



IBM Systems and Technology Group University 2005

# IBM Systems and Technology Group University 2005





IBM Systems and Technology Group University 2005

# Selling Virtualization and CoD on @server p5

Course #: P16


**Nha Dai Vu**  
Marketing Manager, @server p5 Virtualization  
[nhavu@us.ibm.com](mailto:nhavu@us.ibm.com)

**Eric Finkel**  
Consulting Learning Specialist  
[efinkel@us.ibm.com](mailto:efinkel@us.ibm.com)



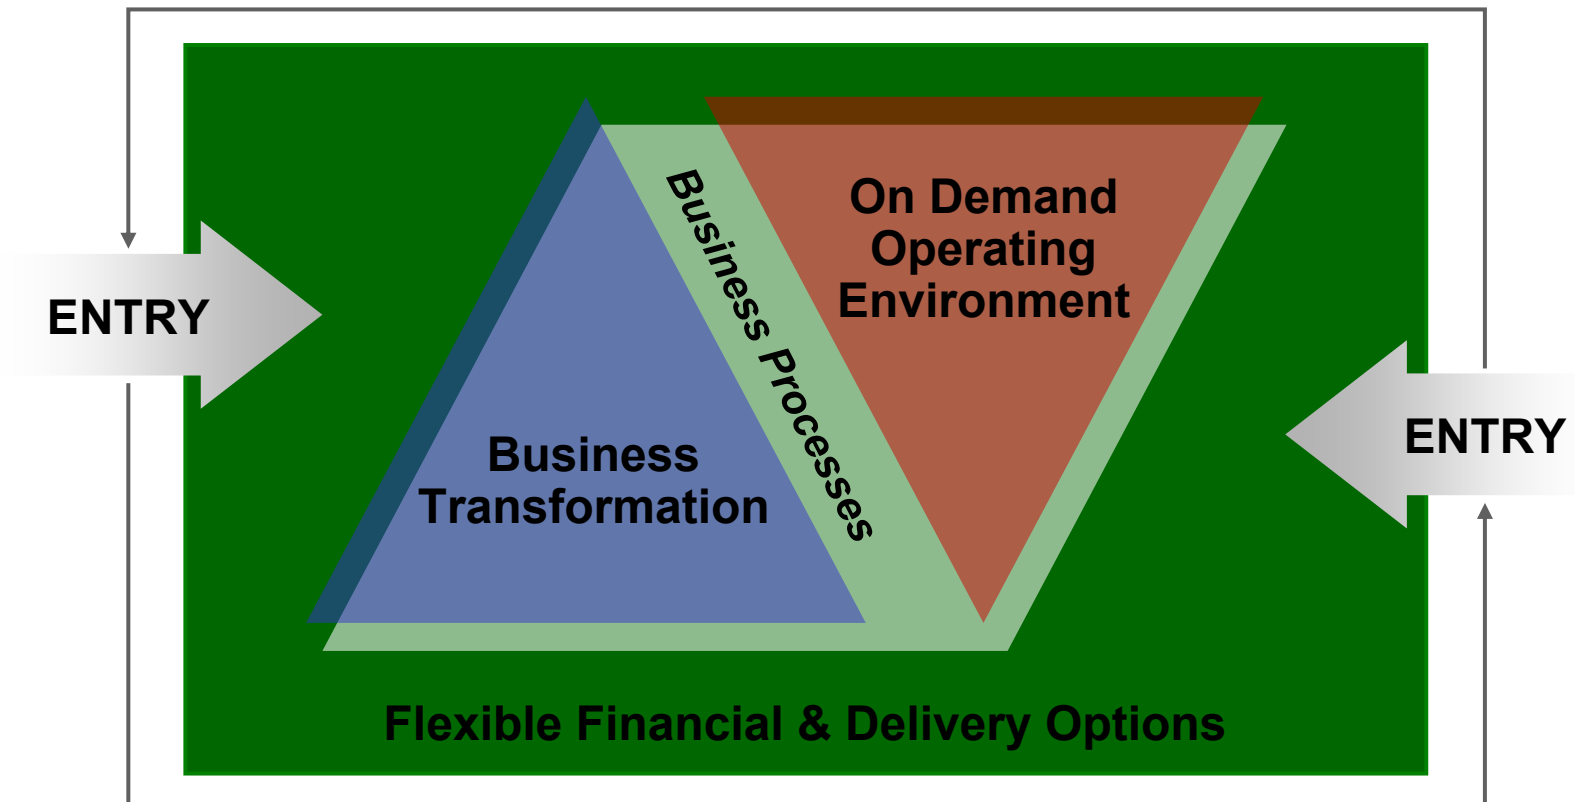
# Agenda

- **Learning Objectives**
- **Introduction and Strategic Context**
- **Offering Description**
  - Virtualization on POWER5
  - Capacity on Demand (CoD)
- **Selling Strategies**
- **Competitive Analysis**



# Road to an On Demand Business via Virtualization Solutions


Where you start depends on YOUR client's priorities.



- Increasing flexibility is the key—business models, processes, infrastructure, plus financing and delivery

# Why IT optimization and Virtualization?

1. Complexity and administrative cost are growing faster than hardware cost.
2. Growing revenue while containing cost remains #1 concern of management.
3. Businesses cannot address the need for integration of people, etc without first addressing the underlying infrastructure.



- ✓ Lower the cost of their existing infrastructure
- ✓ Reduce the complexity of adding to that infrastructure
- ✓ Build heterogeneous infrastructure across multiple Data Centers that are more responsive to their business needs


# Virtualization Solutions: Stages of Deployment

**Orchestrate Infrastructure:** Sense and respond to changes based on business policies




**Virtualize Outside The Enterprise:** Suppliers, partners, customers and external resources

**Secure Cross Enterprises:** Enable internal and external integration and resources.



**Virtualize The Enterprise:** Enterprise wide Grids and Global Fabrics

**Automate Workflows:** Tasks like change/ configuration, ITIL processes



**Virtualize Unlike Resources:** Heterogeneous systems, application based Grids and networks



**Virtualize Like Resources:** Homogenous systems, storage and networks

**Updated IT Governance and Management Processes**


# IBM Virtualization Engine

- *EWLM*
- *Director Multi-platform*
- *System provisioning*
- *Grid Toolbox*
- *VE Console*

- *Dynamic LPAR (DLPAR)*
- *Micro-Partitioning/Shared processor LPAR*
- *Partition Load Manager (PLM)*
- *Virtual LAN*
- *Virtual I/O (VIO)*
  - *Shared Ethernet*
  - *Virtual SCSI*



IBM Virtualization Solutions






- *IBM Dynamic Infrastructure (IDI) for mySAP Business Suite*

## KEY BENEFITS OF VIRTUALIZATION

- **Reduce** operating costs
- Maximize **utilization** of computing resources
- **Simplify** the monitoring and management of infrastructure
- **Ease** the pain of **deploying standardized** resources quickly to improve interoperability and flexibility

# Advanced POWER Virtualization (APV)

- 
  - Micro-Partitioning™**
    - Supports up to 254 partitions, 1/10th of processor granularity
    - Key enabled (like CoD); on for all active processors (not selective); once turned on, cannot be turned off; enablement key installed during manufacture; feature can also be purchased as an MES upgrade
    - Requires AIX 5.3 or Redhat AS 3.3, or SuSE SLES 9
  
- 
  - Virtual I/O (VIO)**
    - Software that ships on a CD
    - Provides Shared Ethernet and Virtual Storage support
    - Supports AIX 5.3 and Linux (RHEL AS 3.3, SLES 9) clients
    - Requires Micro-Partitioning Firmware Key
  
- 
  - Partition Load Manager (PLM)**
    - Software that ships on a CD
    - Supports AIX 5.2 and AIX 5.3
    - Provides cross partition CPU/memory management

- Pricing<sup>1</sup> (per processor)**

Products	Feature Number	Advanced POWER Virtualization	VIO SWMA <sup>3</sup> (1 yr / 3 yr)	PLM SWMA <sup>3</sup> (1 yr / 3 yr)
p5-520 <sup>2</sup>	7940	\$590	\$90 / \$245	\$20 / \$55
p5-550 <sup>2</sup>	7941	\$590	\$90 / \$245	\$20 / \$55
p5-570 <sup>2</sup>	7942	\$990	\$90 / \$245	\$20 / \$55
p5-590	7992	\$0	\$0 / \$155 <sup>4</sup>	\$0 / \$35 <sup>4</sup>
p5-595	7992	\$0	\$0 / \$155 <sup>4</sup>	\$0 / \$35 <sup>4</sup>

1 Prices are current as of 12/31/2004 and are based on U.S. List Prices

2 Prices for p5-520, p5-550, and p5-570 also applies to p5-520 Express, p5-550 Express, p5-570 Express offerings, respectively

3 SWMA prices are for standard 9x5 support; 24x7 support available for additional fee

4 3-year SWMA prices shown are actually for Years 2 and 3 (Year 1 has a price of \$0)



# Advanced POWER Virtualization (APV) vs. Advanced OpenPower Virtualization (AOPV)

## Pricing/packing differences

	eServer p5 APV	OpenPower AOPV
<b>COMPONENTS</b>		
LPAR/DLPAR	NO (already part of base hardware)	YES
VLAN	NO (already part of base hardware)	YES
Micro-Partitioning	YES	YES
VIO	YES	YES
PLM	YES	NO

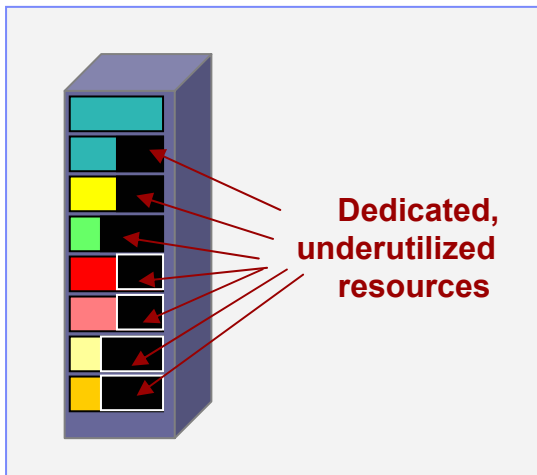
1 Price varies by system (e.g., included in hardware price for p5-590 and p5-595; \$590 per CPU for p5-520 and p5-550; \$990 per CPU for p5-570)

2 1-year SWMA contract (at minimum) is required

NOTE: Prices are current as of 12/31/2004 and are based on U.S. List Prices


# pSeries had industry leadership with POWER 4 DLPARs...now we increase the gap vs. competitors with POWER5/APV

## With Partitioning




### POWER4

- Granular, multiple OS images per system - Single CPU, 256 MB Memory, I/O Adapter level granularity
- LPAR fault and security isolation enforced by Hardware & Firmware
- AIX 5.2 Dynamic Reconfiguration (DR): Processor, memory and adapter DR w/o reboot



### POWER5

- **Increased granularity**
  - Micro-partitions (1/10 CPU)
  - Up to 254 Micro-Partitions
  - Granular CPU entitlements (1/100 CPU)
  - Granular memory entitlements (16MB)
- **Increased flexibility**
  - Shared processor pools
  - Shared I/O
  - VLAN
  - Multi-OS Support (AIX 5.2\*/5.3, Linux, i5/OS)



\* AIX 5L v5.2 partitions must be dedicated LPARS – no access to shared processor pools and Virtual I/O

# POWER 5 Virtualization: Micro-Partitioning AIX 5.3 and Linux 2.6 Exploitation

## (2) MAJOR DEPLOYMENT MODELS

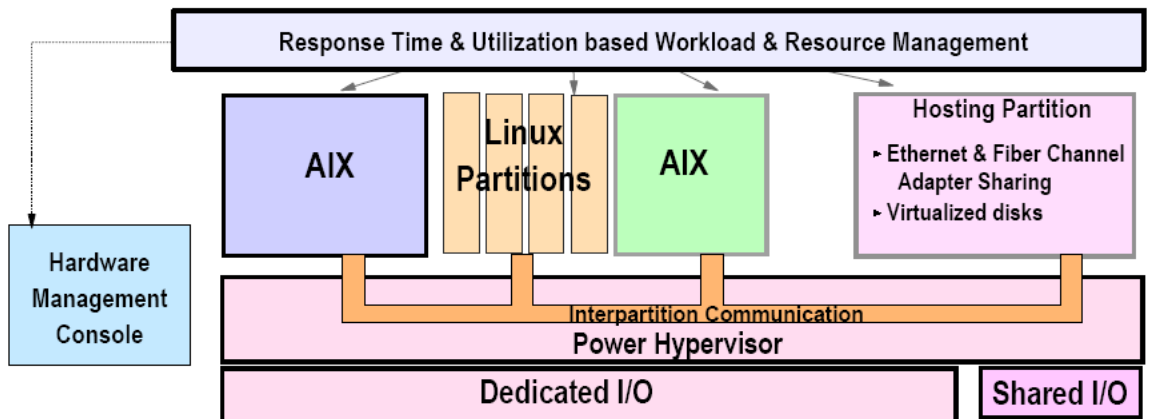
1

### Responsive OS images in Shared Processor LPARs with higher utilization

- Historically, needed to set aside capacity for unexpected peaks; run at low utilization rates
- Rapid increase in capacity (10ms) - high priority workloads get resources when needed; lower priority/batch workloads consume all unused resources

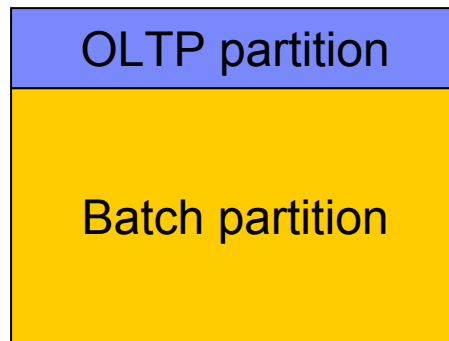
2

### Server consolidation of small underutilized servers using multiple sub-CPU micro-partition




# Shared Processor LPAR - responsive partition example

- Work can be **prioritized** to give **time critical tasks** a high fraction of resource
- **Lower priority partitions** can run with **less dedicated resource**, but can “soak” up spare CPU cycles **automatically**




When interactive load is **low**, batch partition uses available resources



**Sudden surge in demand...**High interactive load gets precedence when needed

# Shared Processor LPAR – configuration detail



## Configuration considerations

### ➤ Number of Virtual processors

- Minimum Virtual CPU,
- Desired Virtual CPU
- Max Virtual CPU

### ➤ Entitled Capacity in Processing Units


- In units of 1/100 of a CPU
  - Minimum processing units
  - Desired processing units
  - Maximum processing units
- Capped or Uncapped

### ➤ Variable Weight

- % share (priority) of surplus capacity
- Only for uncapped

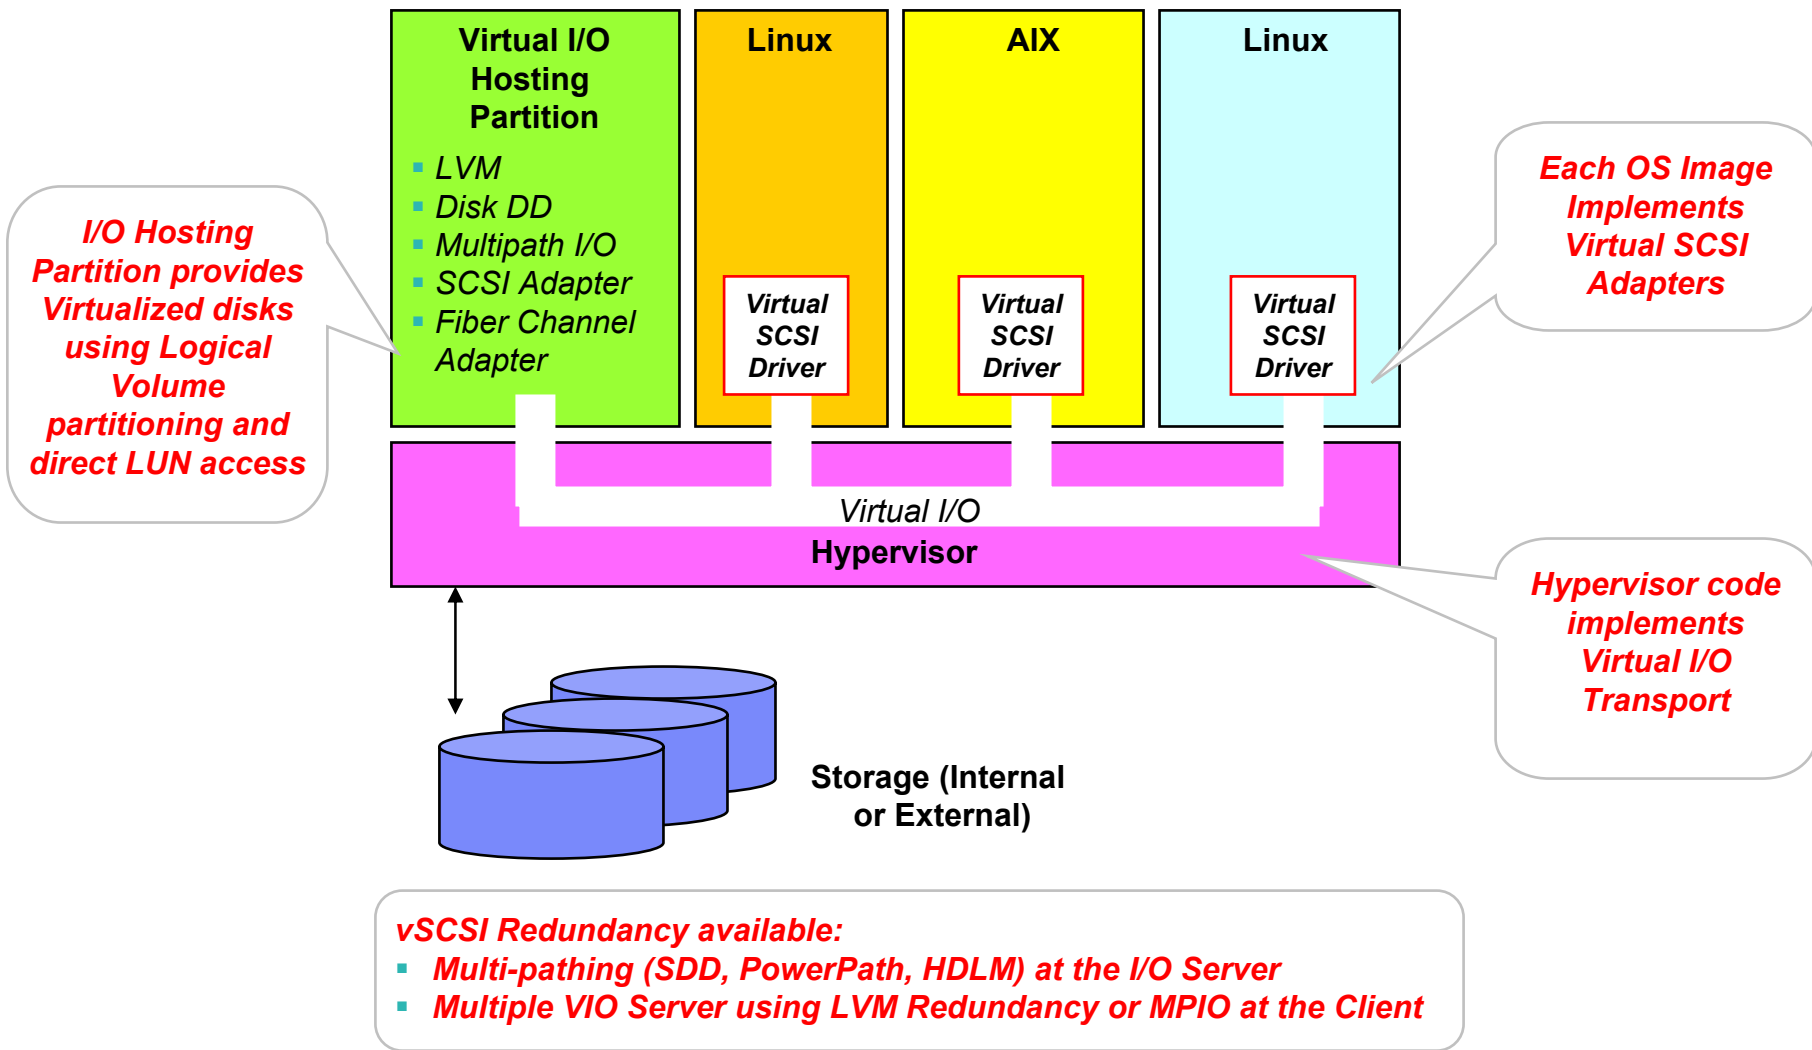
# Virtualizing I/O (1 of 2)

## Virtual Ethernet (VLAN) and Shared Ethernet adapters

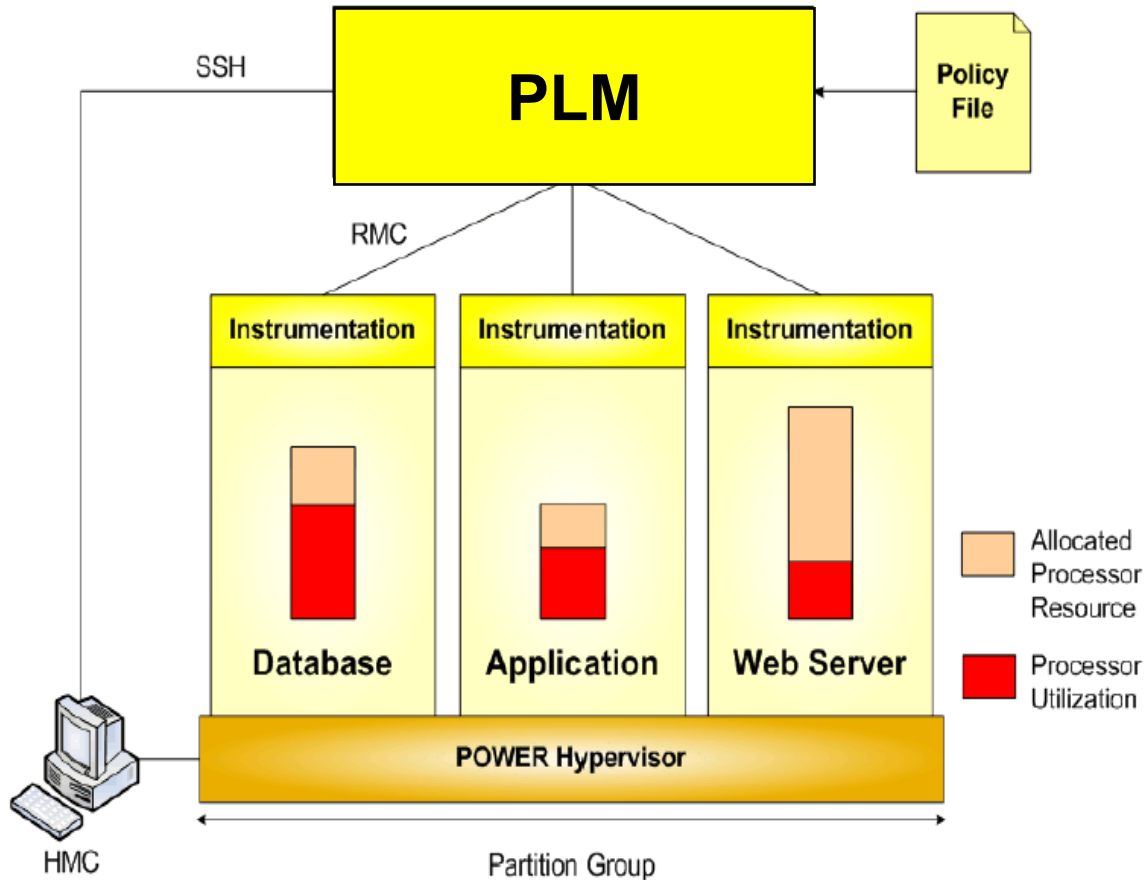


# Virtualizing I/O (2 of 2)

## Shared Fiber Channel Adapters and Virtualized Disks



# Partition Load Manager for AIX 5L @server p5 systems



- **Policy-based**, automatic partition resource tuning
- Dynamically adjust **CPU** and **memory** allocation
- Manages **AIX 5L v5.2** or **5.3**
- Supports both **dedicated partitions** and **micro-partitions**
- Managed partitions are **grouped**; **different instances of PLM** on the same server can manage different groups



# OS Support for Virtualization

Function	AIX 5.2F	AIX 5.3	Linux SLES 9	Linux RHEL 3 U3
<b>DLPAR</b>				
▪ Processor	✓	✓	✓	✗
▪ Memory	✓	✓	✗	✗
▪ I/O	✓	✓	✓	✗
<b>Virtual LAN</b>	✗	✓	✓	✓
<b>Micro-partitions</b>	✗	✓	✓	✓
<b>Virtual Storage</b>	✗	✓	✓	✓
<b>Virtual Ethernet</b>	✗	✓	✓	✓
<b>Partition Load Manager</b>	✓	✓	✗	✗

## i5/OS support available for 1.65GHz. p5-570, p5-590, p5-595

- i5/OS can participate in the shared processor pool, but only to the processor maximum
- I5/OS can utilize VLAN, but canNOT be a client to VIO server

## NOTES:

- DLPAR standard on @server p5; part of AOPV on OpenPower
- Virtual I/O Hosting Partition is an “appliance” – encapsulated partition available on both @server p5 and OpenPower

# Communicating the value of Advanced POWER Virtualization

**Reduce total cost of ownership...**

**...by making better use of IT assets by significantly improving server utilization.**


- Consolidation of under-utilized (e.g., 10-15%) workloads (onto partitions with <1 CPU)
- Approach mainframe-like utilization without sacrificing resource availability
- Economic I/O model (e.g., Shared Ethernet, Shared storage)
- Simplify IT management (e.g., rapid deployment of additional capacity)

**Improve business agility/flexibility..**

**...by dynamically and automatically re-allocating resources to applications as needed -- to better match changing business cycles or handle unexpected surge in demand.**

- Support multiple workloads/departments/applications
- Rapid response (10 ms) to changing workload needs via shared processor pool
- Automatic access to spare capacity via APV/Reserve CoD
- Exploitation of VLAN (e.g., high-speed communication)


# Quantifying the TCO benefits – (3) steps



## Sell initially on raw p5 price/ performance




The capacity to handle the workload of 32 systems and 256 processors -- with room to spare



## Position virtualization as game-changer



Average UNIX server utilization is estimated to be 15%<sup>3</sup>  
Virtualization enables mainframes to achieve more than 85% utilization.



## Estimate TCO benefits of SCON

### System infrastructure costs:

- Hardware maintenance
- Facilities (floor space, power consumption)
- License, update subscription and support fees for systems software and databases
- Personnel costs (salaries, benefits, bonuses, training) for management, operation and technical support

### Costs of other IT functions

- Personnel costs for application development
- Network management
- End-user support
- Technical staff (including architecture and planning)
- General management and administration and other functions that are affected by the efficiency or otherwise of underlying system infrastructures

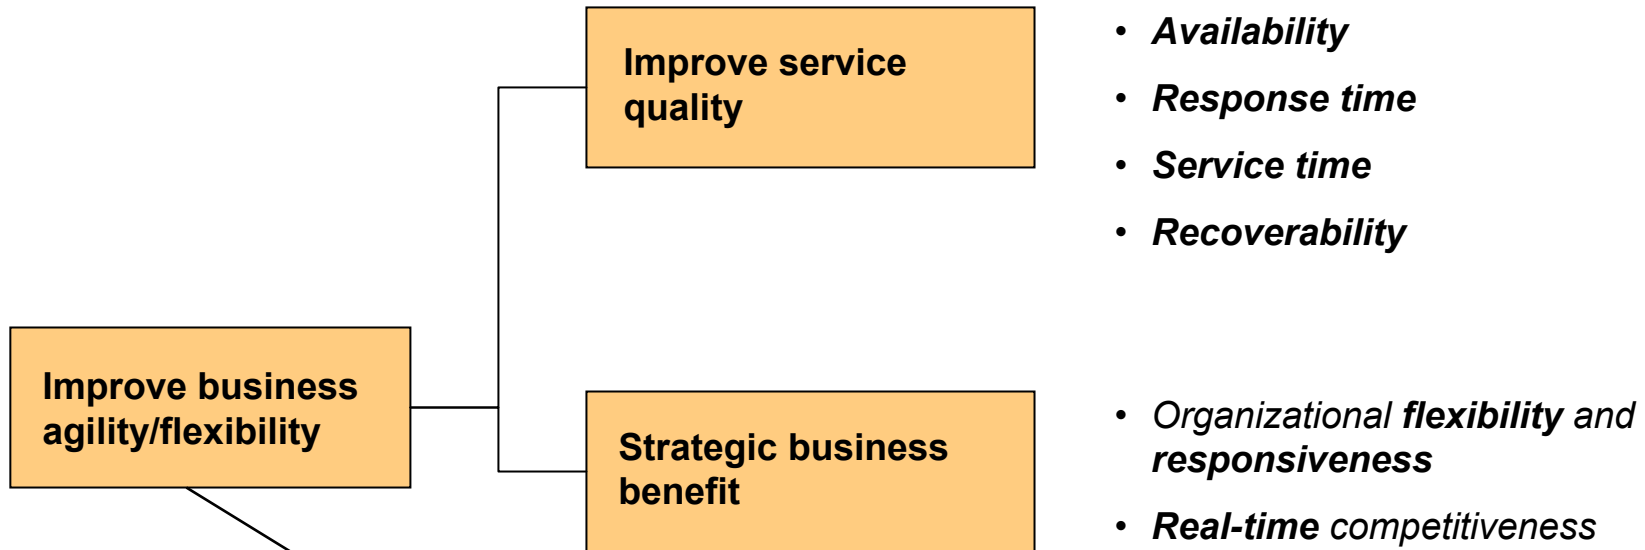

1 Sun UltraSPARC III v480 systems running @ 1.2 GHz; Sun UltraSPARC III F15K system running at 1.2 GHz;

2 HP PA-RISC rp5470 systems running @ 875 MHz; HP PA-RISC Superdome running at 875 MHz

3 Based on data from Project Scorpion

Source: IDEAS International; ITG; team analysis

# Communicating the benefits of improving business agility and flexibility



*...by dynamically and automatically re-allocating resources to applications as needed -- to better match changing business cycles or handle unexpected surge in demand.*

- Support multiple workloads/departments/applications
- Rapid response (10 ms) to changing workload needs via shared processor pool
- Automatic access to spare capacity via APV/Reserve CoD
- Exploitation of VLAN (e.g., high-speed communication)

Source: ITG; team analysis

# Addressing “Overhead FUD”

## Put it in the context of the value delivered:

- Intent is to drive-up throughput (utilization)
- Customers may not reach the "theoretical" maximum gains in throughput
- Equivalent of a luxury tax our competitors would LOVE to have

## “Sizing contingency”...NOT “overhead”


- rPerf and other measures are only an *approximate* indicator of performance between systems
- Uncertainties in workload peaks
- System effects due to Micro-partitioning

## Performance and price/performance

- eServer p5 systems performance are superior to anything offered by the competition
- Even after accounting for any performance efficiency, the eServer p5 systems still offer considerably more processing power
- Compelling story: do far more work with fewer processor resources at high CPU utilization

## Complex subject, simple rules-of-thumb are misleading

- With moderately-sized partitions on a system, "overhead" is not important
- Dependent on individual workload characteristics (cache, SMP, and I/O behaviors), the degree to which the peak demands of different partitions are concurrent, the total number of partitions, and the total number of virtual processors



## Customer scenarios – where APV is likely to have *significant* impact

NOT EXHAUSTIVE

### ***Server consolidation of many, small underutilized servers***

- Legacy AIX platforms (e.g., single app per server, fractional CPUs)
- Web/file/print (e.g., consolidation Linux, IA-32 to sub-CPU partitions)
- Application development/testing (secure, isolation)

### ***Responsive OS images in Shared Processor LPARs with higher utilization***

- Production (e.g., OLTP) with batch systems
- Overlapping production systems (e.g., peak of one; valley of another)
- Virtual blade environment (handle bursts of activity)

### ***Exploitation of Virtual I/O, VLAN***

- Multi-tier applications – (consolidation of web/application/database)
- Using virtual disk to consolidate relatively small boot images
- Consolidation of low-utilized Ethernet adapters

# Positioning/winning against Sun (DSD, N1 Grid Containers)

## What (today and planned)

- Dynamic System Domains (i.e., hard partitions)
- N1 Grid Containers with Solaris 10 (isolated execution environment within a single Solaris instance)

## Key positioning points

### Vs. Physical partitioning:

- Does not eliminate the possibility of system-wide outages; **availability** (e.g., Hypervisor on zSeries) more important than SPOF
- **Inferior technology** (e.g., limited flexibility, poor granularity, no resource sharing)
- Dynamic Processor Deallocation and Processor Sparing

### Vs. N1 Grid Containers:

- Lack of **flexibility** (e.g., no multi-OS, multi-version, multi-release level support, different change management procedures)
- Poor **fault isolation**: OS faults will bring down ALL containers
- Poor **security isolation**: Only user space level isolation; kernel level is still exposed
- (Possibly) **higher ISV software costs?**
- **TCO improvement is overstated**; costs to shift from per-partition to “per Container” level in TCO traditional TCO models

# Positioning/winning against HP (nPars, vPars, VSE)

## What (today and planned)

- nPars (i.e., hard partitions) on PA-RISC and IPF
- vPars (i.e., logical partitions) for HP-UX on PA-RISC; support for IPF servers other OS by 1H/2005?
- Secure Resource Partitions (similar to N1 Grid Containers) for HP-UX
- HP Virtual Server Environment (VSE) for HP Integrity servers:
  - gWLM (SLA-based resource management)
  - Systems Insight Manager (systems management)
  - vPar/VM (sub-CPU partitioning, dynamic adjustment of CPU, virtual memory, virtual I/O) projected for 2H2005 or 1H2006

## Key positioning points

### Vs. Physical partitioning:

- Does not eliminate the possibility of system-wide outages; **availability** (e.g., hypervisor on zSeries) more important than SPOF
- Dynamic Processor Deallocation and Processor Sparing
- Inferior technology (e.g., limited flexibility, poor granularity, no resource sharing)

### Vs. vPars:

- Only available for HP-UX on PA-RISC (**No support for IPF, Linux, Windows**)
- vPAR granularity of 1 CPU
- With movement of I/O slots and/or memory, the affected vPars **require a reboot**

### Vs. vPars/VM and VSE:

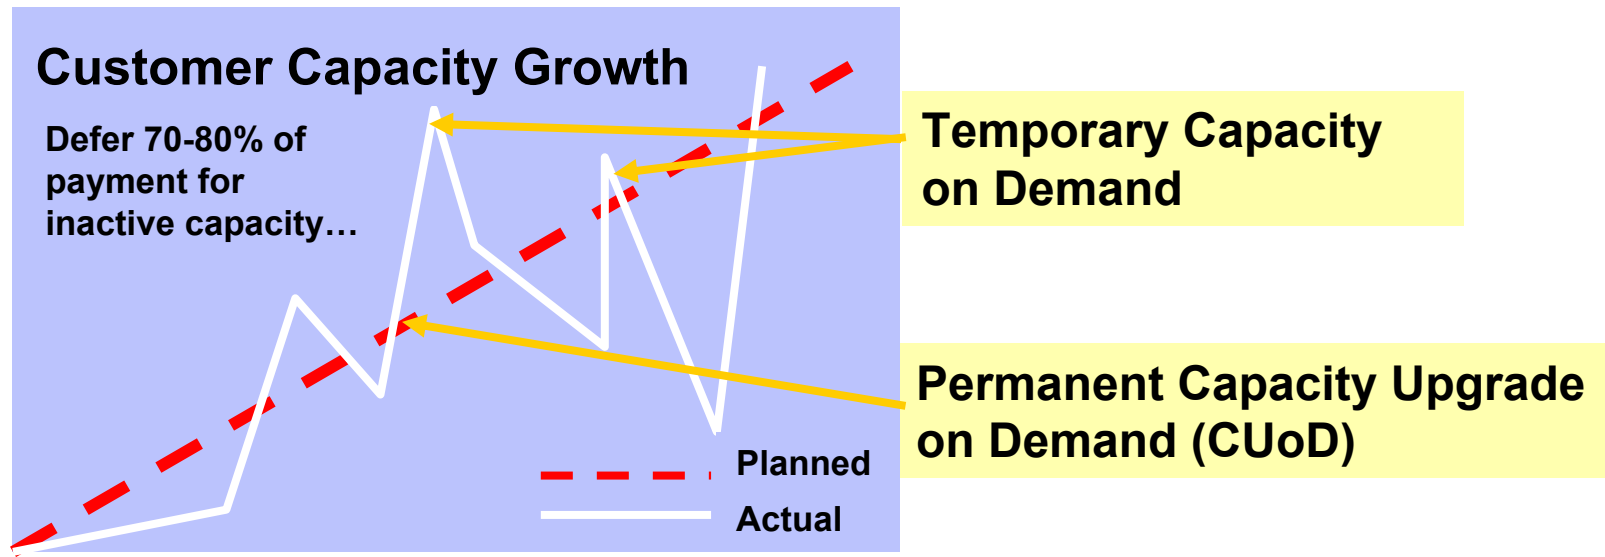
- vPars/VM is a **design goal** (2H2005 or 1H2006?)
- Initial release **limited**: max 2 CPU per VM; late with multi-OS support
- HP VSE is **not** part of a **cohesive strategy**



# Capacity on Demand (CoD) on Software Licensing on @server p5

# Why is On Demand Important to you?

- Server capacity you need, when you need it to help lower **total cost of acquisition**
- Address your planned IT capacity growth needs **quickly, non-disruptively**
- Built-in flexibility to address unplanned, **temporary spikes in demand**
- Increased configuration **flexibility**
- Increased **reliability** with processor and memory sparing
- **Deploy** new **services quickly** to react to changing business requirements
- Gain additional workload throughput and **automated systems performance leveling**



\*Capacity on Demand Features available on selected models of eServer p5

# What is different in POWER5 systems?

Feature	Description	POWER4	POWER5
<b>Permanent Processor Activation</b>	Key enabled, upgrade for inactive processors	✓	✓
<b>Permanent Memory Activation</b>	Key enabled, upgrade for inactive memory	✓	✓
<b>Reserve CUoD</b>	Prepay, debit account approach (currently called On/Off in POWER4)		✓
<b>On/Off CUoD</b>	Post pay, usage billing approach <sup>1</sup>	✓	✓
<b>Trial CoD</b>	30 day trial of all available CoD resources Additional Trial with processor CUoD activation (2 Processors and 4GB Memory)	✓	✓
<b>Capacity Backup</b>	Disaster recovery offering on high-end systems	✓	✓ <sup>2</sup>

1 POWER4 on/off CUoD refers to prepay, debit account approach

2 Currently statement of direction only

NOTE: Capacity on Demand Features available on selected models of eServer p5

## Permanent Capacity Upgrade on Demand (CUoD)

Pay for non disruptive growth available when you need it  
 Rapid deployment of new applications  
 No reboot necessary to utilize newly activated resources

**p5-595**



- ✓ Available on Entry, Midrange, and High End Servers
- ✓ Upgrade memory and processors in increments of 1GB memory and 1 processor
- ✓ Simple, secure, and fast activation
- ✓ No requirement to ever activate in the future
- ✓ No additional contracts or monitoring

**p5-550**



### When to sell:

- ✓ **Staged performance upgrades during server consolidation or new application implementation... helps lower cost of acquisition... pay as you grow**
- ✓ **Response to growing demand encumbrance**
- ✓ **Provide processor and memory sparing**

## Reserve Capacity on Demand

More than just prepaid processing power on demand

Capacity for Unpredictable Demands  
Event peaks - promotions, events, mergers, test...  
Low cost alternative to exploit inactive processors

- ✓ Price/Performance – low cost solution to smooth out intermittent, unpredictable capacity requirements available just when you need it without administrative intervention
- ✓ Significant performance throughput improvement when you place “inactive processors” in the shared processor pool
- ✓ Add processors to shared processor pool only
- ✓ Purchased in 30 processor day increments
- ✓ Requires AIX 5L V5.3 or Linux OS supporting Micro-Partitions and Advanced POWER Virtualization Feature

p5-595



p5-550




### When to sell:

- ✓ In virtualized environments to increase throughput and provide additional computing capacity when needed
- ✓ Response to unplanned, unpredictable capacity requirements without administrator intervention
- ✓ Help maintain end user service level agreements through peak capacity load requirements

## On / Off Capacity on Demand

For the predictable peaks in processing Power requirements **p5-595**

Repeating weekly, monthly, annual demand  
Planned Event peaks - promotions, events, etc.  
Temporary demand requirements



- ✓ Price/Performance – lower cost solution to cover predictable and planned peak processing requirements, available just when you need it for only as long as you need it... for memory and processors
- ✓ Activate processors and memory in dedicated or virtualized environments
- ✓ You precisely control how much resource is turned on, when it gets turned on, and when it gets released or turned off
- ✓ Data is collected by the system, transmitted to IBM and billed quarterly
- ✓ Only processor activation available on p5-550

**p5-550**



### When to sell:

- ✓ **Self managed program for temporary added memory and processors to handle additional capacity requirements**
- ✓ **For AIX 5L V5.2 or V5.3 in both dedicated and shared partition environments**
- ✓ **Help maintain end user service level agreements through peak capacity load requirements**

# Trial Capacity on Demand

One time, 30 days of no additional charge use  
Test out inactive capacity

- ✓ 30 processor day trial period for all the “inactive” memory and processors installed on the server
- ✓ No charge, just get the activation key from the IBM web site
- ✓ You can control the amount of resource you want to test, partial upgrades, or try all that you have installed
- ✓ After the test drive, then exploit one of the CoD upgrade offerings
  - Permanent Activation
  - Reserve CoD
  - On/Off CoD

**p5-595**



**p5-550**



## When to sell:

- ✓ Can be used for benchmarking and evaluating use of extra resources prior to purchasing permanent or temporary activations
- ✓ Can be used for a 30 day peak load requirement
- ✓ After each permanent processor activation, an additional 30 day trial is available with 2 processors and 4 GB of memory

# Software Licensing.... Positioning with POWER5 Systems

## –Not a drastic difference from today...

- Only difference is with products that **use processor-based** licensing models

## –IBM offers clients the flexibility...

- To continue to use software as they are doing so today and take advantage of the greater performance of POWER5 systems in a dedicated partition environment
- To exploit all the advantages of POWER5 virtualization, Micro-Partitions, VIO, PLM, and CoD to consolidate servers

## –To leverage Virtualization, customers are asking for the following requirements...

- Predictable software costs
- Licensing strategy that is fair and does not penalize for exploitation of virtualization
- Licensing strategy that is easy to reconcile and track

## –We are committed to a phased approach... this is the first step...

- Statement of Direction (SOD) for Sub-capacity licensing for IBM middleware, OS, and pSeries LPs\*
- Daily license entitlements for temporary processor activations for **IBM middleware**, OS, and pSeries LPs

## –We have an end game and strategy we are working on...

- Providing a license management tool that helps **Customers** plan and track software usage and entitlement requirements
- Eventual automation of software billing

## –Primary value of Virtualization is about much more than software licensing...

- Virtualization is about more efficient and flexible use of servers
- Providing higher performance with better relative TCO and TCA

\* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.



## IBM Software Licensing Direction\*

- **Requirements for Sub-Capacity licensing**
- **Requirement for Daily Licensing for temporary processor activations**
- **Requirement for tracking and planning software entitlement requirements**
  
- **Software Licensing with POWER5 Virtualization**
  - Announced Sub-Capacity licensing with OS and selected middleware (WebSphere Application Server for iSeries, HACMP)
  - IBM SOD to provide Sub-Capacity licensing on additional IBM middleware products
  
- **Software Licensing with CoD Temporary Processor Activations**
  - IBM announced On/Off Capacity on Demand for selected middleware software licensing in Oct 2003
  - IBM announced daily license entitlements for AIX and pSeries Licensed Program Products in Dec 2004
  
- **IBM Tivoli License Manager**
  - License Manager to be provided to customers for reporting and planning

\* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

# Sub-capacity licensing guidelines\*

## # of license entitlements

---

### Dedicated LPARs



- Based on **allocated CPUs**
- If allocations dynamic, use **peak allocations**

### Shared processor LPARs (capped)



- Based on **maximum entitlement** (processing units)
- **Round up** if needed (no sub-CPU licensing)

### Shared processor LPARs (uncapped)



- Based on **virtual processor**
- **Cannot exceed** the total number of physical processors in the **shared pool** for any given application

### CoD implications

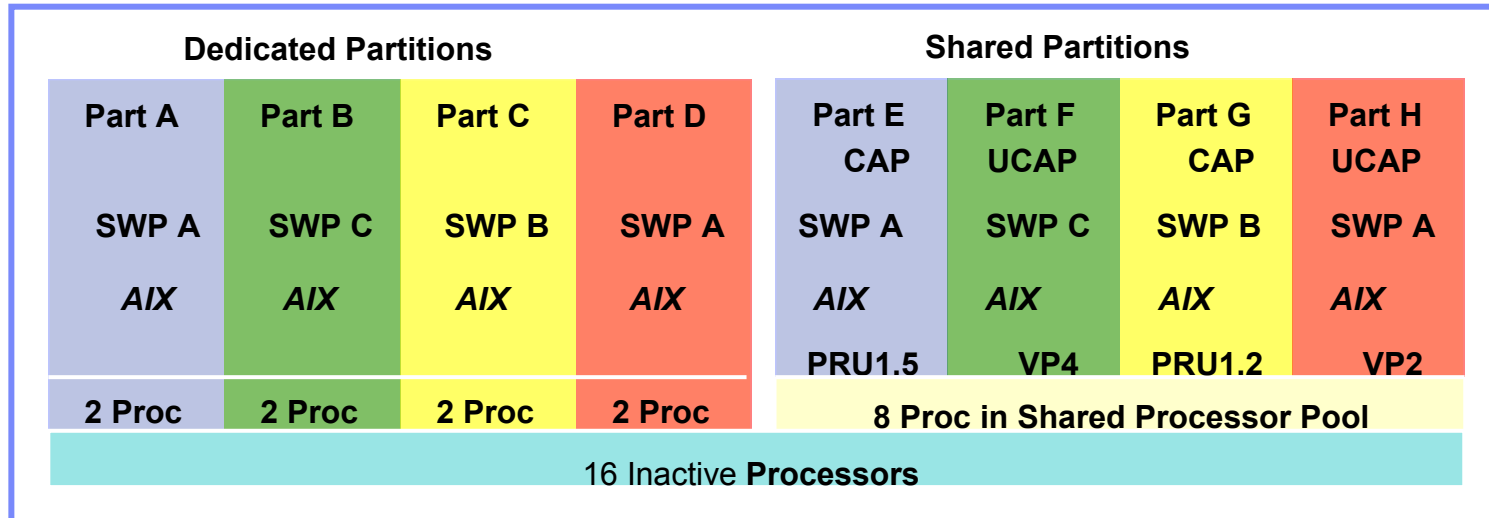


- For **Permanent** activation, requires **additional** license entitlements
- For **Reserve** CoD, the customer **prepays** for 30 days of both processors and OS entitlements
- For **On/Off** CoD, **IBM bills** the customer for both processor days and OS license entitlement days based on the customer's reported usage
- For **Trial** CoD, OS license entitlements are **included** in the trial

\* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

# POWER5 Virtualization and Software Pricing for OS and Middleware

32-way Server with 16 active processors, 16 inactive processors



Total Entitlements
AIX → 16
SWP A → 8
SWP B → 4
SWP C → 6

➤ **Software Licensing with POWER5 Virtualization – Dedicated Partitions**

- OS = total of 8 processors → 8 license entitlements
- SWP A = total of 4 processors → 4 entitlements
- SWP B = total of 2 processors → 2 entitlements
- SWP C = total of 2 processors → 2 entitlements

➤ **Software Licensing with POWER5 Virtualization – Shared Partitions**

- OS Capped Partitions = total of 4 processors, Uncapped Partitions = total of 6 Virt Proc, → 8 entitlements
- SWP A Capped Partitions = total of 2 processors, Uncapped Partitions = total of 2 Virt Proc → 4 entitlements
- SWP B Capped Partitions = total of 2 processors → 2 entitlements
- SWP C Uncapped Partitions = total of 4 Virt Proc → 4 entitlements

# Topic Summary

## Having concluded this topic, you should be able to:

- Describe IBM's Virtualization Engine (VE) technology and its components
- Describe the features and benefits of Advanced POWER Virtualization (APV) and CoD offering on @server p5
- Illustrate the customer uses of virtualization
- Communicating the value proposition of APV
- Identify the competitive advantages of IBM's virtualization offerings

# BACKUP

# PLM vs. POWER5 Hypervisor

PLM Differentiation	Capability	PLM	P5 PHYP
HW Support	POWER4 PLM automates DLPAR adjustment for P4 install base	X	
	POWER5	X	X
OS Support	AIX 5.2 PLM runs on AIX 5.2 on P4 and P5 systems (through PRPQ)	X	
	AIX 5.3	X	X
	pLinux		X
Physical Processor Management	Dedicated PLM runs on AIX 5.2 and/or P4 systems	X	
	Capped shared	X	
	Uncapped shared	X	X
Virtual Processor Management	Virtual processor minimization for efficiency	X	
	Virtual processor adjustment for physical processor growth	X	
Physical Memory Management	Share-based	X	
	Minimum and maximum entitlements	X	
Management Policy	Entitlement-based	X	X
	Goal-based		
	Application/middleware instrumentation required		
Management Domains	Multiple management domains on a single CEC	X	
	Cross platform (CEC)		
Administration	Simple administration	X	X
	Centralized LPAR monitoring (PLM command provides usage stats)	X	
	TOD-driven policy adjustment (PLM command supports new policy load based as TOD)	X	

# Competitive Landscape

Vendor	IBM					HP		Sun	Sun	Fujitsu	Unisys	VMware
Family	xSeries (x4x5)	pSeries	iSeries	zSeries		rp, rx and Superdome		E4900, E6900, E20K and E25K	Solaris based platforms	PRIMEPOWER 900, 1500 and 2500	ES7000	Intel based servers
Solution		LPARs	LPARs	LPARs		nPars	vPars	Domains	Containers	XPARs		
Implementation	HW	Hyper-visor	Hyper-visor	Hyper-visor	SW	HW	SW	HW	SW	HW	HW	SW
Operating system	Windows, Linux	AIX, Linux	OS/400, Linux	z/OS, z/VM, Linux	z/VM, z/OS, Linux	HP-UX		Solaris	Solaris <sup>1)</sup>	Solaris	Windows, SCO UnixWare	Windows, Linux, FreeBSD
Maximum number of partitions	2	20 - 160	32	30	> 10,000	2 - 16	4 - 128	2 - 18	2 - 18	8 - 15	8	Up to 64
Granularity (Minimum number of processors per partition)	2	One-tenth	One-tenth	Small fraction of one		4 (1)	1	2 - 4	Small fraction of one (sharing the same OS instance)	2 or 4	4	Small fraction of one
Dynamic reconfiguration	No	Yes	Yes	Yes		No	Yes	Yes	Yes	Yes	No	Yes
Shared resources	None	CPUs, I/O	CPUs, I/O	CPUs, memory, I/O channels		None	None	None	None	None	None	CPUs, memory, I/O
Interpartition communication	No	No	Yes	Yes	Yes	No	No	No	No	No	No	Yes


<sup>1)</sup> Solaris 10 and above

Source: Gartner, February 2003  
IBM, HP, Sun, Fujitsu, Unisys and VMware 2004

## Reserve CoD with Processor Pooling

### 16-Way Server

- 8 Active Processors Dedicated to Partitions A,B,C
- 8 Inactive CoD Processors



### Two Shared Partitions Started

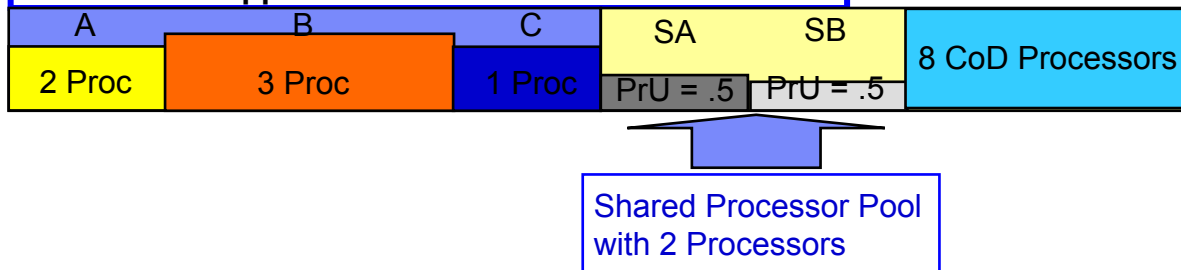
- ✓ Defined with partial processor entitlements
- ✓ Defined with VP = 2 to allow them to utilize both processors
- ✓ Workload from both partitions done on both processors in the shared processor pool

\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature




# Reserve CoD with Processor Pooling

**Administrator changes processor allocation**  
 -- allocates 2 processors to Shared Processor Pool  
 -- starts 2 uncapped DLPAR in Shared Processor Pool



Shared Processor  
Units Used



With the assignment of 2 permanent processors to the shared processor pool and starting 2 partitions, up to 2 processors units can be used

Two Shared Partitions Started

- ✓ Defined with partial processor entitlements
- ✓ Defined with VP = 2 to allow them to utilize both processors
- ✓ Workload from both partitions done on both processors in the shared processor pool


\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

# Reserve CoD with Processor Pooling


**The Administrator purchases a 30 day Reserve CoD Activation Feature**

- Enters activation key at HMC
- Allocates 8 CoD processors to the Shared Processor Pool
- Starts partition "SC" in the Shared Processor Pool

**30 Processor Days available**



Shared Processor Pool with 10 Processors



Reserve CoD Activated with 30 Processor Day Allocation

- ✓ 8 CoD processors assigned to the Shared Processor Pool... total of 10 processors in the Shared Processor Pool
- ✓ Workload from partitions SA, SB, SC spread across all 10 processors... increased throughput for these partitions
- ✓ So long as shared partition workload remains below 2 processor units, no deduction from Reserve CoD Activation Allocation


When the Reserve CoD activation key is entered and 8 CoD processors are assigned to the shared processor pool, total of 10 processors available in the shared processor pool

\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

# Reserve CoD with Processor Pooling


**Workload Increases by More Than 1/10 of a Processor Unit**

- One processor day deducted from the Reserve CoD
- Workload returns below 2 processor units within the 24 hour period from the time it exceeded 2 processor units



Shared Processor Pool with 10 Processors

**29** Processor Days available



Reserve CoD – Shared Partition Workload Exceeds 2 Processor Units


- ✓ When workload exceeds 2 processor units by 1/10, one processor day deducted from Reserve CoD activation allocation
- ✓ Continue to use additional one processor unit for 24 hours... no additional deduction from allocation
- ✓ Ending processor day allocation is 29 days

Workload exceeds 2 processor units by 1/10 of a processor unit... one processor day deducted


\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

# Reserve CoD with Processor Pooling

**Workload in Next Peak Increases by 4 Processor Units**  
 -- Additional processing requirements automatically handled  
 -- Workload of 4 additional processor units continues for 36 hours



**21 Processor Days available**



Reserve CoD – Shared Partition Workload Exceeds by 4 Processor Units

- ✓ New peak requirement handled with additional processor units
- ✓ Workload continues for 36 hours... 8 processor days deducted, 4 processor units for the first 24 hour period, and 4 processor units for the second 24 hour period
- ✓ Ending processor day allocation is 21 days

Workload exceeds 2 processor units by 4 processor units... 4 processor days deducted for each 24 hour period utilized


\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

# Reserve CoD with Processor Pooling

**Reserve CoD Activation Allocation Reduced to Below Number of Processors in the Shared Processor Pool**

-- Processors are removed from the shared processor pool

-- Message sent to HMC



7 Processor  
Days available

**HMC**

1 processor being removed from the shared pool

Reserve CoD Activation Allocation Reduced Below 10 Processor Days

- ✓ Processors are removed from the shared processor pool
- ✓ Message sent to HMC to alert Administrator
- ✓ Continue to remove processors from the shared processor pool until the allocation is zero... then only permanently active processors in the shared pool


\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

# Reserve CoD with Processor Pooling

**Reserve CoD Activation Allocation Reduced to Below Number of Processors in the Shared Processor Pool**

-- Processors are removed from the shared processor pool

-- Message sent to HMC



**0 Processor**  
Days available

**HMC**

**1 processor being removed from the shared pool**

- Reserve CoD Activation Allocation is Zero
- ✓ All 8 CoD processors are removed from the shared processor pool
  - ✓ Message sent to HMC to alert Administrator as each is removed
  - ✓ Shared partitions continue to run, but with only the remaining 2 permanent processors in the shared processor pool
  - ✓ A partition will never be stopped because the CoD processors are being removed as the Reserve CoD activation allocation is reduced to zero

\*Reserve CoD may only be activated with AIX 5L V5.3 and requires the Advanced Power Virtualization hardware feature

## Software Group SOD on Sub-Capacity Pricing\*

Currently, IBM provides limited sub-capacity pricing for selected products. As part of our July 2004 IBM eServer p5 and eServer i5 announcement, IBM has provided the following Statement of Direction:

To continue to lead the industry in delivering On Demand Business solutions, IBM intends to provide flexible pricing and financing which enable our customers to exploit new technology. IBM has already taken important steps in this area with our announcement of the complementary On/Off Capacity on Demand offerings for hardware and software.

IBM considers sub-capacity licensing to be a strategic offering within the on demand portfolio. Sub-capacity licensing enables the licensing of a software product for use on less than the full capacity (in processors) of the machine when used within a logical partition. IBM currently offers sub-capacity licensing for selected software products, including WebSphere Application Server running on iSeries.

Within the next 12 months, IBM intends to announce sub-capacity licensing terms and conditions for additional middleware products, such as WebSphere Application Server and DB2 Enterprise Server Edition, on other platforms, such as AIX, i5/OS, OS/400, and Linux on POWER. IBM also plans to deliver license management tools to assist customers in tracking and reporting their software assets to take advantage of these flexible pricing models. Please contact an IBM Software Representative if you need further information.

\* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

# End of Presentation



## Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml): AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM, Micro-Partitioning, Hypervisor

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation  
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries  
UNIX is a registered trademark of The Open Group in the United States and other countries.  
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.  
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.  
Intel is a registered trademark of Intel Corporation  
\* All other products may be trademarks or registered trademarks of their respective companies.

## NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.