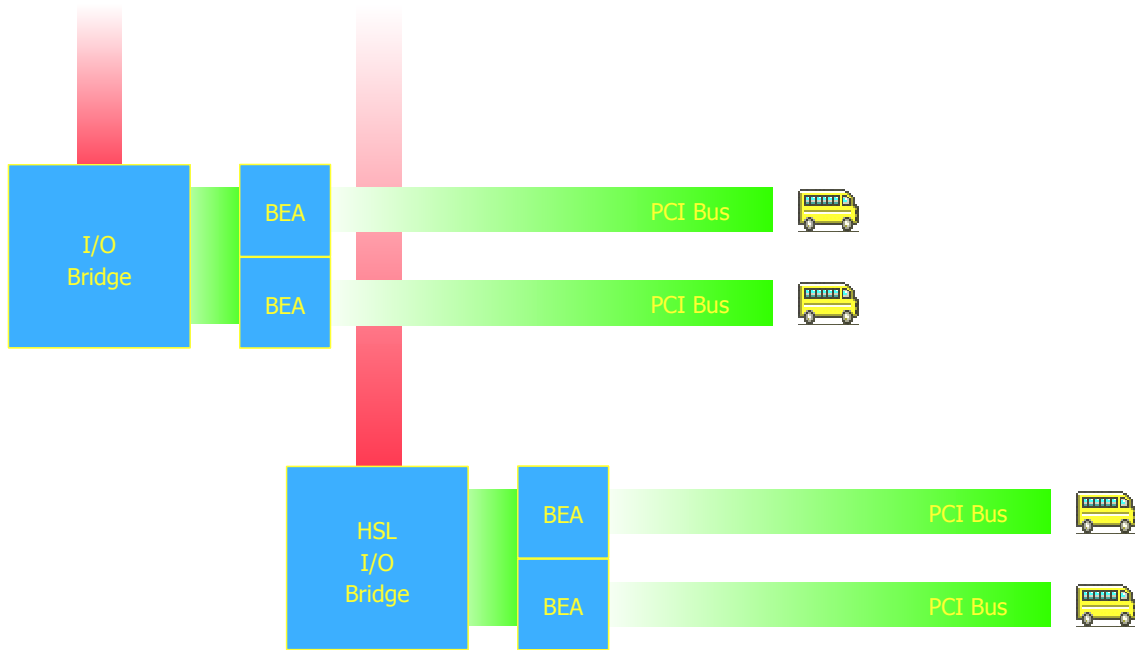
I/O Hubs and I/O Bridges



@ddy Pasteger / REAL Solutions

iSeries Hierarchy of Microprocessors


The PCI Buses



@ddy Pasteger / REAL Solutions

iSeries Hierarchy of Microprocessors

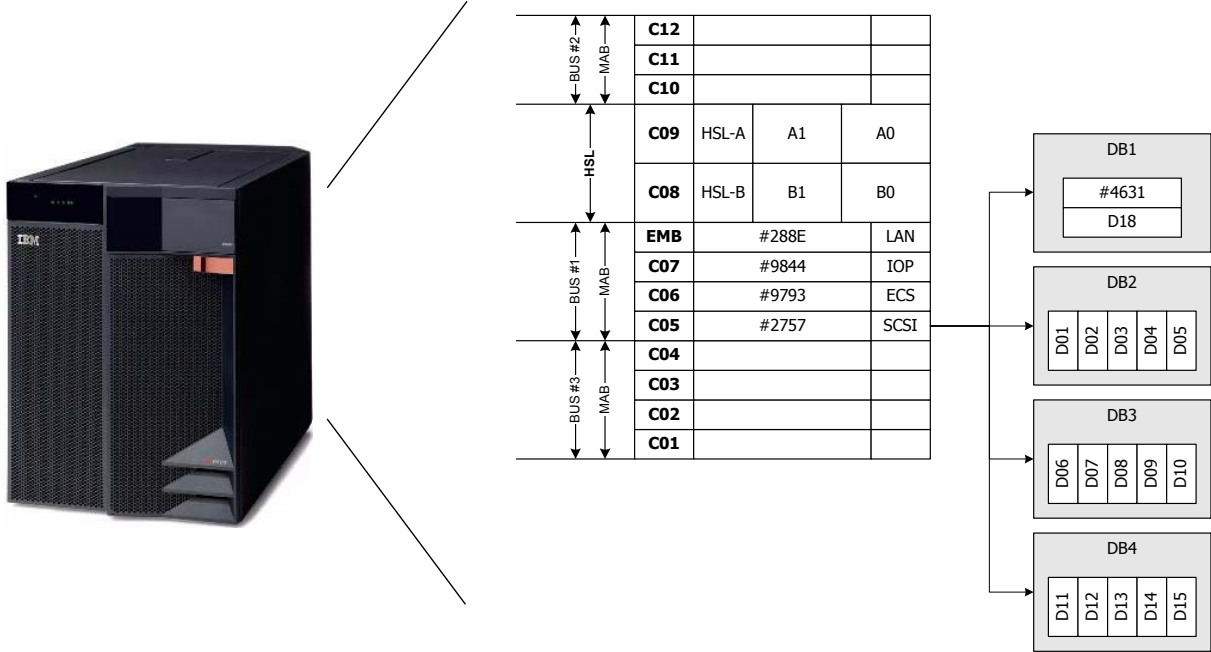
MABs, IOPs and IOAs



@ddy Pasteger / REAL Solutions

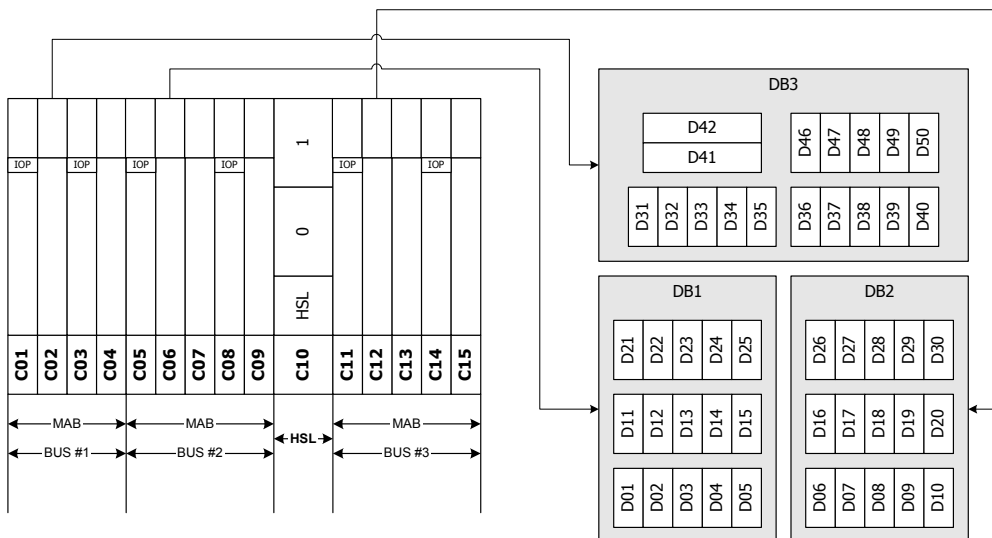
iSeries Hierarchy of Microprocessors

Example : an iSeries model 825 system



iSeries Hierarchy of Microprocessors

Example : a #5094 PCI-X Expansion Tower



Chapter 2. The LPAR Concepts

The LPAR Concepts

Logical Partitioning

- The ability to make a single iSeries Server run as if it were many independent systems
- Each Logical System is called a "Partition"
- Each Partition operates as an independent system

The LPAR Concepts

Primary Partition

- Starting V4R4, every OS/400 is configured with a "Primary Partition"
- The Primary Partition initially owns all the resources available on the machine
 - Processors
 - Main Storage
 - Buses
 - IOPs
- The Primary Partition function as one of the logical system
- The Primary Partition provides functions on which all other partition are dependent
 - Power Management
 - Virtual Operations Panel
 - Logical Partition Definition
 - Integrated Hypervisor
- The Primary Partition is a single point of failure

The LPAR Concepts

Secondary Partition

- Secondary Partitions are created and managed from the Primary Partition
- Secondary Partitions function as independent systems
- Secondary Partitions have their own hardware
 - Processors
 - Main Storage
 - Buses
 - IOPs
- Secondary Partitions have their own software
 - SLIC, OS/400, LPPs, PTFs
 - Primary/Secondary languages
 - System Values, including time-of-day
 - User Profiles
- Secondary Partitions can independently :
 - Power On/Off
 - Dump Main Storage

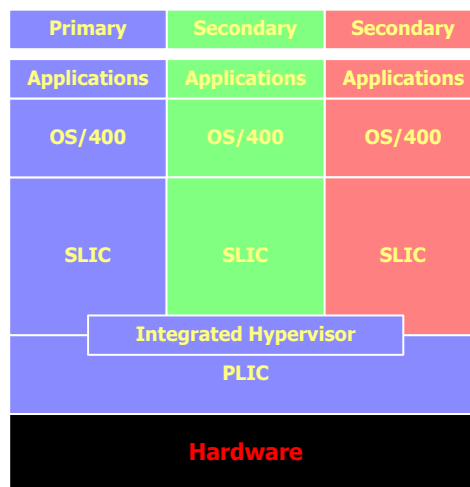
The LPAR Concepts

Integrated Hypervisor

- **PLIC**
 - Allocating resources to a Partition
 - Installing an Operating System in a Partition
 - Starting and Stopping the Operating System in a Partition
 - Dumping Main Storage of a Partition
 - Communications between Partitions
- **SLIC**
 - Main Storage management
 - Task management
 - Heap management
 - I/O Subsystem management

The LPAR Concepts

How Does It Work ?



The LPAR 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
- Rules & Requirements :
 - A least one dedicated processor
 - Granularity of movements : one processor
 - Processor moves are dynamic, without any IPL while within configured ranges
 - Unallocated processors are not used




@ddy Pasteger / REAL Solutions

The LPAR Concepts

Shared Processors

- A *Shared Processor* allows to assign partial processors to a partition
- Physical processors are assigned in a "Shared Processor Pool"
- Rules & Requirements :
 - A least 0.10 processor
 - Granularity of movements : 0.01 processor
 - Processor moves are dynamic, without any IPL while within configured ranges
 - Unallocated processors or partial processors are not used



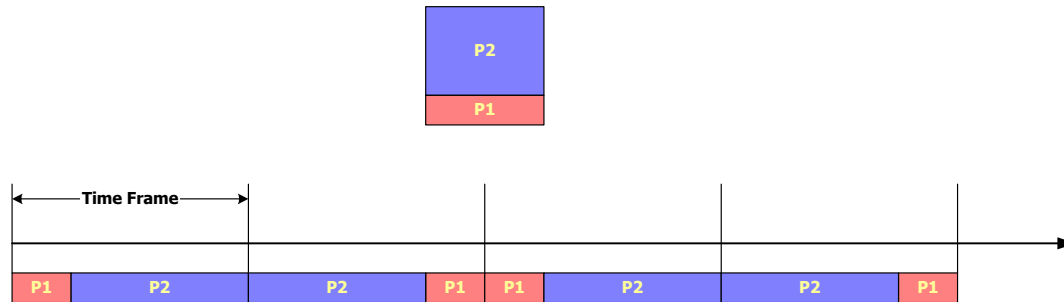
@ddy Pasteger / REAL Solutions

The LPAR Concepts

Shared Processors

- Example : Single CPU iSeries System with two partitions

- Partition #1 receive 0.25 processor
- Partition #2 receive 0.75 processor



The LPAR Concepts


Virtual Processors

- The number of *Virtual Processors* represent the number of parallel threads of execution
- The processing power is spread equally across its virtual processors
- Rules & Requirements
 - The number of virtual processors must be less or equal than the number of processors in the shared pool
 - Changes are dynamic, without any IPL while within configured ranges

The LPAR Concepts

Virtual Processors

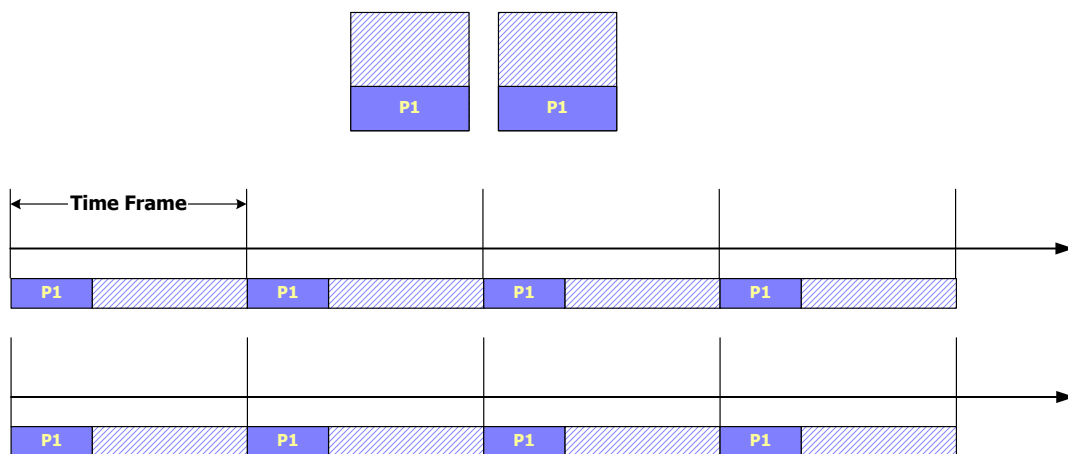
- Example : 0.75 processing unit running on 1 virtual processor



The LPAR Concepts

Virtual Processors

- Example : 0.75 processing unit running on 2 virtual processors



Interactive Capacity

- Determine the percentage available for each partition
- Rules & Requirements
 - Minimum : 0%
 - Granularity : 1%
 - Maximum : depends on allocated processing power
 - Changes are dynamic, without any IPL while within configured ranges
 - Unallocated interactive capacity is not used

Main Storage

- Processors use memory to temporarily hold information
- Memory is added into *BASE pool
- Memory is removed from *BASE pool
 - The system will keep the minimum required (determined by system value QBASPOOL)
 - The system will "flush" any data from memory pages to disk before making them available
- Rules & Requirements
 - A partition must have a whole number of Megabytes (1 MB = 1024 x 1024 bytes)
 - Primary partition needs a minimum of 256 MB
 - Secondary partitions needs a minimum of 128 MB
 - Granularity of movements : 1 MB (256 pages of 4 KB)
 - Memory moves are dynamic, without any IPL while within configured ranges
 - Unallocated memory is not used

The LPAR Concepts

I/O Partitioning

- BUS-Level
- IOP-Level

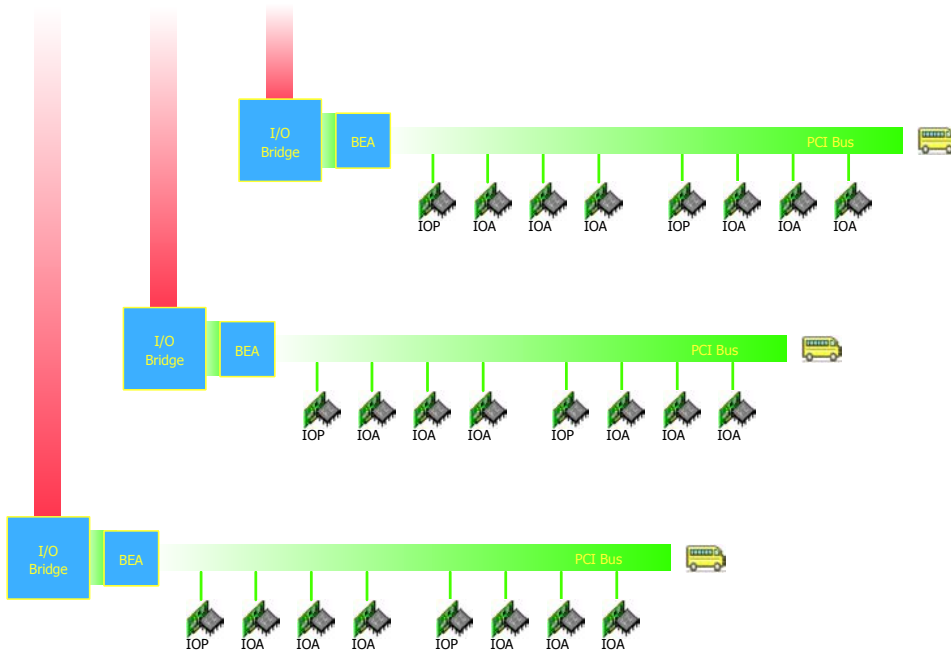
The LPAR Concepts

BUS-Level I/O Partitioning

- A bus and all attached IOPs and devices are dedicated to a single partition
- BUS-Level Advantages
 - Better problem isolation
 - Better availability
 - Better performance
 - More simple hardware management
- BUS-Level Disadvantages
 - Lot of additional hardware required
 - Resources cannot be shared between partitions

The LPAR Concepts

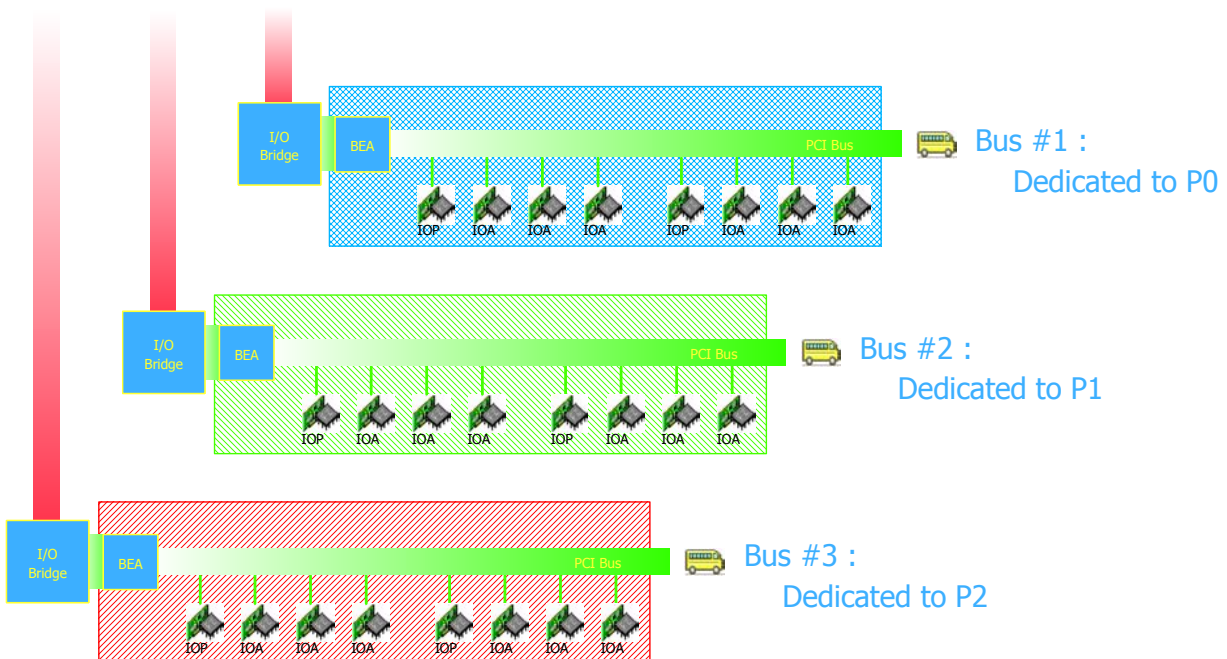
BUS-Level I/O Partitioning



@ddy Pasteger / REAL Solutions

The LPAR Concepts

BUS-Level I/O Partitioning



@ddy Pasteger / REAL Solutions

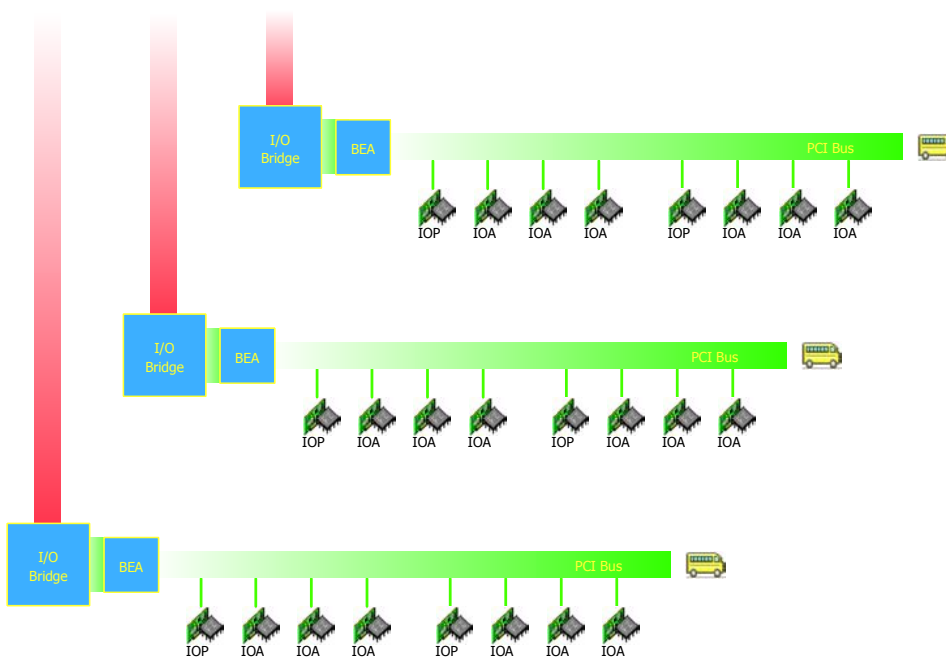
The LPAR Concepts

IOP-Level I/O Partitioning

- Share the bus and divide resources by IOP
- Bus will be defined as :
 - DEDICATED :
 - One partition uses all IOPs and IOAs resources
 - No other partition can use any resource on that bus
 - OWN SHARED :
 - One partition owns the bus and its hardware resources
 - The owner partition will allow selective use of resources by another partition.
 - USE SHARED :
 - Before using resources on a shared bus, a partition must configure that bus in its configuration
- IOP-Level Advantages
 - Greater flexibility
 - Cost reduction
 - Optimization of hardware resources
 - Ability to dynamically switch IOPs between partitions
- IOP-Level Disadvantages
 - Requires more depth hardware management skills
 - Requires good record keeping

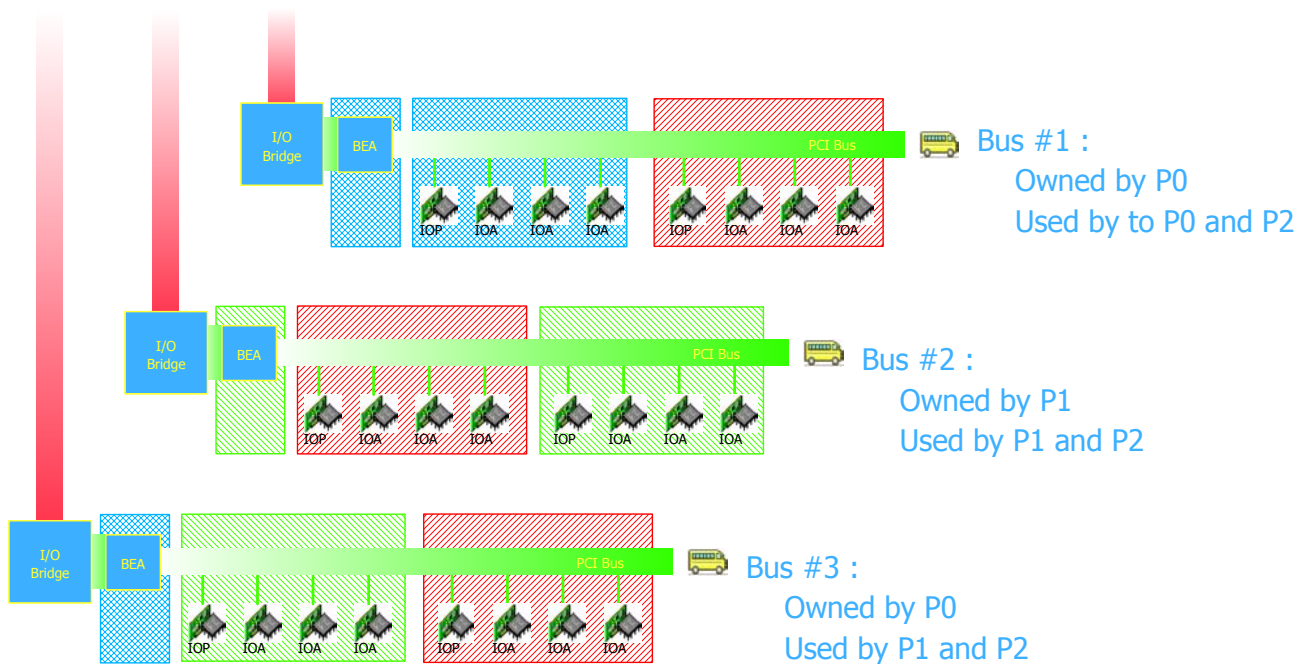
The LPAR Concepts

IOP-Level I/O Partitioning



The LPAR Concepts

IOP-Level I/O Partitioning



@ddy Pasteger / REAL Solutions

The LPAR Concepts

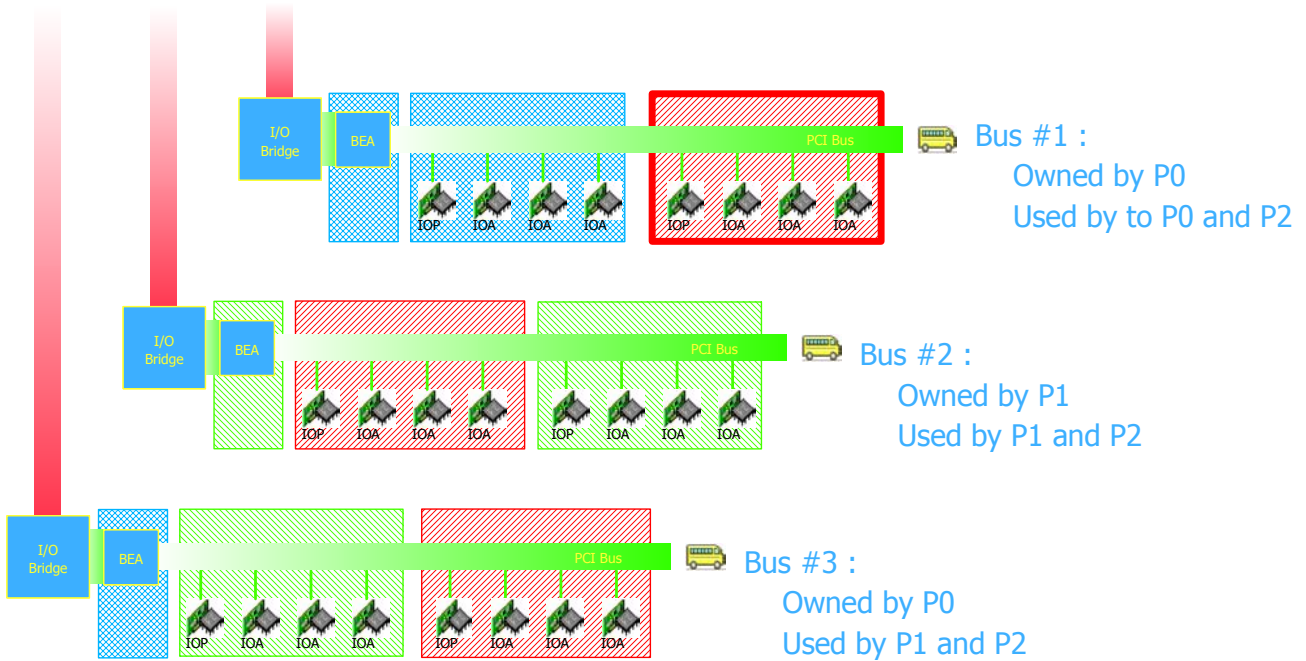
IOP and Device Switching

- Each IOP and all its attached resources can be switched by partitions that share the bus
 - Requires IOP-Level partitioning
- An IOP cannot be used by two partitions at the same time
- Switching Operations
 - Remove the IOP from its current partition
 - Add IOP to a different partition
- Candidates IOPs are :
 - IOPs which controls high cost devices (Tape Libraries, ...)
 - IOPs which controls low utilization devices (CD-ROM, ...)
 - IOPs which controls communications

@ddy Pasteger / REAL Solutions

The LPAR Concepts

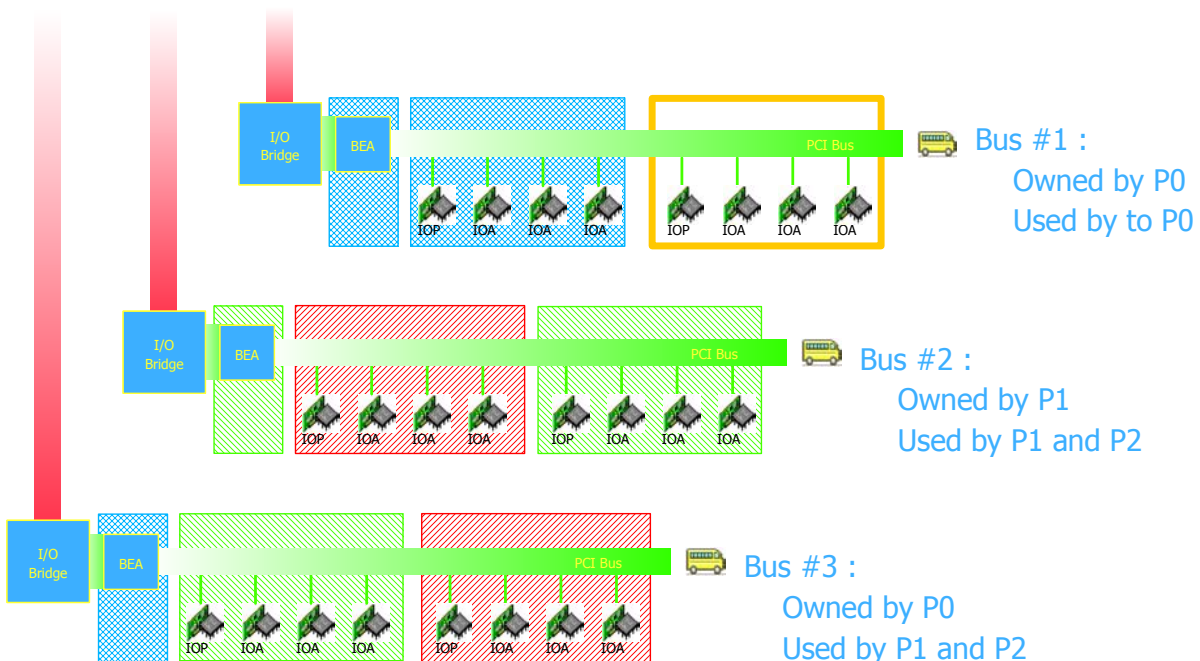
IOP and Device Switching : Identify resource



@ddy Pasteger / REAL Solutions

The LPAR Concepts

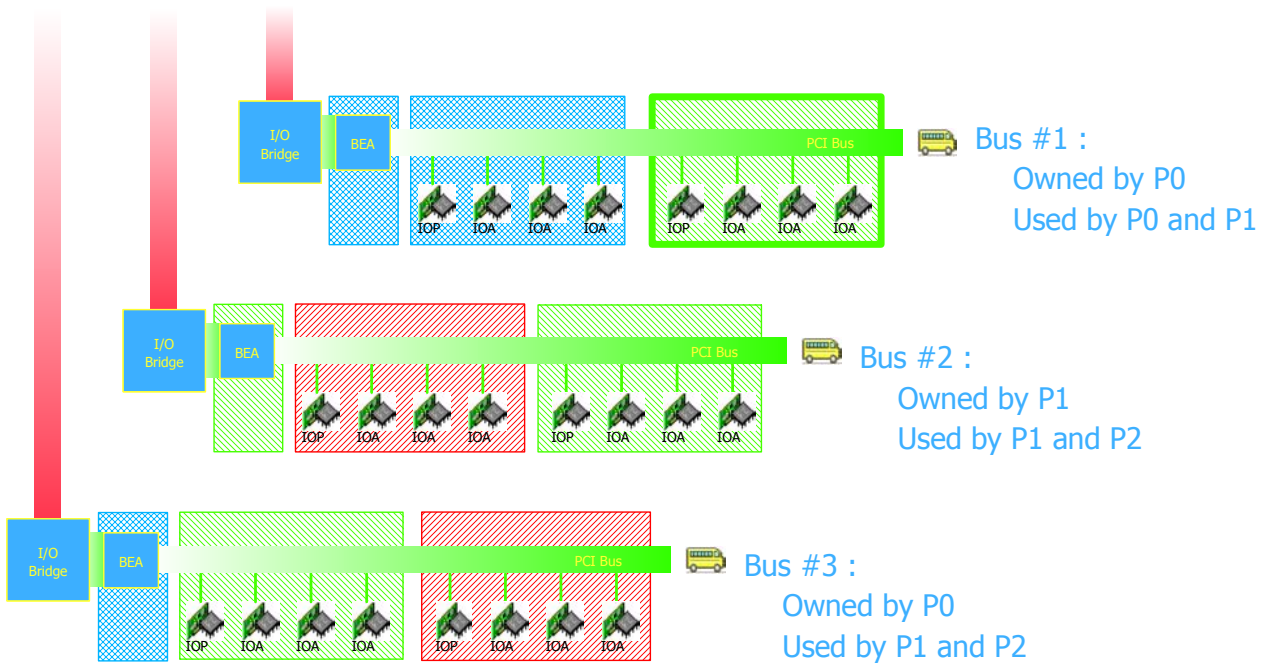
IOP and Device Switching : Deallocate resource



@ddy Pasteger / REAL Solutions

The LPAR Concepts

IOP and Device Switching : Allocate resource

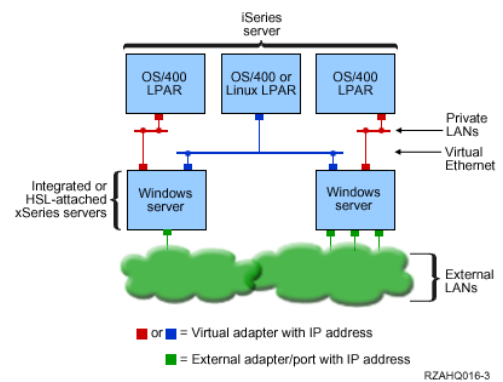


@ddy Pasteger / REAL Solutions

The LPAR Concepts

Other PLIC Services

- Shared Processor Pool
- Shared Memory Resources
- Virtual LAN
 - High performance communication between partitions
 - Up to 16 different virtual LAN connections thru virtual 1Gbps ethernet IOA
 - Partitions assigned to the same LAN are able to communicate thru that link
 - No additional hardware or software
- Virtual OptiConnect
 - Very high performance communication between partitions
 - Optional OS/400 feature
- HSL OptiConnect
 - Very high performance communication between LPARed systems
 - Optional OS/400 feature



RZAHO016-3

@ddy Pasteger / REAL Solutions

Partitioning offerings from various vendors

| | iSeries | Sun | HP | Unisys |
|--|---------|-------------------------|----------|---------------------|
| Operating System | OS/400 | Solaris | HP-UX | Windows 2000 |
| Partitioning type | Logical | Physical (board) | Physical | Physical |
| Maximum number of partitions | 32 | 16 | 16 | 4 |
| Minimum number of processor per partition | 0.1 | 1-4 | 4 | 4 |
| Processor increment above minimum | 0.01 | 1-4 | 4 | 4 |
| Memory increment above minimum | 1 MB | 512 MB | 512 MB | 8 GB |
| I/O increment above minimum bridge | IOP | 2 buses | 1 bus | Direct I/O |
| Independent movement of resource | Yes | No | CPU Only | No |
| Dynamic resource movement | Yes | Partial | Planned | Yes |
| High-speed internal communications | Yes | Yes (routing partition) | No | Yes (shared memory) |

Understand Requirements (1/2)

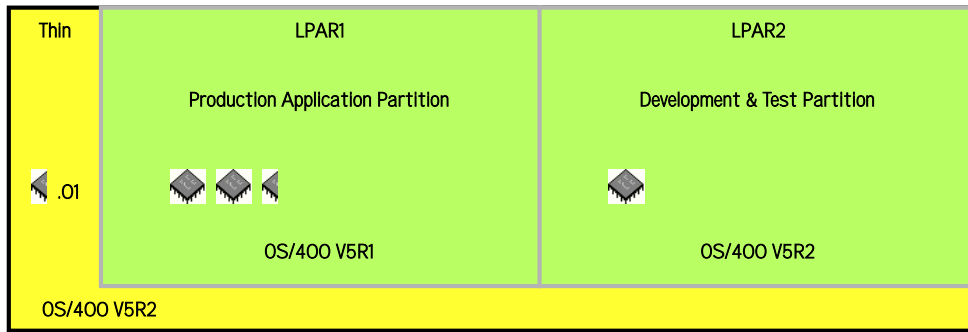
- Rules for SStar Processors
 - Up to 4 OS/400 partitions per processor
 - Up to 10 Linux partitions per processor
 - Up to 10 partitions per processor
 - Up to 32 partitions per system
- Rules for Power4 Processors
 - Up to 10 OS/400 partitions per processor
 - Up to 10 Linux partitions per processor
 - Up to 10 partitions per processor
 - Up to 32 partitions per system

Understand Requirements (2/2)

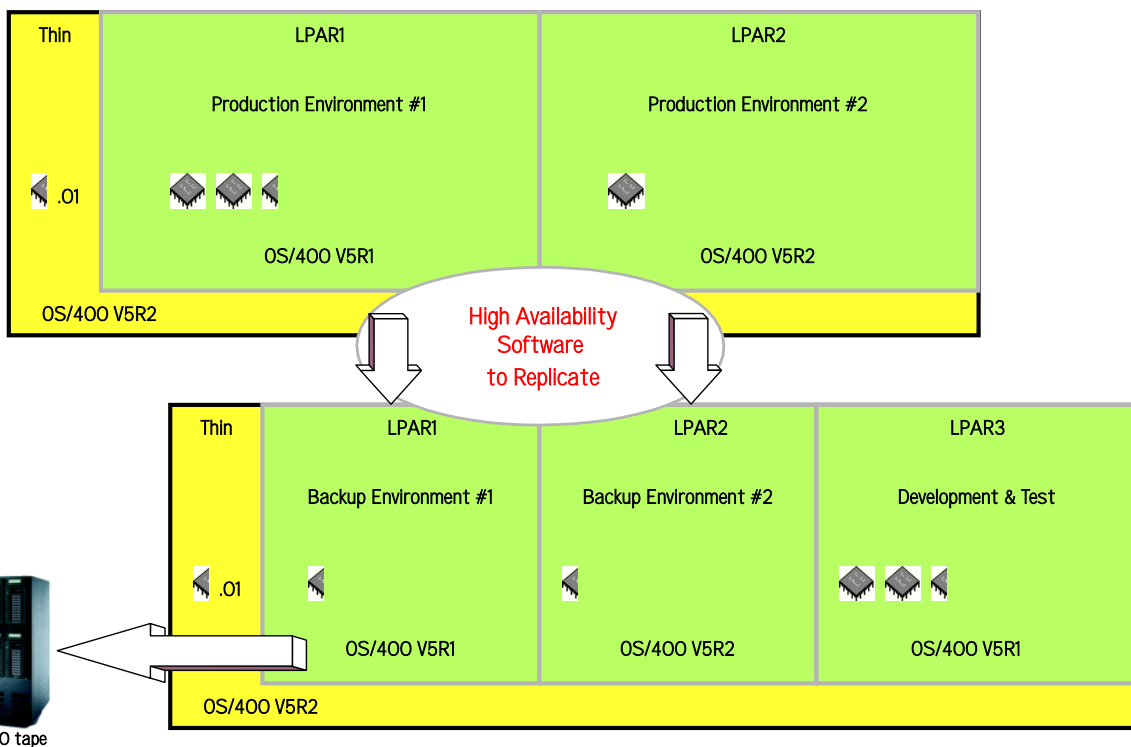
- In order to operate, each Partition must have :
 - Dedicated or Shared Processor
 - Interactive Capacity
 - Main Storage
 - Load Source Unit
 - Available Console
 - Access to Alternate-IPL device
- Consoles
 - Twinax
 - Operations Console, Direct Attach
 - Operations Console, LAN Attach

Chapter 3. The LPAR Benefits

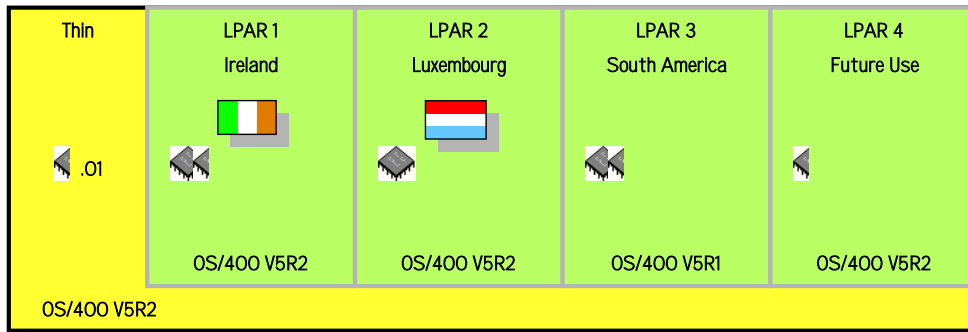
The LPAR Benefits



The LPAR Benefits

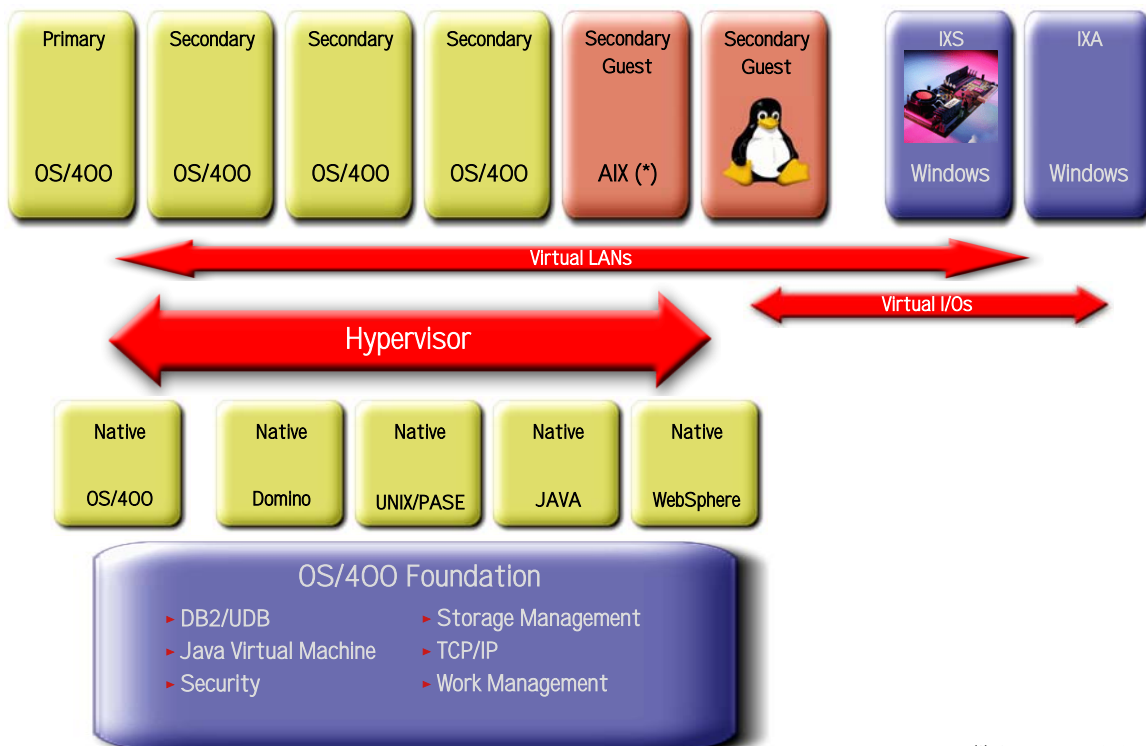


The LPAR Benefits



Eddy Pasteger / REAL Solutions

The iSeries Benefits



(*) Statement of Direction

Eddy Pasteger / REAL Solutions

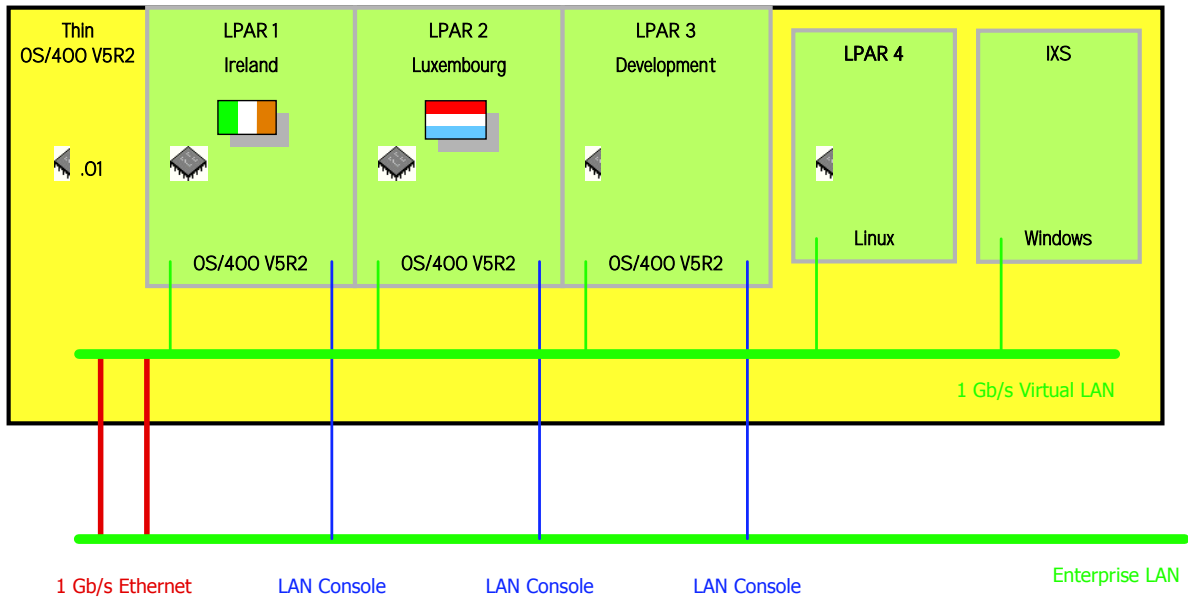
Chapter 4. Build and LPARed System by Example

Build and LPARed System by Example

LPAR Planning

- Perform proper capacity planning for each logical system
 - Determine the number of partition needed
 - Determine the size of each
- Call IBM or IBM Business Partner for support
 - Assist with LPAR planning
 - Assist architecting LPAR solution
- Complete the system design phase
 - LPAR Validation Tool (LVT)
 - Configuration planning
 - Work Sheets
- Conduct a Solution Assurance review with IBM or IBM Business Partner
- Order the necessary hardware and software
 - Based on the validated solution
- Install
- Document

Build and LPARed System by Example



Build and LPARed System by Example

Planning - System Selection

| | | |
|----------------------|----------------|---------------|
| System Model | i825 | |
| Processor Package | 2473/0873/7416 | Enterprise |
| Interactive Feature | 7418 | 100% |
| System Memory | 8 GB | |
| Number of Partitions | 5 | |
| Number of Processors | 3/6 | 3300/6600 CPW |

Build and LPARed System by Example

Planning - Partition Specifications

| Partition | OS Version | Shared | # Processors | Batch CPW | Memory (MB) | Int % | Int CPW |
|----------------------|------------|--------|--------------|-----------|-------------|-------|---------|
| P0 - Primary | V5R2M0 | Yes | 0.1 | 116 | 512 | 0 | 0 |
| P1 - Ireland | V5R2M0 | Yes | 1.0 | 1166 | 3072 | 34 | 1166 |
| P2 - Luxembourg | V5R2M0 | Yes | 1.0 | 1166 | 3072 | 34 | 1166 |
| P3 - Development | V5R2M0 | Yes | 0.7 | 816 | 1280 | 20 | 700 |
| P4 - Linux (HOST=P0) | Linux | Yes | 0.2 | 233 | 256 | N/A | N/A |
| Unallocated | - | - | 0.0 | 3 | 0 | 12 | |
| Windows (HOST=P0) | Windows | - | - | - | - | - | - |

Eddy Pasteger / REAL Solutions

Build and LPARed System by Example

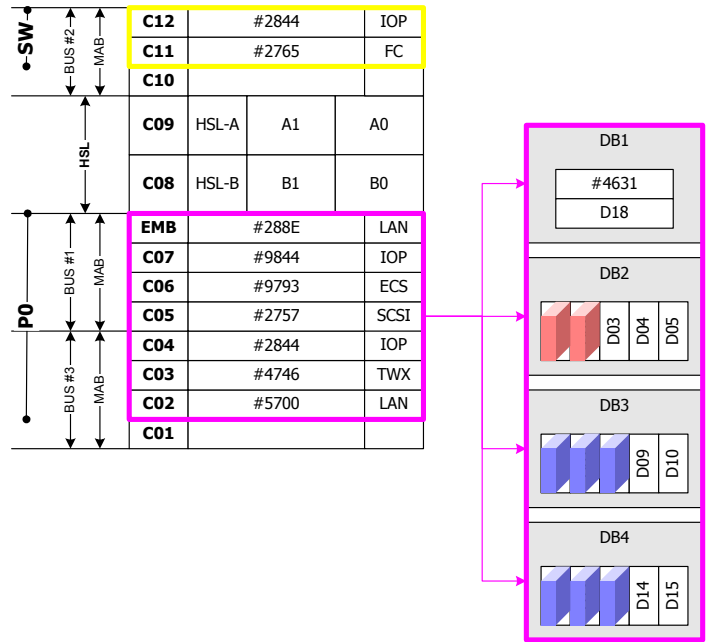
Planning - Partition Specifications

| Partition | DASD | Protection | #4326 | Backup | OPTMLB | LAN | WAN |
|----------------------|--------|------------|-------|-----------|--------|---------|-----------|
| P0 - Primary | - | MIRRORING | 2 | DVD-RAM | No | 1 Gbps | Yes - ECS |
| P1 - Ireland | 300 GB | RAID-5 | 9 | FC TAPE | Yes | Virtual | Yes |
| P2 - Luxembourg | 300 GB | RAID-5 | 9 | FC TAPE | Yes | Virtual | Yes |
| P3 - Development | 500 GB | RAID-5 | 15 | FC TAPE | No | Virtual | No |
| P4 - Linux (HOST=P0) | 100 GB | RAID-5 | 3 | SCSI TAPE | No | Virtual | No |
| Unallocated | | | | | | | |
| Windows (HOST=P0) | 100 GB | RAID-5 | 3 | SCSI TAPE | No | Virtual | No |

Eddy Pasteger / REAL Solutions

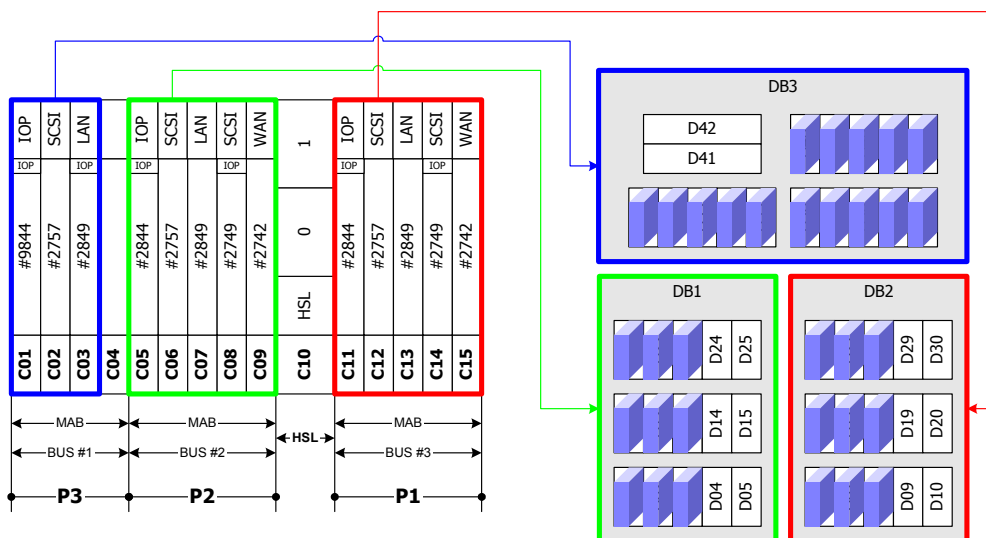
Build and LPARed System by Example

Planning - iSeries 825 System Unit



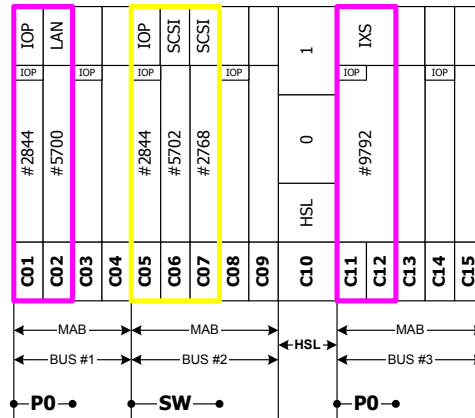
Build and LPARed System by Example

Planning - #5094 PCI-X Expansion Tower



Build and LPARed System by Example

Planning - #5088 PCI-X Expansion Tower



Build and LPARed System by Example

Planning - LPAR Validator Tool

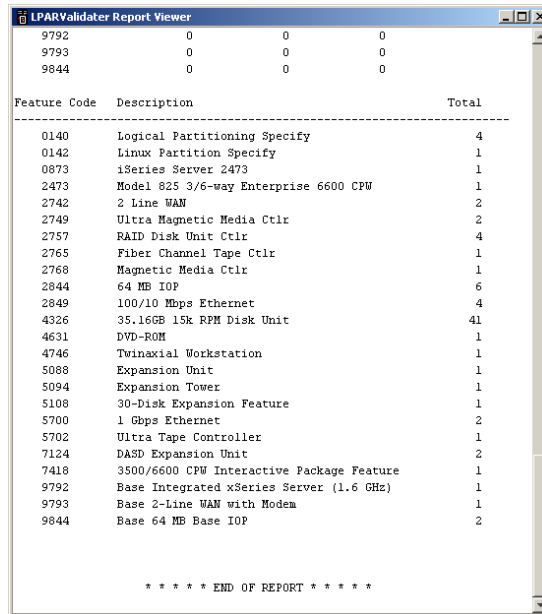
The screenshot shows the LPARValidator tool interface. The main window displays a list of hardware components and their configurations:

| Add/Remove | Slot | IOP/IOA/Dev | Partition | Description |
|------------|--------|-------------|-----------|------------------------------|
| Remove | DB4 | 4326x3 | Primary | 35.16GB 15k RPM Disk Unit |
| Add | DB4(a) | | Primary | Disk Drives |
| Add | DB4(b) | | Primary | Disk Drives |
| Remove | DB3 | 4326x3 | Primary | 35.16GB 15k RPM Disk Unit |
| Add | DB3(a) | | Primary | Disk Drives |
| Add | DB3(b) | | Primary | Disk Drives |
| Remove | DB2 | 4326x2 | Primary | 35.16GB 15k RPM Disk Unit |
| Add | DB2(a) | | Primary | Disk Drives |
| Add | DB2(b) | | Primary | Disk Drives |
| Add | D18 | | Primary | Internal Tape/CD-ROM/DVD-RAM |
| Remove | D17 | 4631 | Primary | DVD-ROM |
| Remove | C12 | 2844 | Sw-LPAR | 64 MB IOP |
| Remove | C11 | 2765 | Sw-LPAR | Fiber Channel Tape Ctr |
| Add | C10 | | Sw-LPAR | IOA |
| Remove | C07 | 9844 | Primary | Base 64 MB Base IOP |
| Remove | EMB | 288E | Primary | 100/10 Mbps Ethernet |
| Remove | C06 | 9793 | Primary | Base 2-Line WAN with Modem |
| Remove | C05 | 2757 | Primary | RAID Disk Unit Ctr |
| Remove | C04 | 2844 | Primary | 64 MB IOP |
| Remove | C03 | 4746 | Primary | Twinaxial Workstation |
| Remove | C02 | 5700 | Primary | 1 Gbps Ethernet |
| Add | C01 | | Primary | IOA |

At the bottom of the main window, a red message states: "All partition requirements have been satisfied." A large yellow star with the word "Yes!" is overlaid on this message.

An "About LPARValidator" dialog box is open in the bottom right corner, showing the version (3.0) and copyright information (© Copyright 2001 IBM Corporation).

Planning - LPAR Validator Tool



| Feature Code | Description | Total |
|--------------|---|-------|
| 9792 | | 0 |
| 9793 | | 0 |
| 9844 | | 0 |
| 0140 | Logical Partitioning Specify | 4 |
| 0142 | Linux Partition Specify | 1 |
| 0873 | iSeries Server 2473 | 1 |
| 2473 | Model 825 3/6-way Enterprise 6600 CPW | 1 |
| 2742 | 2 Line WAN | 2 |
| 2749 | Ultra Magnetic Media Ctrlr | 2 |
| 2757 | RAID Disk Unit Ctrlr | 4 |
| 2765 | Fiber Channel Tape Ctrlr | 1 |
| 2768 | Magnetic Media Ctrlr | 1 |
| 2844 | 64 MB IOP | 6 |
| 2849 | 100/10 Mbps Ethernet | 4 |
| 4326 | 35.16GB 15k RPM Disk Unit | 41 |
| 4631 | DVD-ROM | 1 |
| 4746 | Twinaxial Workstation | 1 |
| 5088 | Expansion Unit | 1 |
| 5094 | Expansion Tower | 1 |
| 5108 | 30-Disk Expansion Feature | 1 |
| 5700 | 1 Gbps Ethernet | 2 |
| 5702 | Ultra Tape Controller | 1 |
| 7124 | DASD Expansion Unit | 2 |
| 7418 | 3500/6600 CPW Interactive Package Feature | 1 |
| 9792 | Base Integrated xSeries Server (1.6 GHz) | 1 |
| 9793 | Base 2-Line WAN with Modem | 1 |
| 9844 | Base 64 MB Base IOP | 2 |

***** END OF REPORT *****

Implementation & Operations

Interfaces available for installation and configuration

- Via 5250 displays thru DST or SST
- Via iSeries Navigator

Implementation & Operations

New Logical Partition - Welcome

Welcome to the New Logical Partition wizard!

This wizard helps you complete the steps to create a new logical partition. You can learn more about logical partitioning in the Information Center. To access the Information Center, click Help.

You can click Cancel at any time to leave the wizard. Any changes you have made will be lost.

← Back **Next** → ✓ Finish ✗ Cancel

New Logical Partition - Type of Partition

What type of logical partition would you like to create?

- Create new OS/400 partition**
Creates a fully functioning logical partition for use with OS/400.
- Create new partition for future use**
Creates a partition with only memory and no other hardware resources. The partition can not be used until the required hardware resources are moved to it. You can avoid full system restarts later by reserving partitions.

Note: The Information Center has information about logical partition minimum requirements. To access the Information Center, click Help.

← Back **Next** → ✓ Finish ✗ Cancel ? Help ?

Implementation & Operations

New Logical Partition - Name

Each partition is required to have a name. Using the same name for both the partition name and the partition's network attributes is recommended.

What is the name of your new logical partition?

Partition name:

Partition ID:

← Back **Next** → ✓ Finish ✗ Cancel

New Logical Partition P4 - Memory

Each partition requires memory. Move at least the minimum memory to the new partition. For information about minimum memory requirements, click Help.

Available memory:

| Partition | Current | Pending | Minimum | Maximum |
|---------------------|---------|---------|---------|---------|
| Primary (0) | 512 | 512 | 512 | 6144 |
| P1 (1) | 3072 | 3072 | 512 | 6144 |
| P2 (2) | 1024 | 1024 | 512 | 6144 |
| P3 (3) | 1024 | 1024 | 512 | 6144 |
| Unassigned Hardware | 512 | 0 | 0 | 6144 |

Memory to move: MB

Move

Memory for partition P4:

| From Partition | Moved Memory |
|---------------------|--------------|
| Unassigned Hardware | 512 |

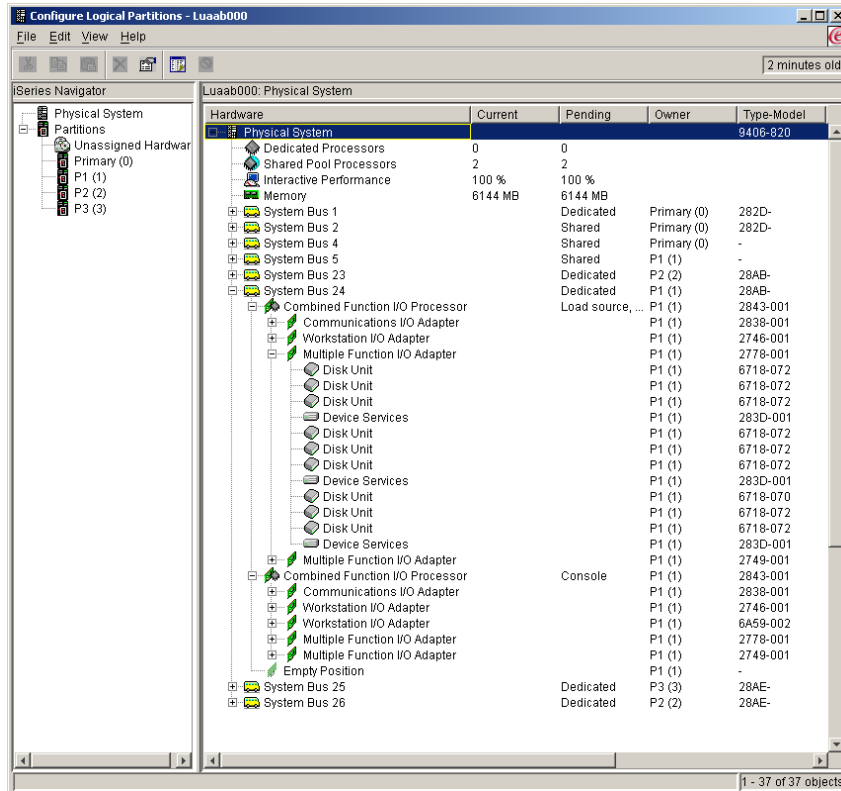
Remove

Total memory for partition P4: MB

Set Minimum/Maximum...

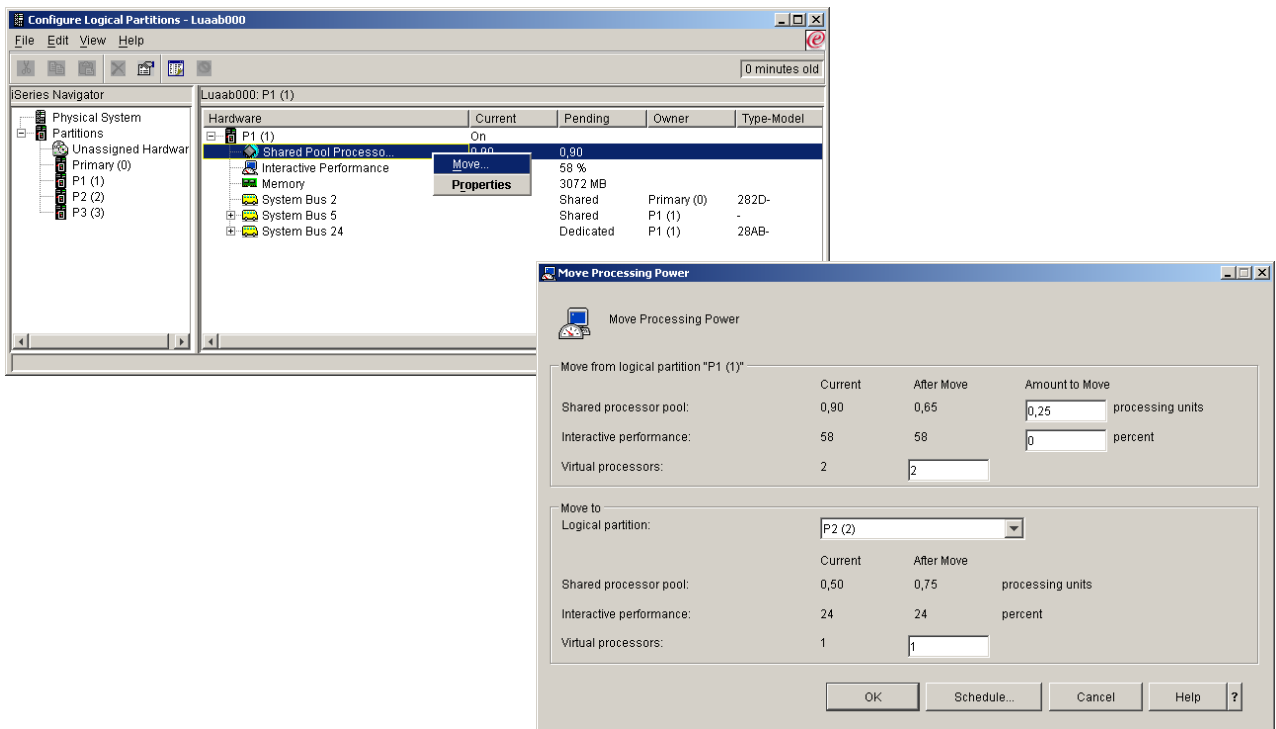
← Back **Next** → ✓ Finish ✗ Cancel ? Help ?

Implementation & Operations

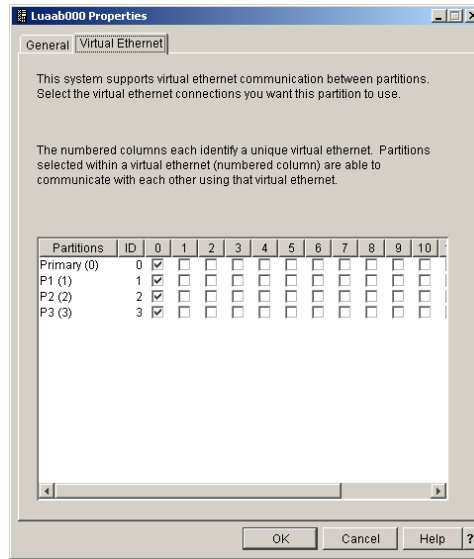


Eddy Pasteger / REAL Solutions

Implementation & Operations



Eddy Pasteger / REAL Solutions



Chapter 5. Linux on iSeries LPAR

Requirements

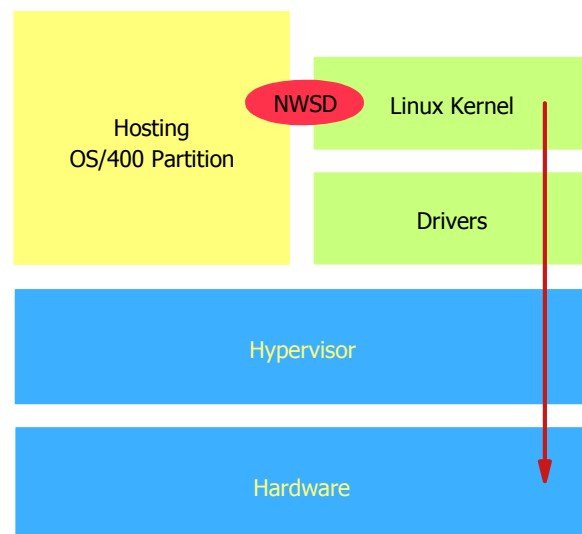
- Primary partition must be V5R1 or higher
 - V5R2 recommended to benefit of dynamic CPU and memory allocation
- Hosting partition must be V5R1 or higher
- Linux distribution supported
 - RedHat, SuSE or Turbolinux
 - Power PC Edition (cf. Apple MacIntosh)
 - 32-bits or 64-bits ... Power4 processors requires a 64-bits edition

Two ways to implement

- Direct I/O
- Virtual I/O

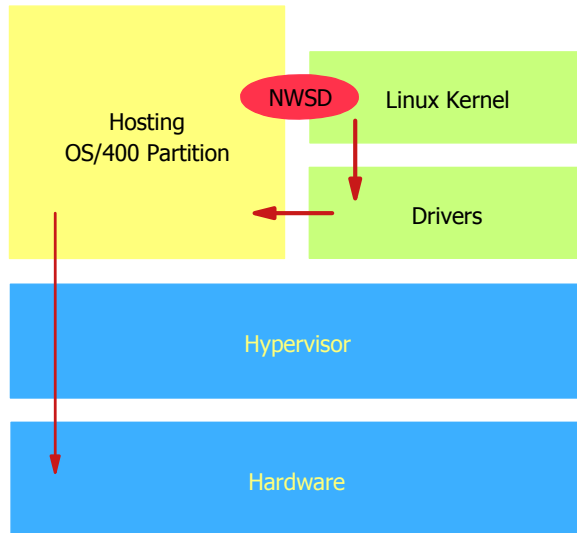
Direct I/O

- IOPs are not needed
 - Linux doesn't know what an IOP is !
- Linux owns its (supported) IOAs
 - Drivers are provided to handle them
 - Fibre Channel, Disk Controller, Network Adapter



Virtual I/O

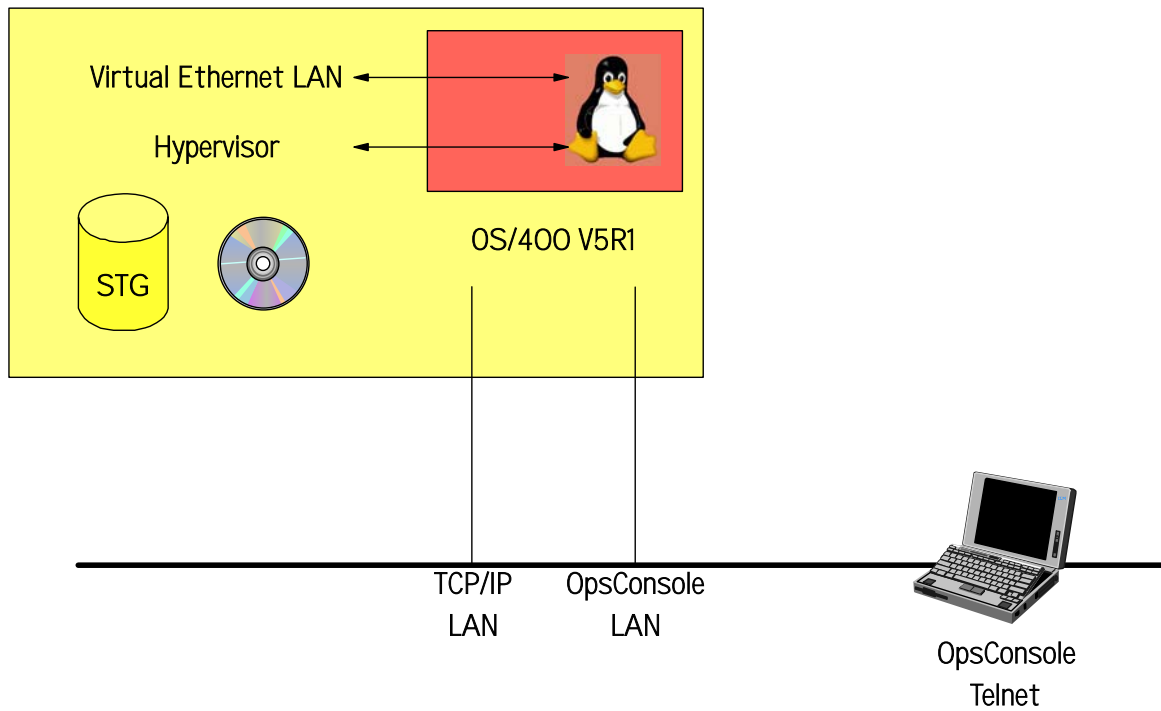
- No additional hardware is required
 - I/Os are redirected to the hosting OS/400 partition
- Virtual ...
 - CD-ROM, Disk, LAN adapter, Tape



Operations

- Prepare a Linux Console
 - External workstation with a simple telnet client
 - Telnet the hosting OS/400 partition on port 2301 and select the guest partition to work with
- Create the Guest Linux Partition
 - Allocate resources (processor, memory)
- Create a Network Server Description (NWSD)
 - Network Server Type = *GUEST, specify Partition ID
 - IPL source is an IFS stream file (boot file)
- Create Network Server Storage Spaces (NWSSTG)
 - For virtual disk I/Os only
- Vary On NWSD
 - Partition starts on IFS stream file
 - Installation continues on CD-ROM
- Process to a regular Linux installation

Linux on iSeries LPAR



@ddy Pasteger / REAL Solutions

iSeries Dynamic Logical Partitioning

Directions

- AIX on an iSeries Partition (2004 as IBM SoD)
- Hardware Management Console
- Dynamic Partition Sizing
- ...

@ddy Pasteger / REAL Solutions

References

- **Fortress Rochester : the inside story of IBM iSeries**
 - Frank G. Soltis, NEWS/400 Books, 2001
- **Capacity Planning for Logical Partitioning on the IBM eServer iSeries Server**
 - SG24-6209-00, IBM RedBooks, 2001
- **Linux on iSeries, an implementation guide**
 - SG24-6232-00, IBM RedBooks, 2001
- **IBM eServer iSeries System Handbook**
 - GA19-5486-23, IBM RedBooks, 2003
- **Internet**
 - <http://www.ibm.com/servers/eServer/iSeries/LPAR>

- **E-mail**
 - eddy.pasteger@real.lu

Questions ?

Thank You !