IBM Systems and Technology Group University 2005

ON DEMAND BUSINESS™

IBM Systems and Technology Group University 2005

Partitioning Comparison: zSeries, pSeries, VMware

Session ZT35

Reed A. Mullen
Programming Consultant

Learning Objectives

- **At the conclusion of this material, you should be able to:**
  - Articulate the business value of partitioning and virtualization
  - Help customers understand the business and technology issues that need to be considered when selecting a partitioning solution for virtual server deployments
  - Describe the unique partitioning and virtualization benefits offered by the zSeries platform

Agenda

- **Partitioning, the Mainframe Charter, and on demand Computing**
- **Server partitioning basics and business value**
- **Platform overview of partitioning technology**
  - POWER5
  - Intel (VMware)
  - zSeries (z/VM)
- **zSeries partitioning differentiation**
- **Conclusion and wrap-up**

IBM Systems and Technology Group University 2005

Partitioning, the Mainframe Charter, and on demand Computing

This educational piece is intended for your use in selling. It is NOT a deliverable for your customers.

The Mainframe Charter - Providing a Strategic Framework

IBM is committed to delivering innovative solutions to meet our customers' on demand business requirements

It is our intention to continue to:

Innovation

- **Provide leadership in innovation to enhance the use of IBM eServer zSeries to support increasingly integrated and flexible business processes for the on demand business.**
- **Maintain zSeries' position as a benchmark for flexible, efficient, and responsive platforms for highly complex, integrated environments running a wide range of mission-critical workloads.**
- **Improve the autonomic and self-managing capabilities of the zSeries while working to simplify user processes and system administration tasks.**

Value

- **Enhance the value proposition and lower the cost of computing of zSeries solutions in a way that is compelling, clear, and consistent.**
- **Extend the on demand characteristics of zSeries servers, highlighting its strengths as an environment for usage-based computing.**
- **Increase the ability to account for allocation and use of zSeries resources in an on-demand environment.**

Community

- **Support programs designed to foster vitality in the zSeries community, helping to promote a strong application portfolio and world-class support services.**
- **Provide the skills and expertise to assist customers in designing, developing, and deploying on demand solutions built on a foundation whose cornerstone is zSeries.**
- **Leverage key open standards and common structures to enhance the use of zSeries in large, heterogeneous environments.**

These principles help guide IBM's investment priorities in zSeries systems today and far into the future and demonstrate IBM's commitment to provide value to its zSeries customers.

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

Mainframe Charter: Delivering New On Demand Capabilities

z/VM is a Key Component of the Mainframe Charter

Innovation

Proof Points

Provide leadership in innovation to enhance the use of IBM eServer zSeries to support increasingly integrated and flexible business processes for the on demand business.

- Integrated Facility for Linux (IFL)
- zAAP for Java workloads
- Hi-speed HiperSockets connectivity
- GDPS/PPRC multiplatform resiliency
- Linux Communications Server
- Enhanced z890 and z990 scalability

Improve the autonomic and self-managing capabilities of the zSeries while working to simplify user processes and system administration tasks.

- eWLM for z/OS and Linux (future)
- IBM Director Multiplatform (future)
- New z/VM publication: "Getting Started with Linux on zSeries"

Value

Proof Points

Enhance the value proposition and lower the cost of computing of zSeries solutions in a way that is compelling, clear, and consistent.

- New z/VM V5 pricing model
- Capacity on Demand offerings
- IT Optimization Solution offering
- End-to-End Systems Management Services
- Offerings for Resiliency & Security (e.g., GDPS/PPRC Implementation Services)
- 12-step Strategic Virtualization Assessment offering
- z/OS sub-capacity pricing for key platform software (also adopted by BMC & CA)
- Application transformation and integration services

Community

Proof Points

Support programs designed to foster vitality in the zSeries community, helping to promote a strong application portfolio and world-class support services.

- A growing list of new ISVs and applications
- zSeries training (IBM Learning Services, workshops, conferences)
- Support for University programs
- Linux Community Development System (LCDS)

Leverage key open standards and common structures to enhance the use of zSeries in large, heterogeneous environments.

- z/VM support for FCP/SCSI disks
- Virtual Server Management APIs
- OSA and z/VM virtual network switching (Layer 3 and Layer 2)

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

The On Demand Operating Environment

Think of it as a Global Fabric

Where everything enterprise wide becomes:

- Reusable, connected, and integrated…
- Resilient and highly secure…
- Scaleable and responsive..
- Simplified and optimized…
- Virtualized and cheaper to run…
- Unified and flexible to support the business model

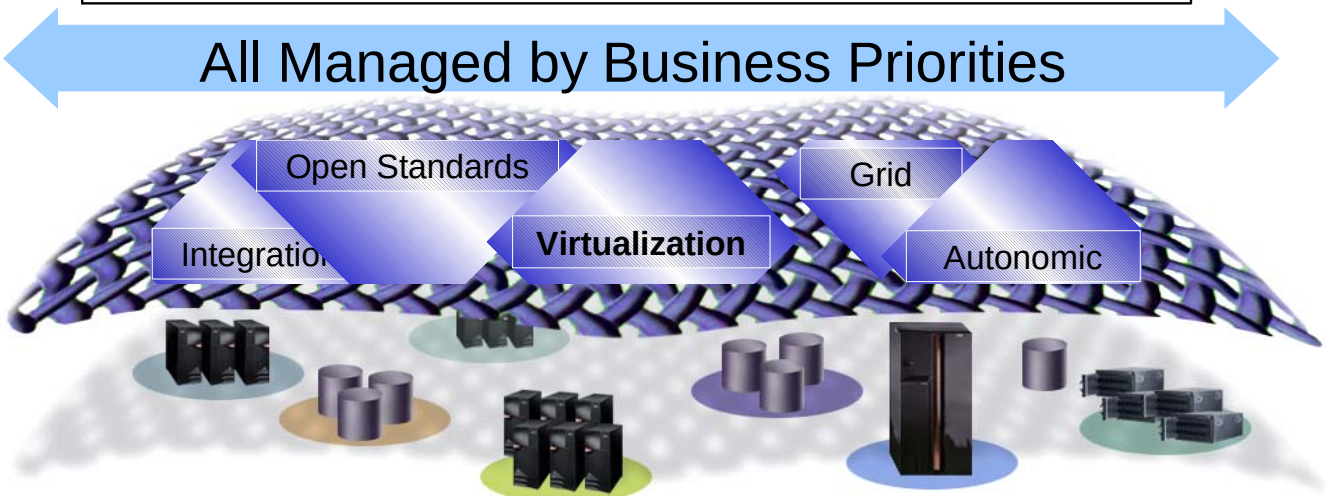

Enhance Your Global Fabric

By Leveraging Your Most Capable Asset: Your Mainframe

As our customers understand the IT requirements for on demand business, they cite a strong synergy to zSeries capabilities

IBM's vision is to leverage zSeries leadership capabilities around:

- ★*Resource Virtualization*
- ★*Business Resiliency and Security*
- ★*Intelligent Workload Management*
- ★*Business Integration*

IBM Systems and Technology Group University 2005

Server Partitioning Basics and Business Value

This educational piece is intended for your use in selling. It is NOT a deliverable for your customers.

Server Partitioning and Virtualization

An Important Consideration for All Areas of Server Design and Deployment

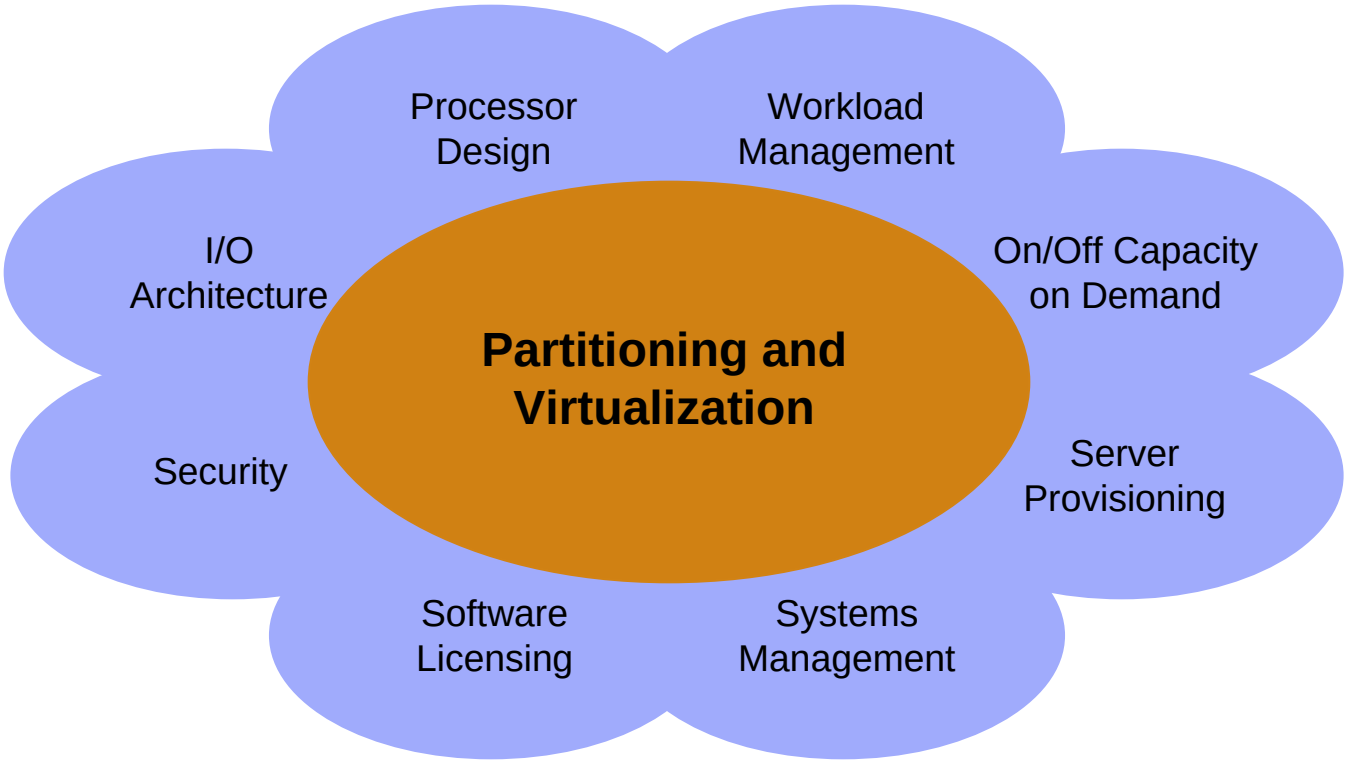

Partitioning Defined

- **Partitioning is the division of a single server's resources* into multiple, independent, isolated systems capable of running their own operating system**
- **Three types of partitioning:**
  - *Hardware* – resources are allocated to partitions on a one-to-one basis with the underlying physical hardware (no sharing among partitions)
  - *Logical* – resources are managed by hardware firmware and allocated to partitions with a finer granularity than hardware partitioning (resource sharing among partitions)
  - *Software* – resources are managed by a software layer, aggregated into shared resource pools, and apportioned to users as *virtual* system resources, separating the presentation of the resources from the actual physical entities

* Resources include: processors, memory, I/O adapters and devices, networking interfaces, co-processors

Hypervisor Technologies

“Trapping and Mapping” method

- Guest OS runs in user mode
- Hypervisor runs in privileged mode
- Privileged instructions issued by guest operating system(s) are trapped by hypervisor
- IA-32 complications:
  - Some instructions behave differently in privileged and user modes
    - For example, “POPF” treatment of the interrupt enable flag
  - User mode instructions that access privileged resources/state cannot be trapped; instruction must be changed to something that can be trapped
- Some guest kernel binary translation may be required
- Originally used by mainframes in 1960s and 1970s (VM/370)
- Used today by VMware

Hypervisor Technologies – Trapping and Mapping

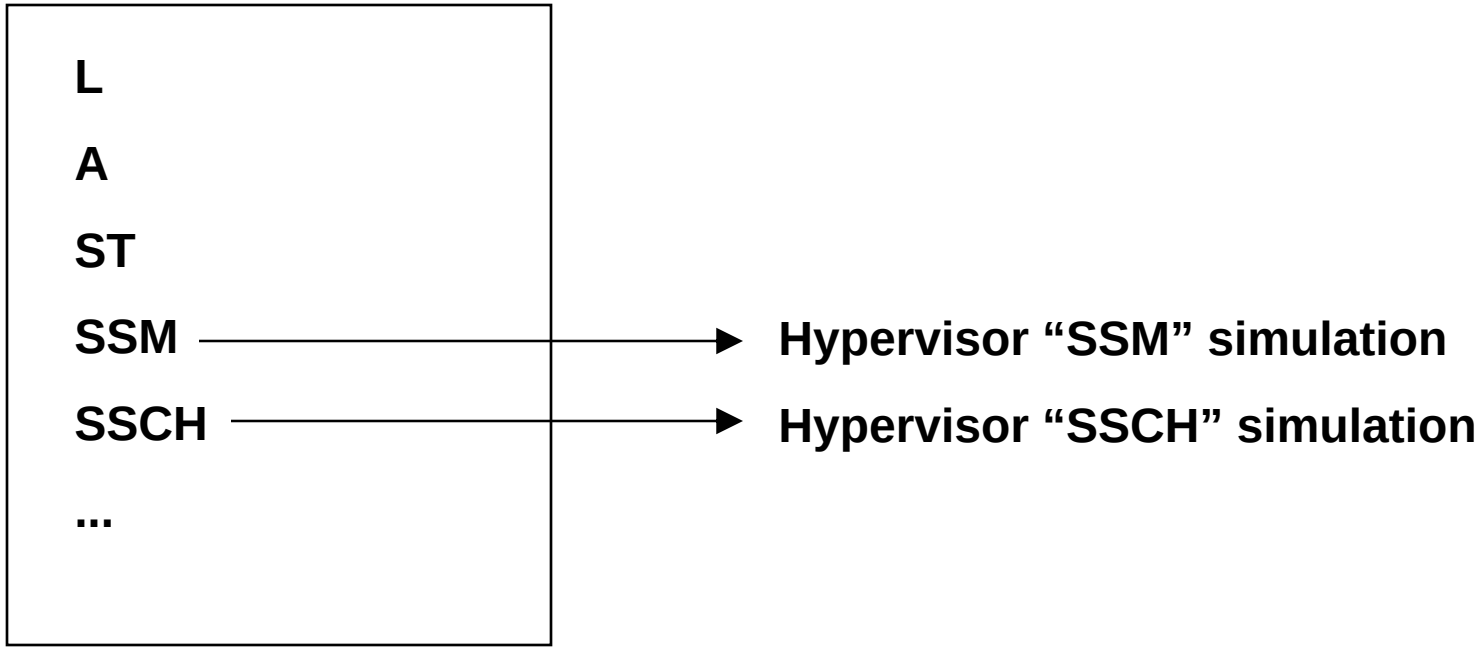

Simulation of guest instructions may create timing problems for software running in the guest environment.

Hypervisor Technologies

Hypervisor Call method

- Guest OS runs in privileged mode
- Hypervisor runs in super-privileged mode
- Guest OS kernel (e.g., AIX, i5/OS, Linux) is modified to do hypervisor calls for I/O, memory management, yield rest of time slice, etc.
- Memory mapping architecture is used to isolate guests from each other and to protect the hypervisor
- Used by POWER5 today

Hypervisor Technologies – Hypervisor Call

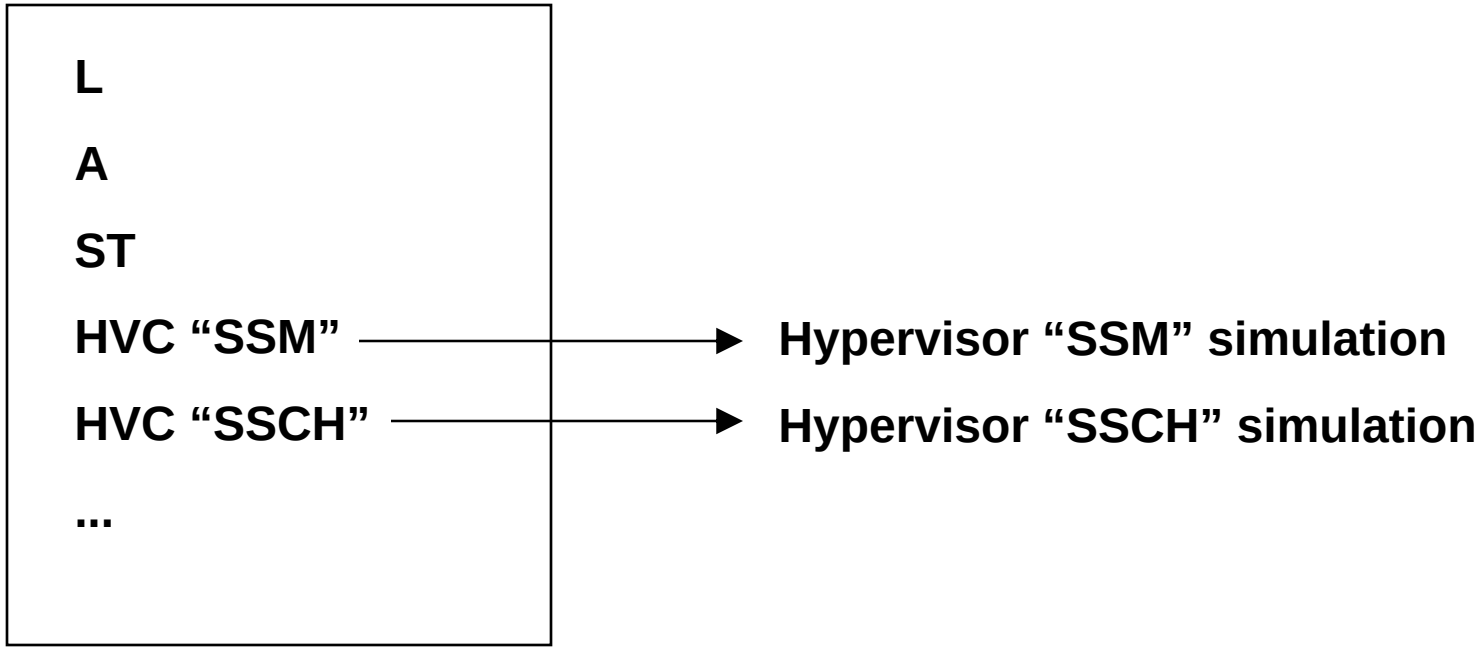

Hypervisor Technologies

Direct Hardware Support method

- Guest OS runs in privileged mode
- Guest OS can be run unmodified, but can issue some hypervisor calls to improve performance or capability
  - I/O (z/VM)
  - Yield time slice (PR/SM™ and z/VM)
- Extensive hardware assists for hypervisor (virtual processor dispatching, I/O pass-through, memory partitioning, etc.)
- Used by zSeries (PR/SM™ and z/VM)

Hypervisor Technologies – Direct Hardware Support

L

A

ST

SSM

SSCH

...

Hardware processes guest instruction with no hypervisor simulation

What Business Value Does Partitioning Offer?

- **Higher utilization rates for hardware assets**
  - Share the CPUs, I/O adapters, disk drives, network cards, etc.
  - Get more work done with the same amount of resources
- **Less time required to create and configure (or reconfigure) server images**
  - Respond to business conditions *now*, not later today, tomorrow, or next week
- **Workload management at the server level**
  - Control resource allocations at the server level
  - Nicely complements workload management at the application level
- **Helps save money on software fees**
  - Software licensing is based on real CPU capacity, not the number of server images

Partitioning Technologies* Designed to Add Business Value

- **Highly granular allocation of hardware assets**
  - Add another server without impacting other server images
- **Large-scale server hosting**
  - Run more than a thousand server images
- **Resource consumption recording / reporting**
  - Capture data at the hypervisor level
  - Useful for charge-back, capacity planning
  - Problem determination *and* fix verification
- **Hot stand-by without the hardware expense**
  - Idle backup image(s) ready to run (or be booted) if primary server(s) fail
- **Transfer a running virtual machine to another real server**
  - Another option for failover / zero-downtime
- **Autonomic disk failover to back-up storage subsystem**
- **Architecture simulation**
  - Satisfy server configuration requirements without suffering expense of real hardware
- **In-memory application sharing**
  - Share program executables among multiple server images
  - Files stored in hypervisor memory and accessed via Linux execute-in-place file system
- **Server-memory-cached disk I/O**
  - High-speed read access to files on disk
- **Virtual Disks in Storage**
  - High-speed read and write access to files in memory (excellent swap devices for Linux)
- **Built-in console message routing**
  - Route messages from all virtual servers to a single virtual machine
  - Deploy system automation solutions that interrogate and respond to console messages
- **“Hands free” auto-logon of server images**
- **Initiate operating system shutdown from “outside” the server image**
  - Without requiring agent code running on guest operating system

* Not all platforms offer these technologies

IBM Systems and Technology Group University 2005

Virtualization Platform Technology Overview

eServer POWER5 Systems Virtualization Architecture

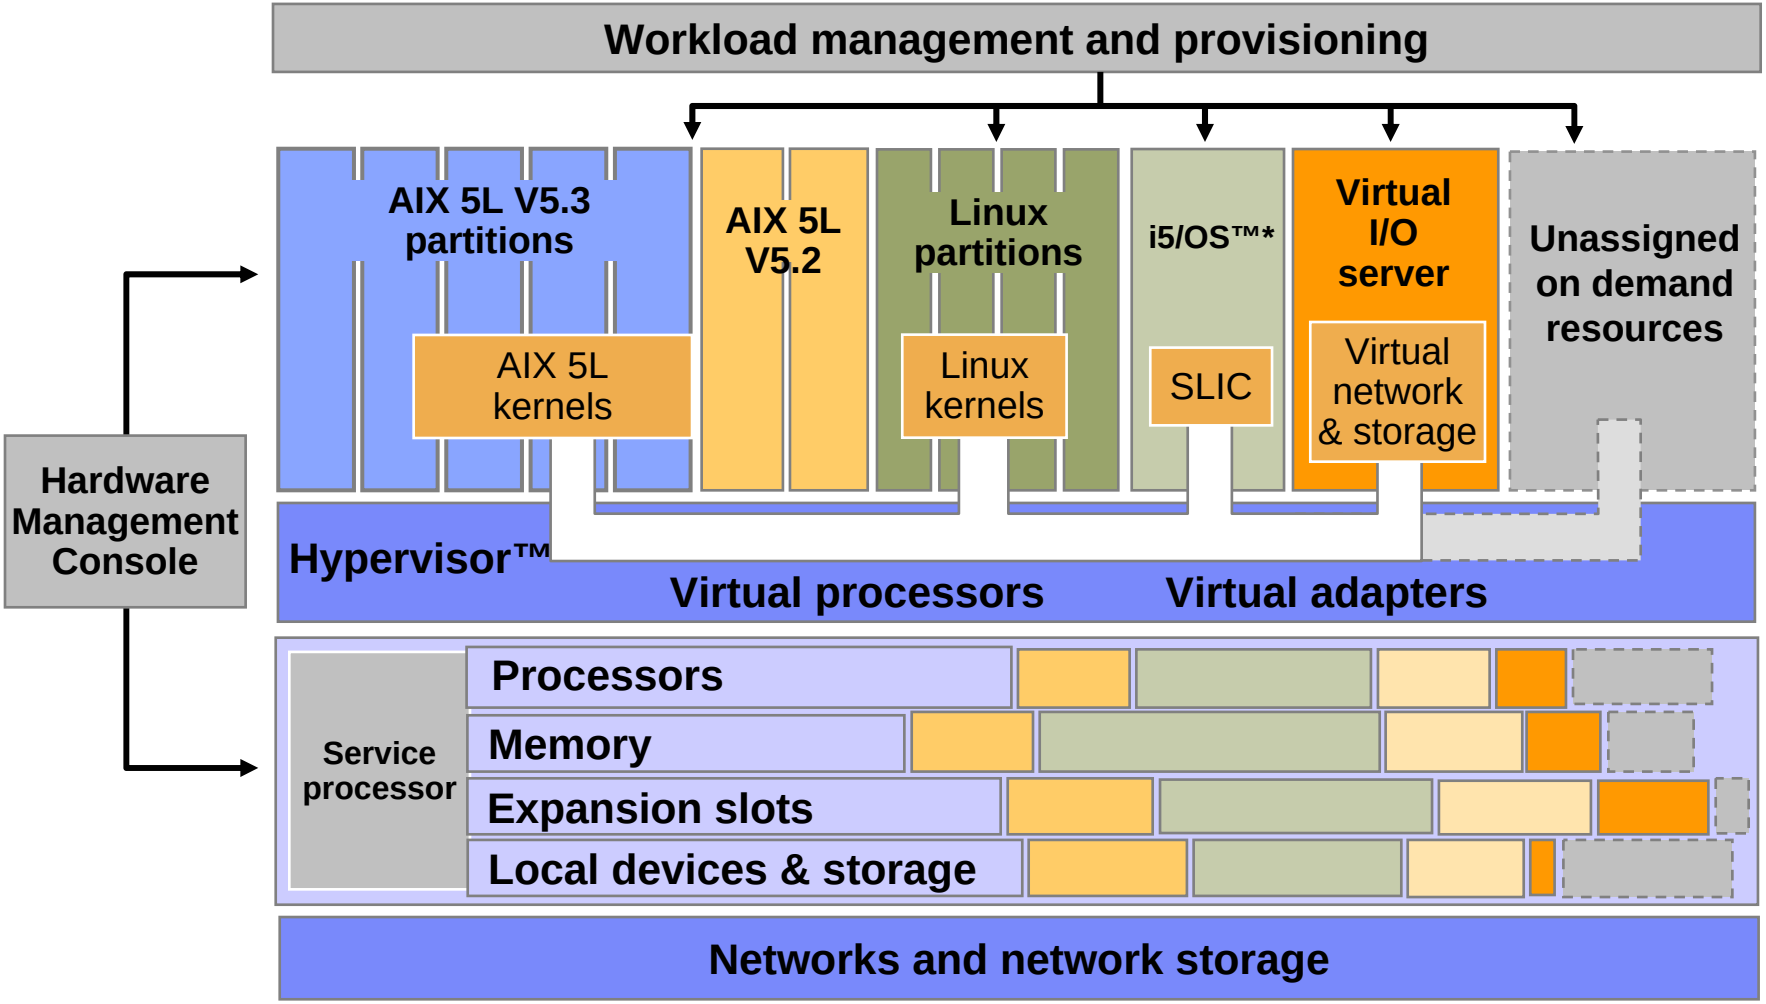

*Available on 1.65 GHz POWER5 p5-570, p5-590 and p5-595 models

eServer POWER5 Advanced Virtualization Options

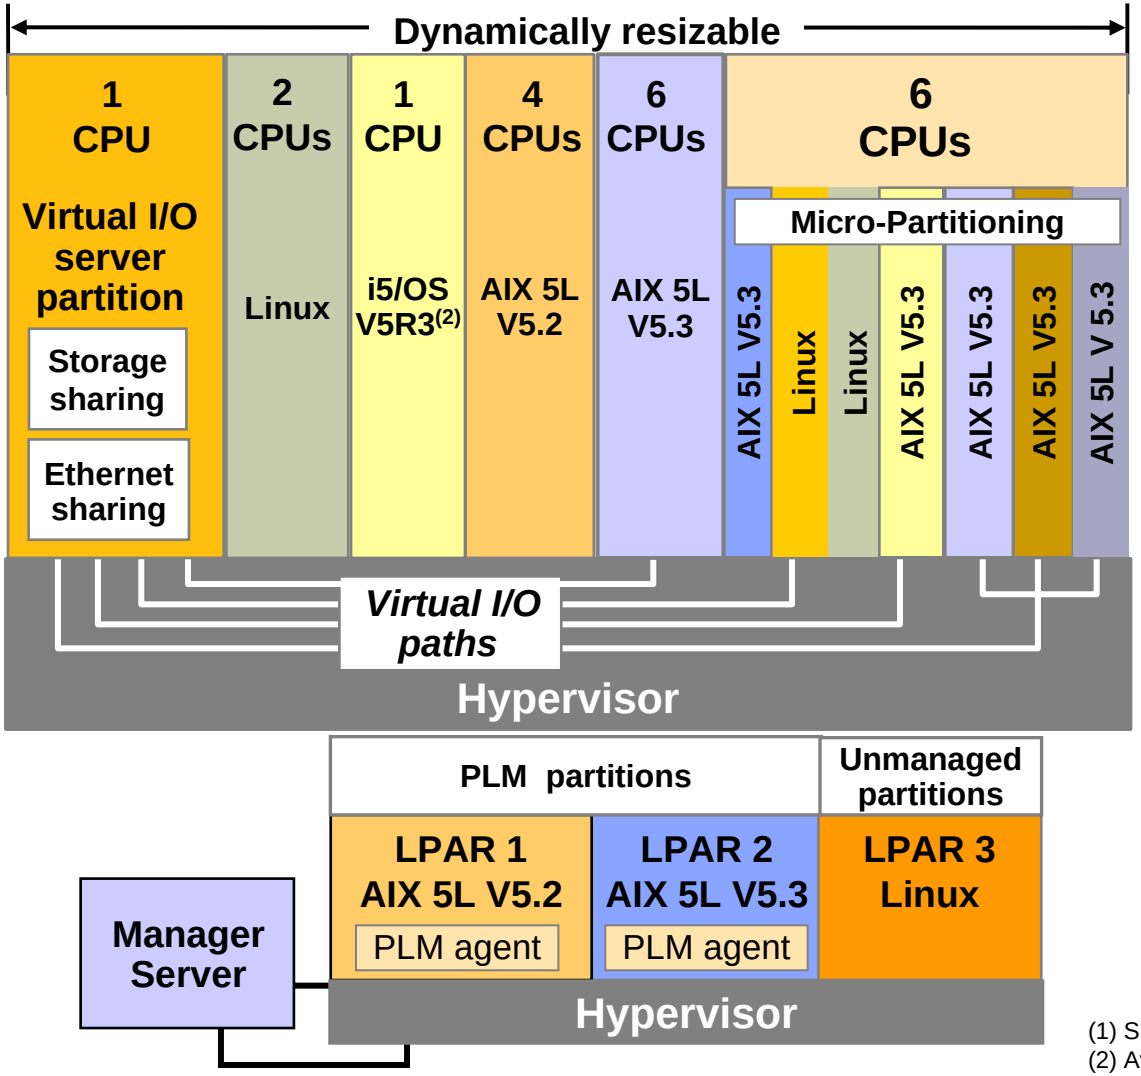

Virtual I/O server

- Shared Ethernet
- Shared SCSI & Fibre Channel attached disk subsystems
- Supports AIX 5L V5.3 and Linux(1) partitions

Micro-Partitioning

- Share processors across multiple partitions
- Minimum partition size: 1/10th of one CPU
- AIX 5L V5.3, Linux(1), or i5/OS(2)

Partition Load Manager

- Cross-partition workload management for AIX 5L V5.2 and V5.3
- Balances processor and memory requests

Managed via HMC

(1) SLES 9 or RHEL AS3
(2) Available on 1.65 GHz POWER5 p5-570, p5-590 and p5-595 models

eServer POWER5 Virtual I/O Server Disk Sharing

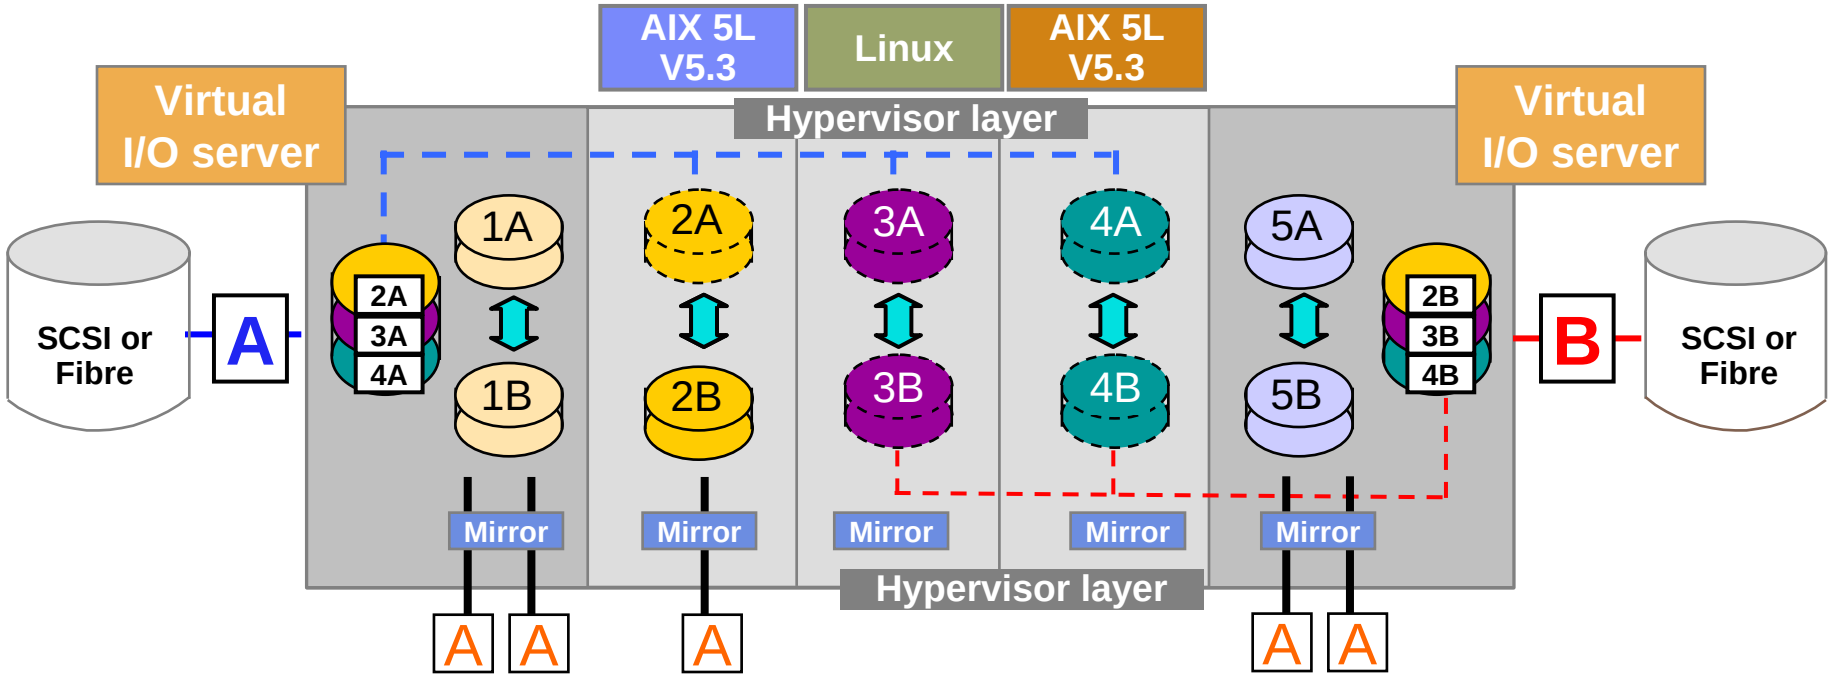

- One physical drive can appear to be multiple logical drives
  - LUNs appear as individual logical drives
  - LAN type latency on I/O operations
- Minimizes the number of adapters
- Can have mixed configuration (virtual and real adapters)
- SCSI and Fibre supported
- Supports AIX 5L V5.3 & Linux partitions

eServer POWER5 Virtual I/O Server Ethernet Sharing

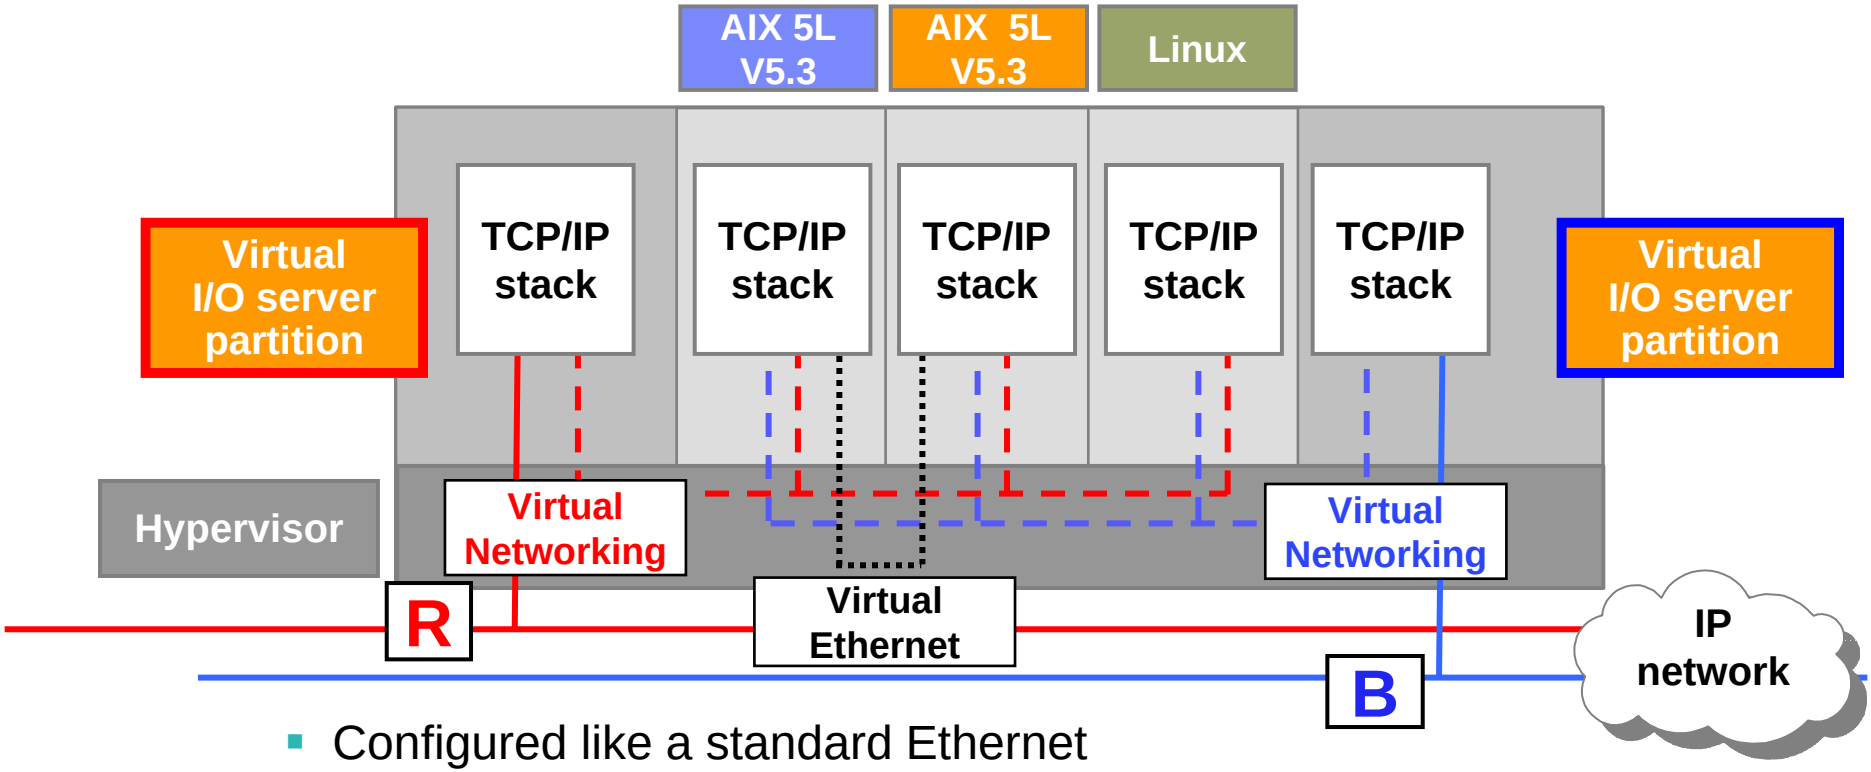

- Configured like a standard Ethernet
- IP forwarding / bridging provided by I/O server partition
- Can have multiple connections per partition
- Virtual “MAC” addressing
- Each adapter can support up to 16 virtual Ethernet LANs

eServer POWER5 Micro-Partitioning

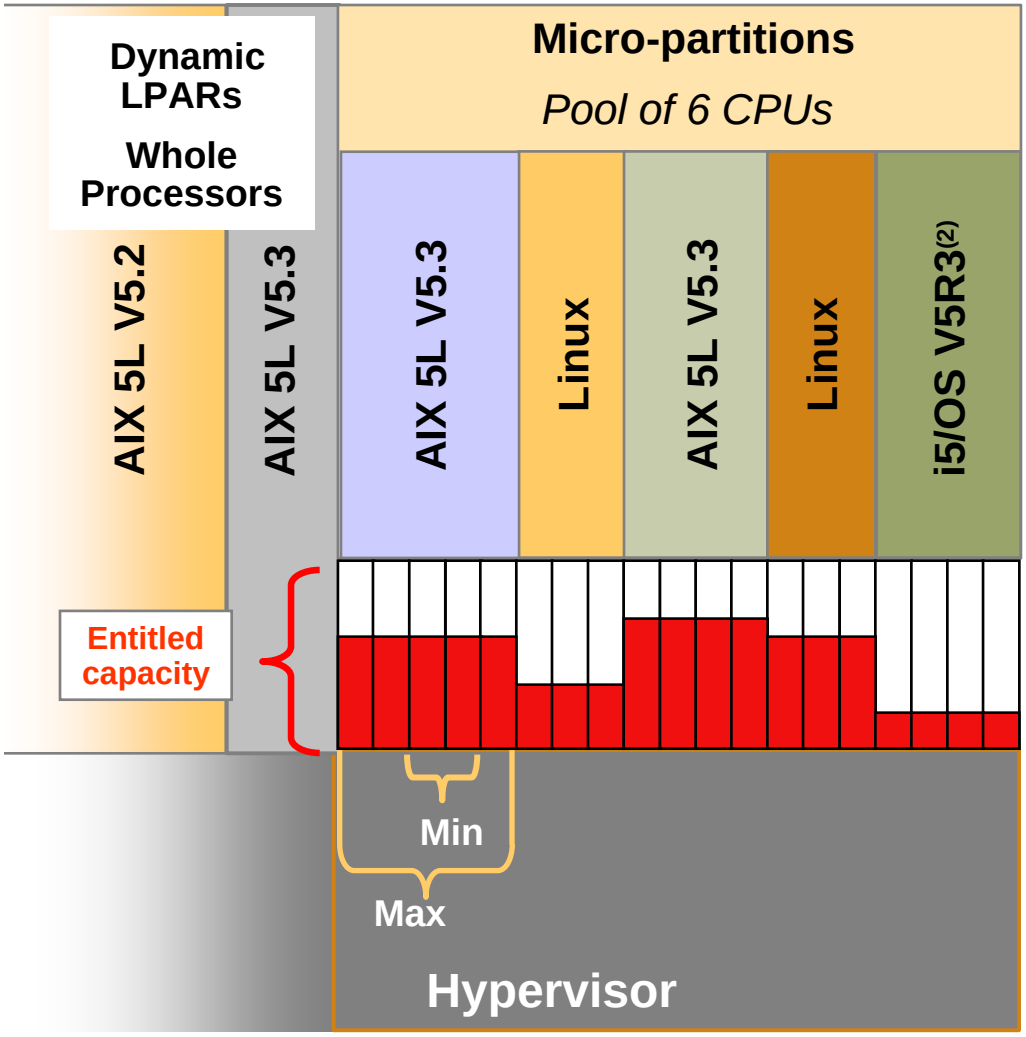

Increased number of LPARs
- Micro-Partitions: Up to 254(1)
- Dynamic LPARs: Up to 32(2)

Configured via the HMC

Number of logical processors
- Minimum/maximum

Entitled capacity
- In units of 1/100th of a CPU
- Minimum: 1/10th of a CPU

Variable weight
- % share (priority) of surplus capacity

Capped or uncapped partitions

(1) On p5-590 and p5-595 with a minimum of 26 active processors
(2) Available on 1.65 GHz POWER5 p5-570, p5-590 and p5-595 models

Partition Load Manager for AIX 5L eServer p5

- **Policy-based, automatic partition resource tuning**
- **Dynamically adjust CPU and memory allocation**

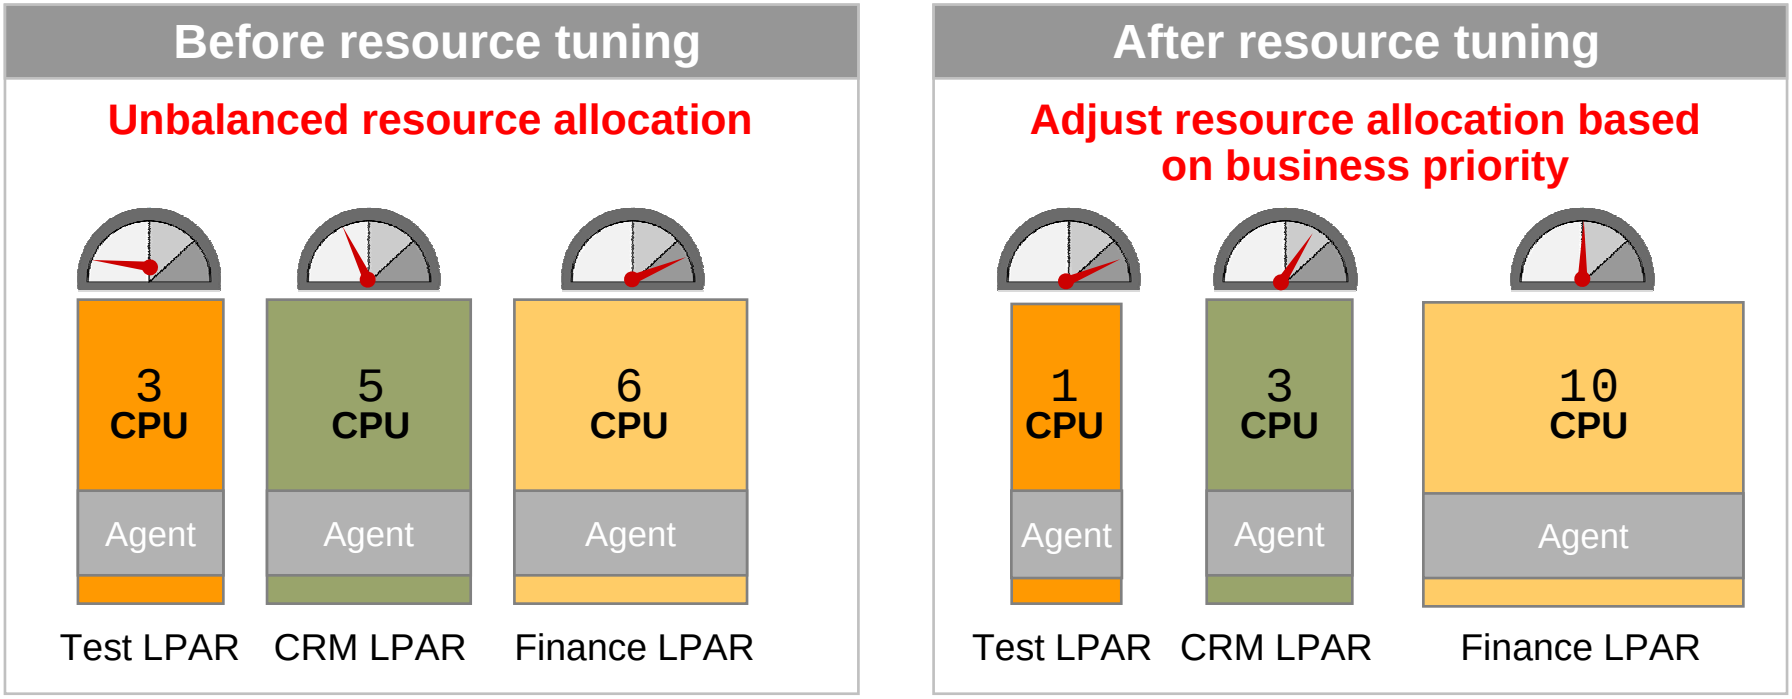

Put the Management Power of AIX 5L to Work

Designed to automatically handle day-to-day activities and empower you to make better infrastructure decisions

Workload management features are designed to free staff from repetitive activities

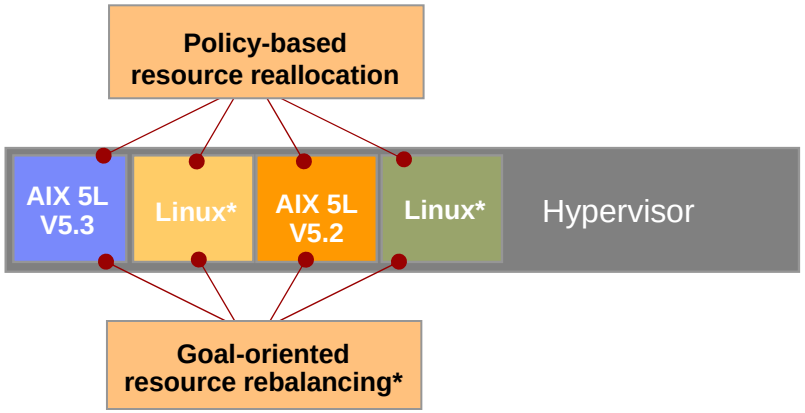

PLUS: New Advanced Accounting Features are designed to give management key data for improved resource utilization and planning

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

* Future support planned

Enhanced Infrastructure Simplification with eServer i5

- **Integrated xSeries Server (IXS) and Integrated xSeries Adapter (IXA) allows eServer i5 systems to host Windows workload in conjunction with i5/OS, Linux, and AIX 5L**
  - Consolidate Intel-based Windows and Linux servers with i5 infrastructure
  - Extend your Linux application options to include the Intel architecture

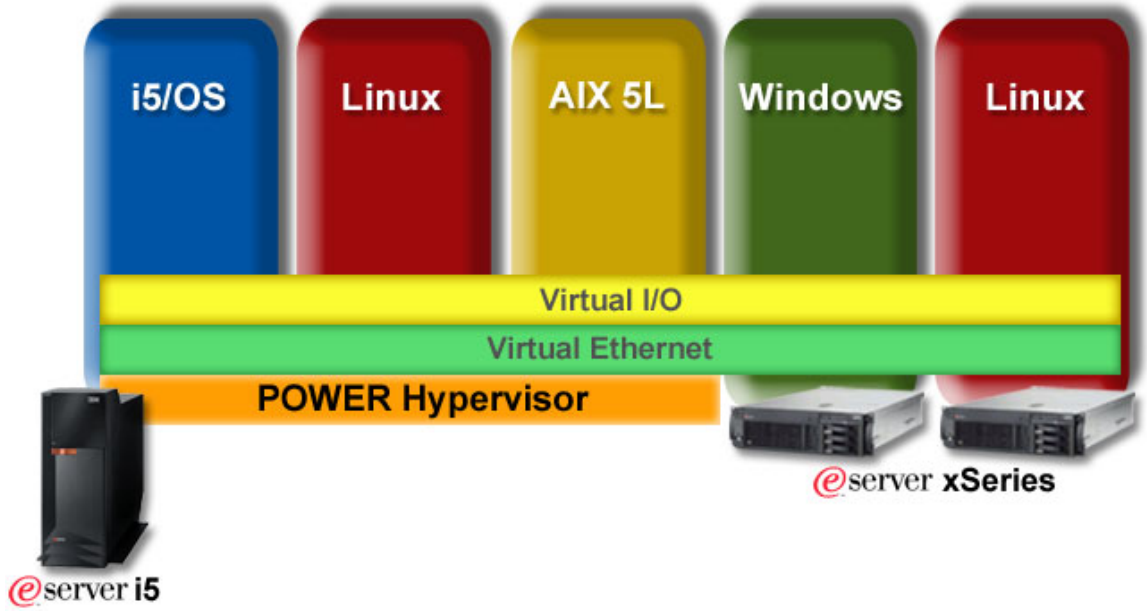

Intel Partitioning with VMware ESX Server

- **VMware ESX runs directly on the hardware (no hosting operating system required)**
- **Creates multiple virtual machines on a single Intel system**
- **Manages resource allocations**
- **Strong fault and security isolation (uses CPU hardware protection)**
- **Shared data cluster-ready**
- **Scalable to large virtual machines and high performance**
- **Direct I/O passthru**
- **Encapsulation**

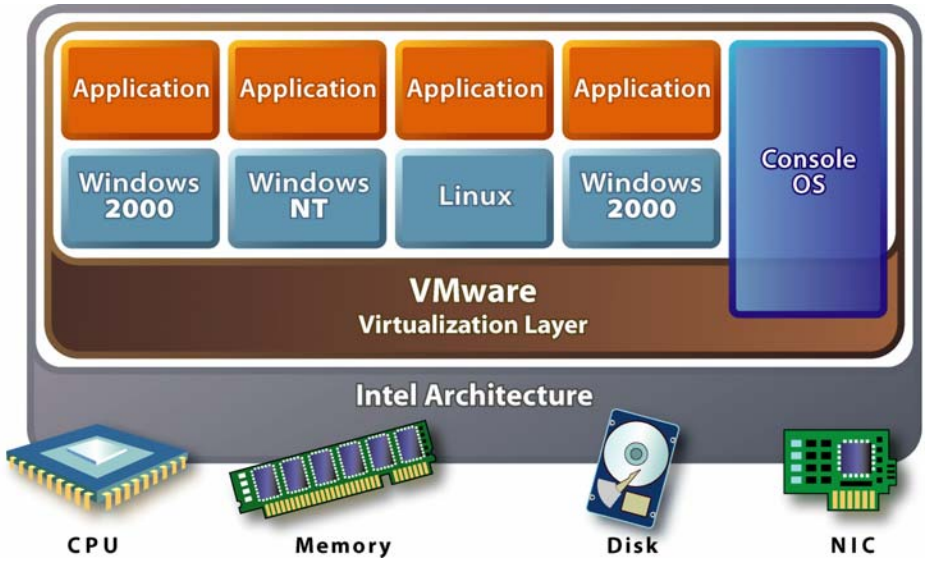

Encapsulation of Virtual Machines

- **Entire state of virtual machine is encapsulated**
  - Memory, disk images, I/O device state
- **Virtual machine state can be saved to a file**
  - Enables “Suspend / Resume” support
- **VM state can be transferred**
  - Allows a running virtual machine to be moved to another VMware system

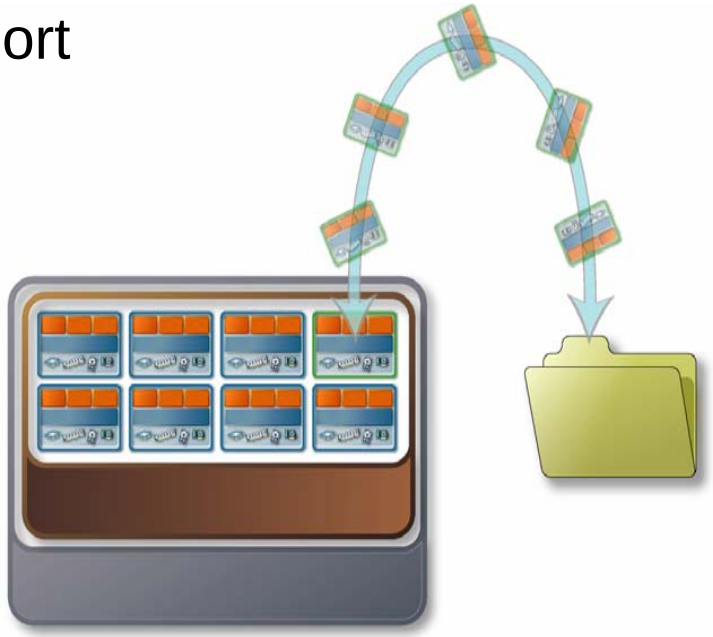

Key Features of VMware ESX Server 2

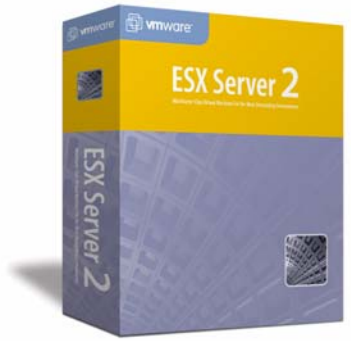

- **Scalability**
  - Architecture supports up to 80 virtual machines per Intel server
  - VMware Virtual SMP supports a maximum of 2 virtual CPUs per virtual machine
  - Optimized for NUMA architecture
- **Availability – Zero Reboot**
  - NIC Teaming
  - Integrated SAN support
  - Advanced file system
- **New Platforms**
  - Supports blade servers from IBM and HP
  - Broad range of operating system support
- **Better Usability**
  - New management user interface
  - Better resource controls
  - Improved Windows support
- **Zero-downtime Management**
  - VMware Control Center
  - VMotion Technology

Zero-Downtime Management

- **Support for VMware Control Center**
  - Manages Intel hardware as a single pool of computing resources
  - Streamlines server provisioning and management
  - Monitors system availability and performance
- **Support for VMotion technology**
  - A key component of VMware Control Center
  - Allows movement of live, stateful applications across distributed systems without server interruption
  - Perform zero-downtime, rolling hardware upgrades
  - Migrate virtual machines to new hosts in response to hardware failure

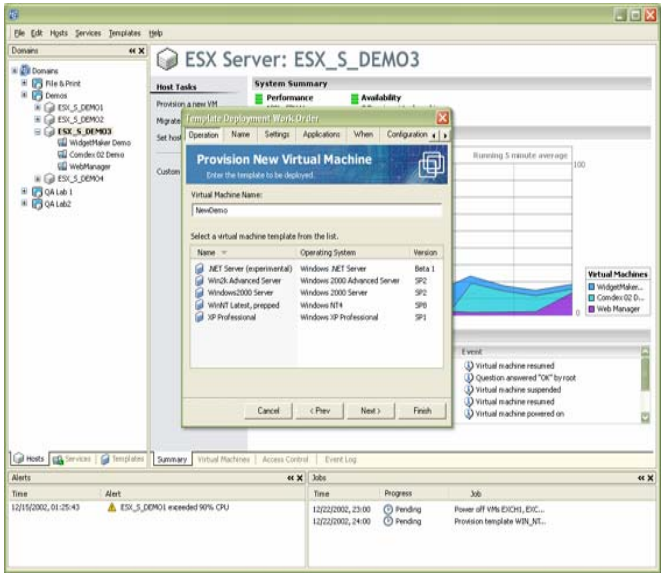

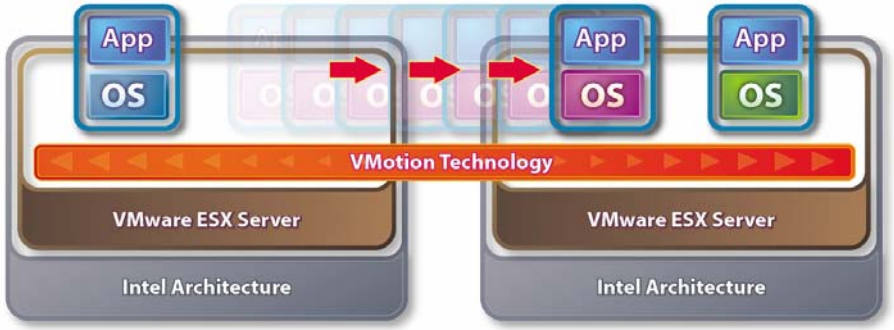

VMware Virtualization Management Architecture

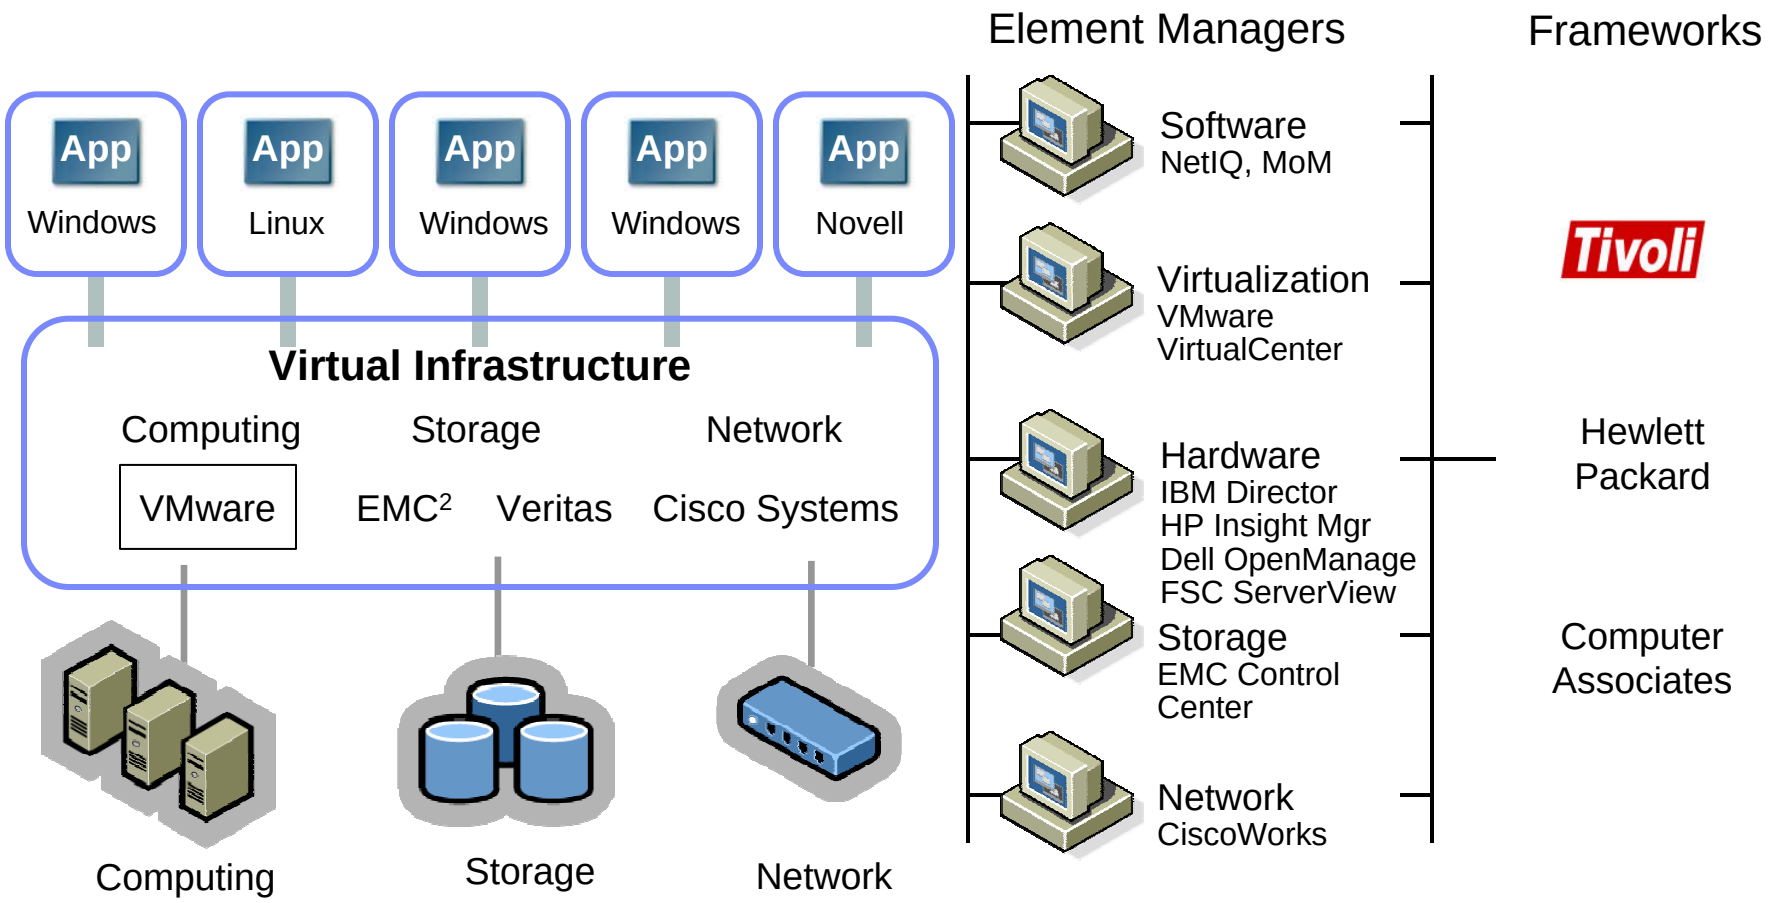

Physical to Virtual Server Consolidation with VMware

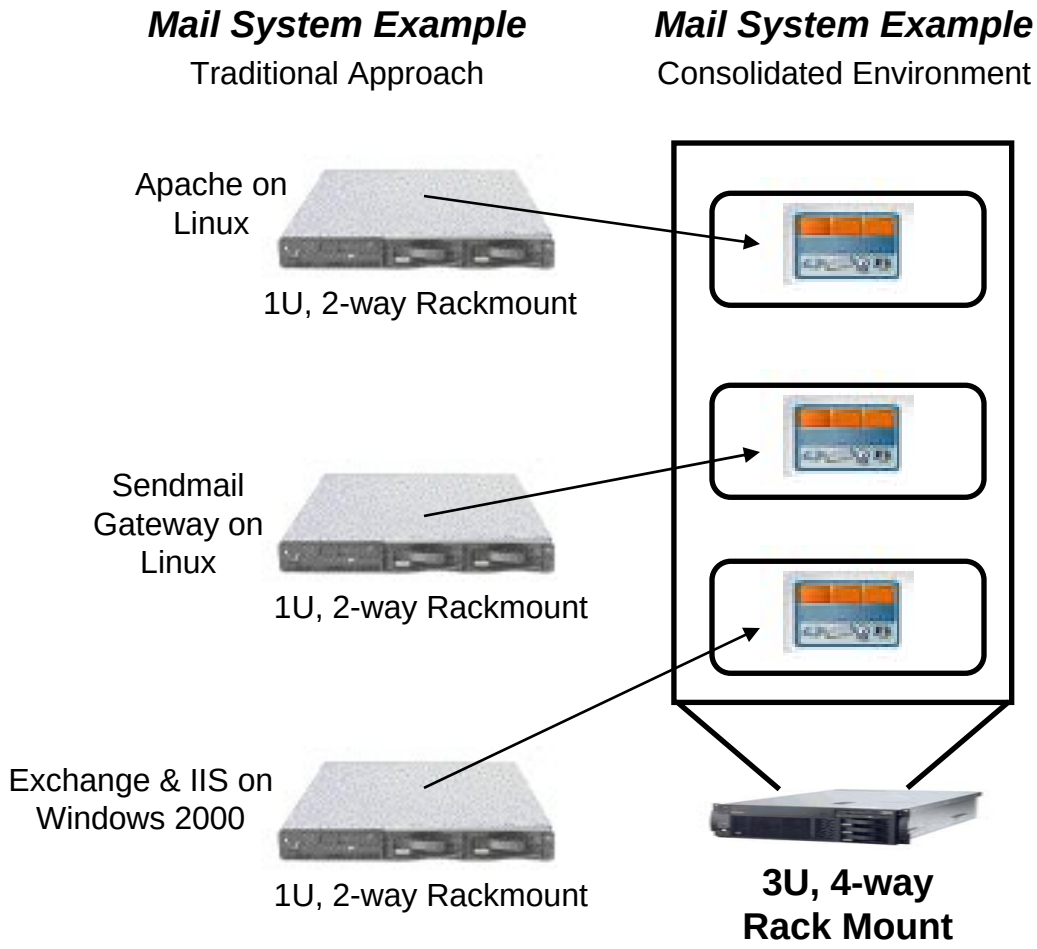

Customer Benefits

- No need to standardize on one OS
- Removes concerns about application interaction or dependencies
- VMware resource governing guarantees that critical apps will get the resources they need
- Reduced system setup and management costs

eServer zSeries Virtualization Architecture

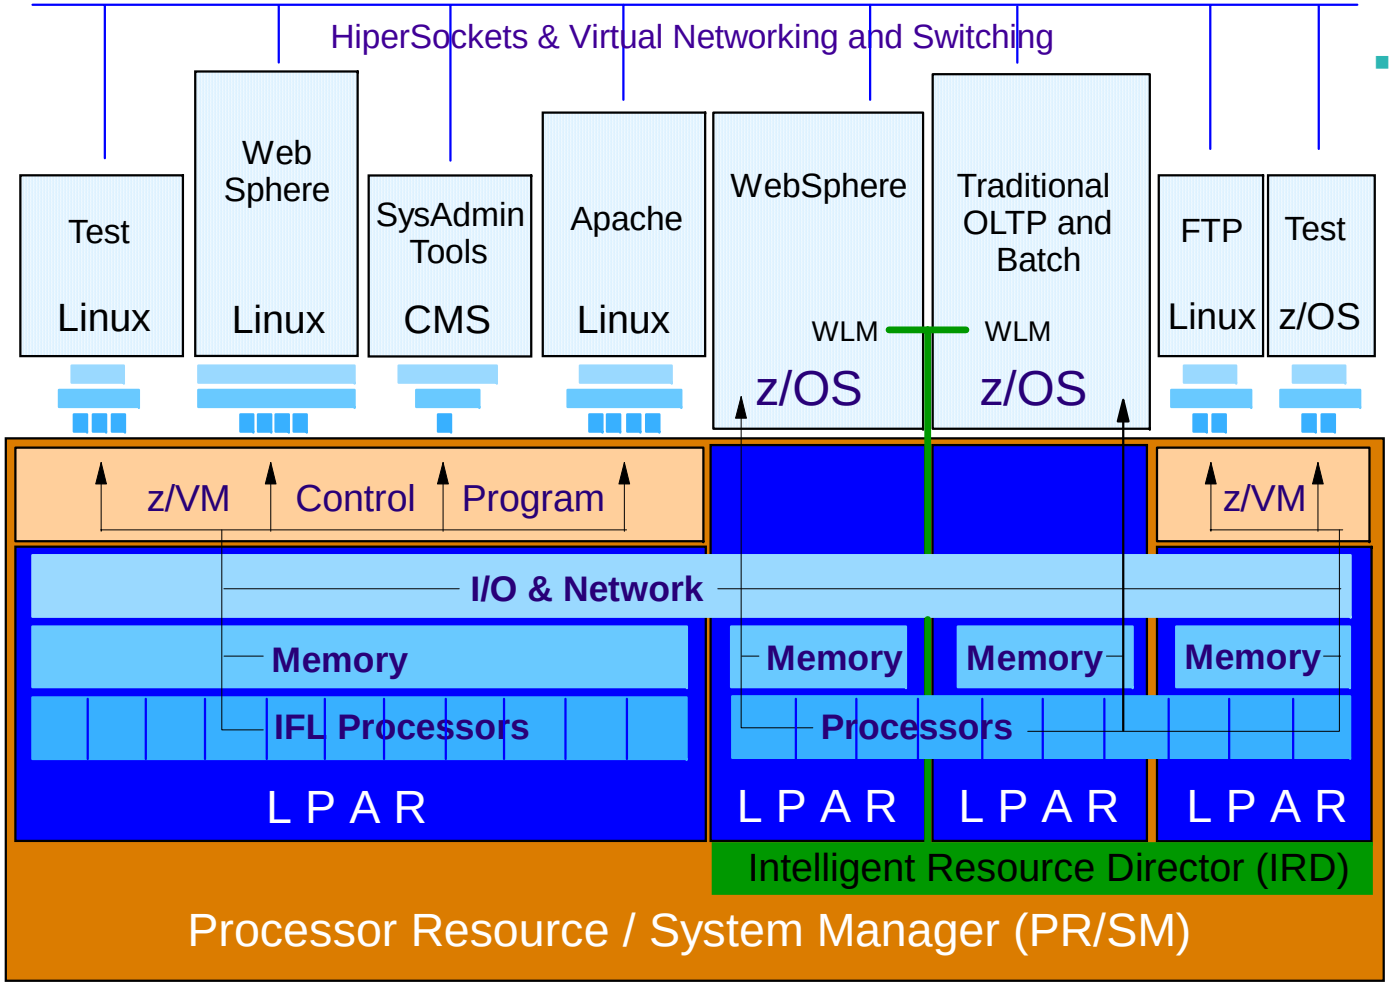

- **Multi-dimensional virtualization technology**
  - zSeries provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM™ enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS LPARs*

* Excluding non-shared resources like Integrated Facility for Linux processors

zSeries LPAR and z/VM: World-class Server Virtualization

- **zSeries Logical Partitioning (LPAR), introduced in 1988, has provided years of business-critical, high-performance server partitioning for the world's largest corporations**
- **z/VM, commercially available since 1972, has supported mixed workloads that require minimal hypervisor overhead, massive scalability, and exceptional levels of availability**
- **Both LPAR and z/VM employ hardware and firmware innovations developed over the years that make virtualization part of the basic componentry of the zSeries platform**

zSeries Interpretive Execution

Advanced Technology for Virtual Server Hosting

- **Start Interpretive Execution (SIE) instruction**
  - Operand is a state descriptor for an LPAR or virtual machine
  - Accommodates fixed-storage and pageable guests
  - Interception controls allow hypervisor intervention
  - Reduces context switch time
- **zSeries implements two levels of SIE**
  - No performance penalty for running z/VM in an LPAR
  - No shadow page tables required for DAT-on guests
  - Considerable architectural and hardware investment required
    - Potential instruction behavioral differences at each level
    - Multiple control register sets
- **Comparisons:**
  - POWER5 Hypervisor
    - Virtual processors run out their time slice regardless of events
    - Larger working sets if threads are independently dispatched; loss of efficiency if paired (one thread can become idle)
  - VMware
    - Intel does not provide assists for virtualization

Additional zSeries Virtualization Facilities

- **Zone Relocation**
  - SIE capability that provides multiple zero-origin storage regions (logical partitions) on one system
  - Enables I/O subsystem to access partition memory directly, without requiring hypervisor intervention
- **Translation Lookaside Buffers (TLBs)**
  - Large allocation of microprocessor space for TLBs directly benefits virtual server scalability
  - z990 provides a TLB arrangement which advantageously uses two buffers
    - Second-level TLB feeds address translation information to the first-level TLB when the desired virtual address is not contained in the first-level TLB
- **Multiple Image Facility (MIF)**
  - Enables channel sharing among multiple LPARs
  - I/O devices on shared channel paths can be accessed simultaneously by sharing LPARs (or restricted to a subset of sharing LPARs)
- **Logical Channel Subsystems (LCSS) support**
  - Allows a z990 to be configured with up to 1024 channels (512 channels for z890)
  - 256 channels can be configured for each LPAR, with selected channel sharing among LPARs possible

Additional zSeries Virtualization Facilities

- **I/O Priority Queueing**
  - Allows high-priority workloads to receive preferential access to I/O subsystem
  - Supported by Intelligent Resource Director and virtualized by z/VM
- **HiperSockets**
  - High-speed, security-rich TCP/IP connectivity among LPARs
  - Memory speed communications
- **Adapter Interruption Pass-Through support**
  - OSA (Ethernet) and FCP (SCSI) virtual machine I/O can be performed while z/VM guest image is running in SIE mode
  - “Thin” interrupt passed to z/VM Control Program when I/O operation belongs to an idle guest system
- **Layer 2 (MAC) and Layer 3 (IP) network switching**
  - OSA and z/VM support enables virtual IP and MAC network switching without requiring a hosting partition

z/VM CPU Resource Controls

Highly Granular Sharing of System Resources

- Allocate system resources among virtual machines using SET SHARE command
  - This is a highly flexible and self-managed function of the z/VM Control Program
  - Reserve CPU capacity for peak usage
    - Use it when needed
    - Relinquish the processor cycles for other servers when not needed
  - "Absolute guests" receive top priority
  - The Virtual Machine Resource Manager can be used to monitor and adjust remaining capacity allocated to "Relative guests"

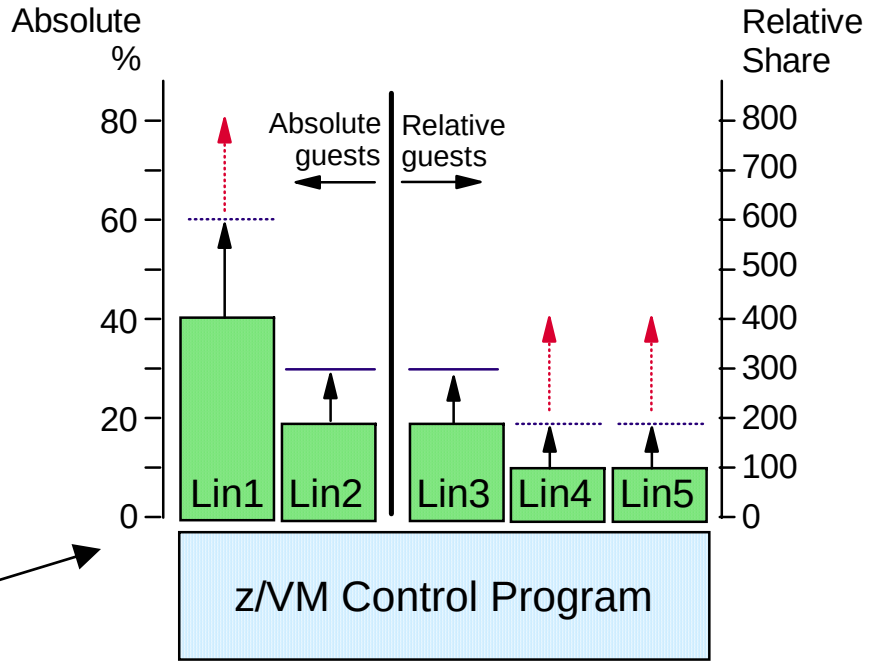

z/VM Directory Entries

```
SHARE Lin1 ABSOLUTE 40% ABSOLUTE 60% LIMITSOFT
SHARE Lin2 ABSOLUTE 20% ABSOLUTE 30% LIMITHARD
SHARE Lin3 RELATIVE 200 RELATIVE 300 LIMITHARD
SHARE Lin4 RELATIVE 100 RELATIVE 200 LIMITSOFT
SHARE Lin5 RELATIVE 100 RELATIVE 200 LIMITSOFT
```

Notes:

..... = limit can be exceeded if unused capacity is available (limitsoft)

——— = limit will not be exceeded (limithard)

z/VM Disk Virtualization Technology

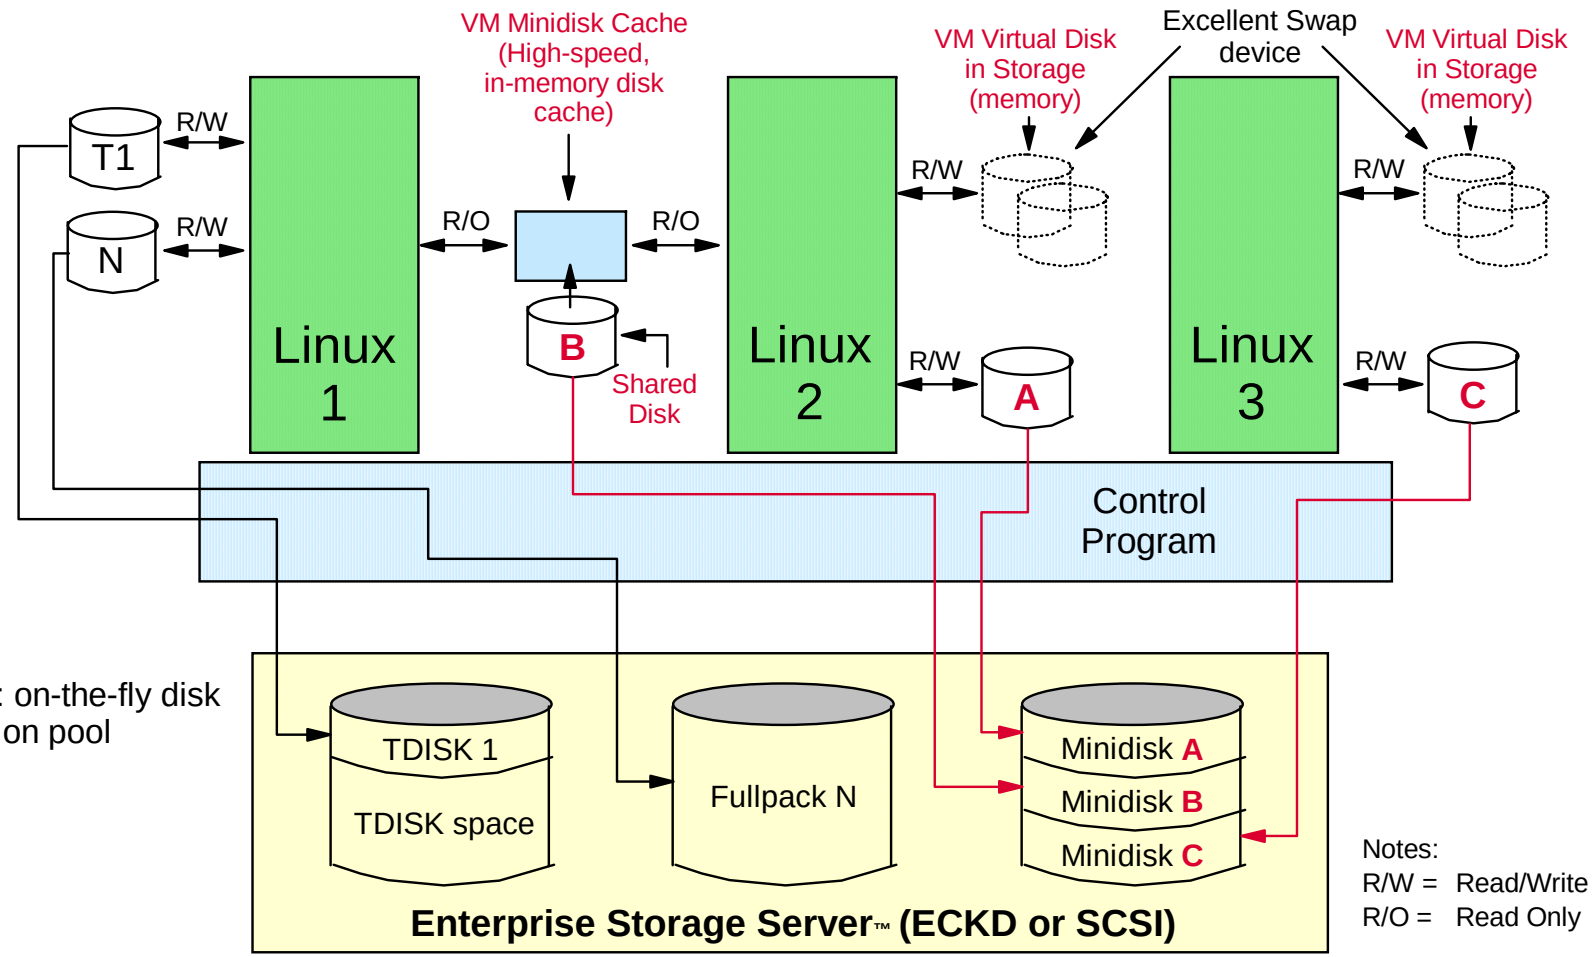

TDISK: on-the-fly disk allocation pool

Notes:
R/W = Read/Write
R/O = Read Only

Linux and z/VM Technology Exploitation

Exploitation of z/VM Support for Discontiguous Shared Segments (DCSS)

- DCSS support is a z/VM exclusive
  - Share a single, real memory area among multiple virtual machines
  - High-performance data access
  - Can reduce real memory utilization
- Linux exploitation support available today
  - Execute-in-place (xip2) file system
  - DCSS memory locations can reside outside the defined virtual machine configuration
  - Access to file system is at memory speeds; executables are invoked directly out of the file system (no data movement required)
  - Avoids duplication of virtual memory and data stored on disks
  - Enables throughput benefits for Linux guest images and enhances overall system performance and scalability

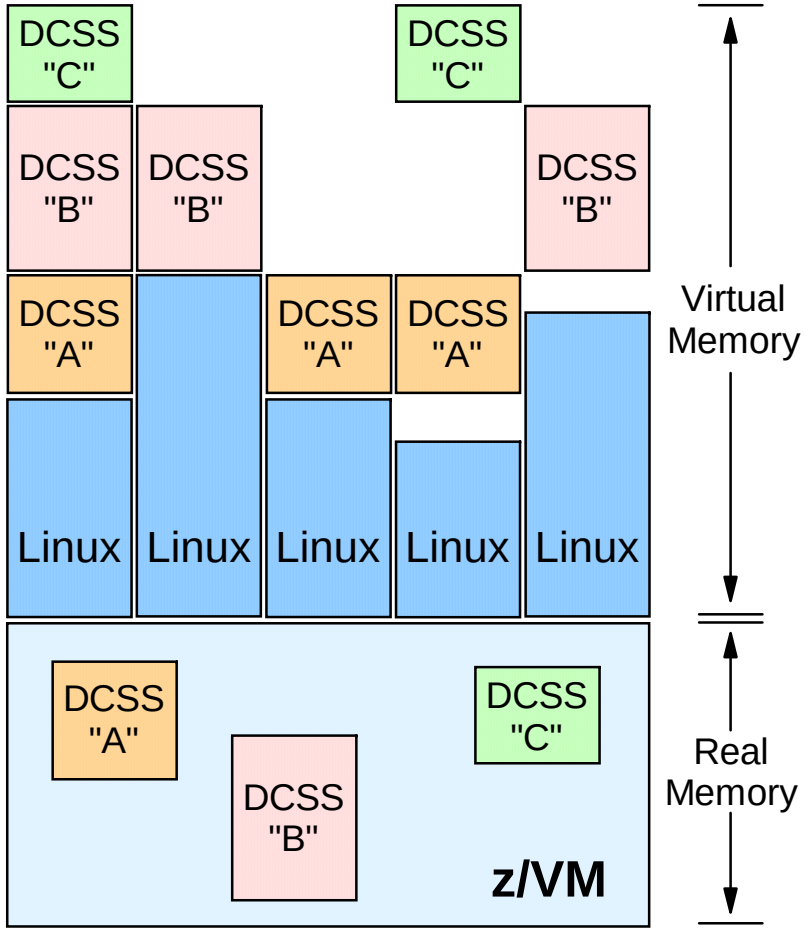

Monitoring Linux on zSeries

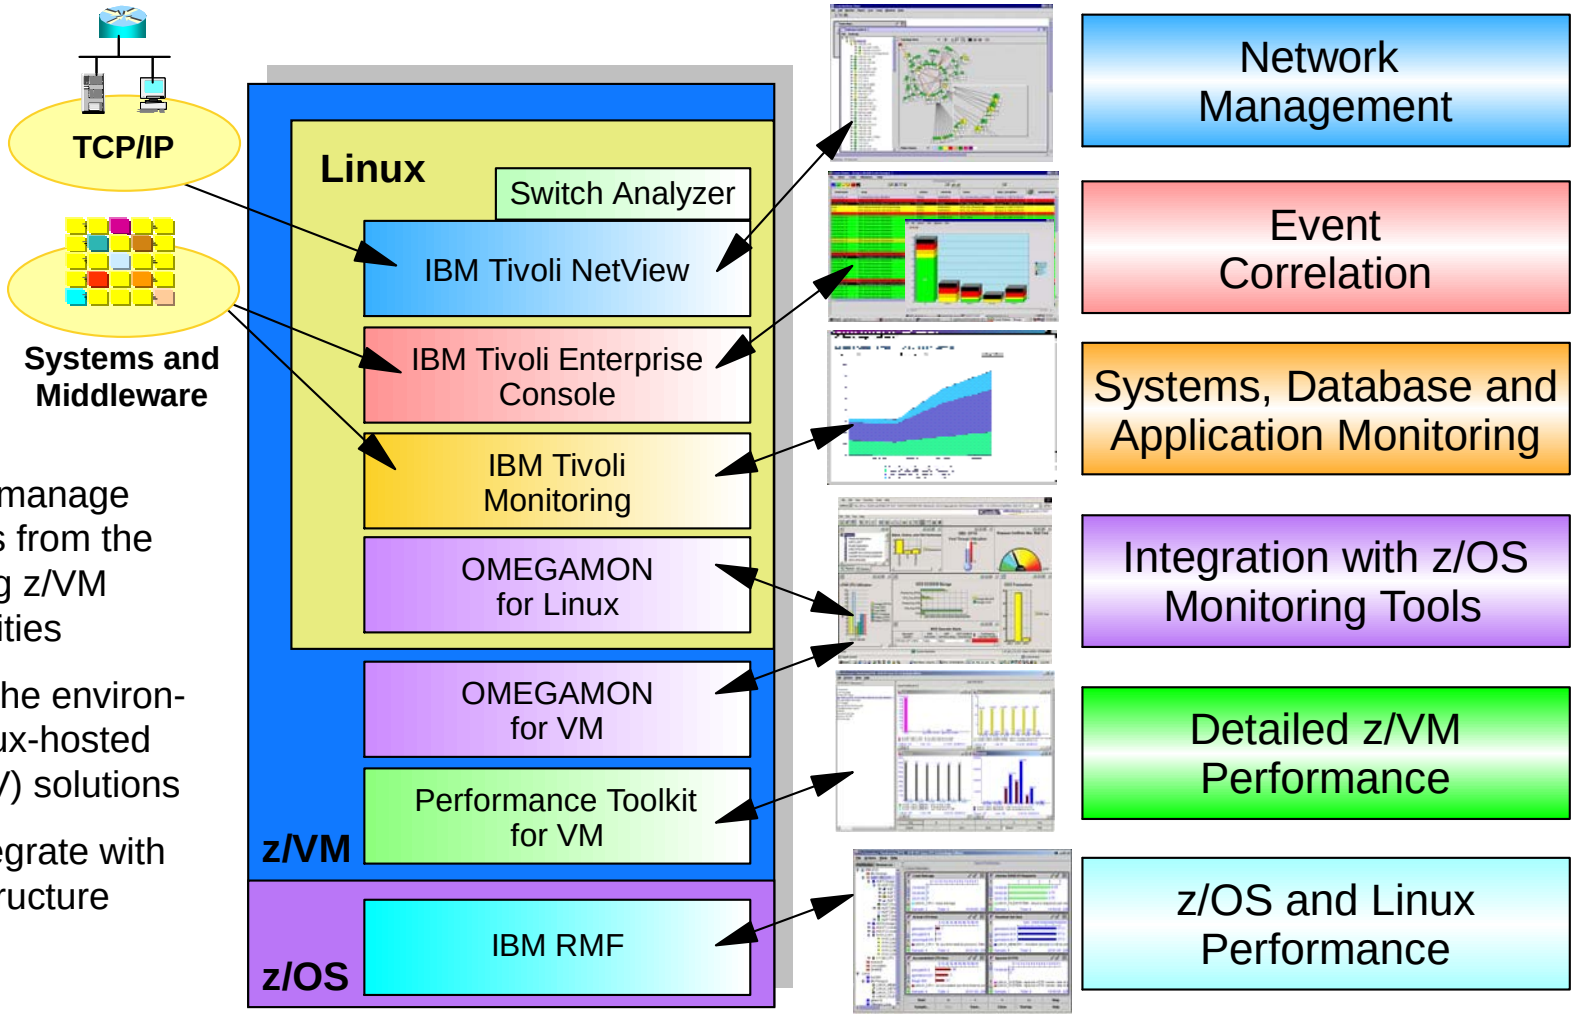

Measure and manage guest systems from the "outside" using z/VM tools and facilities

Complement the environment with Linux-hosted Tivoli (and ISV) solutions

Optionally integrate with a z/OS infrastructure

GDPS/PPRC Multiplatform Resiliency for zSeries

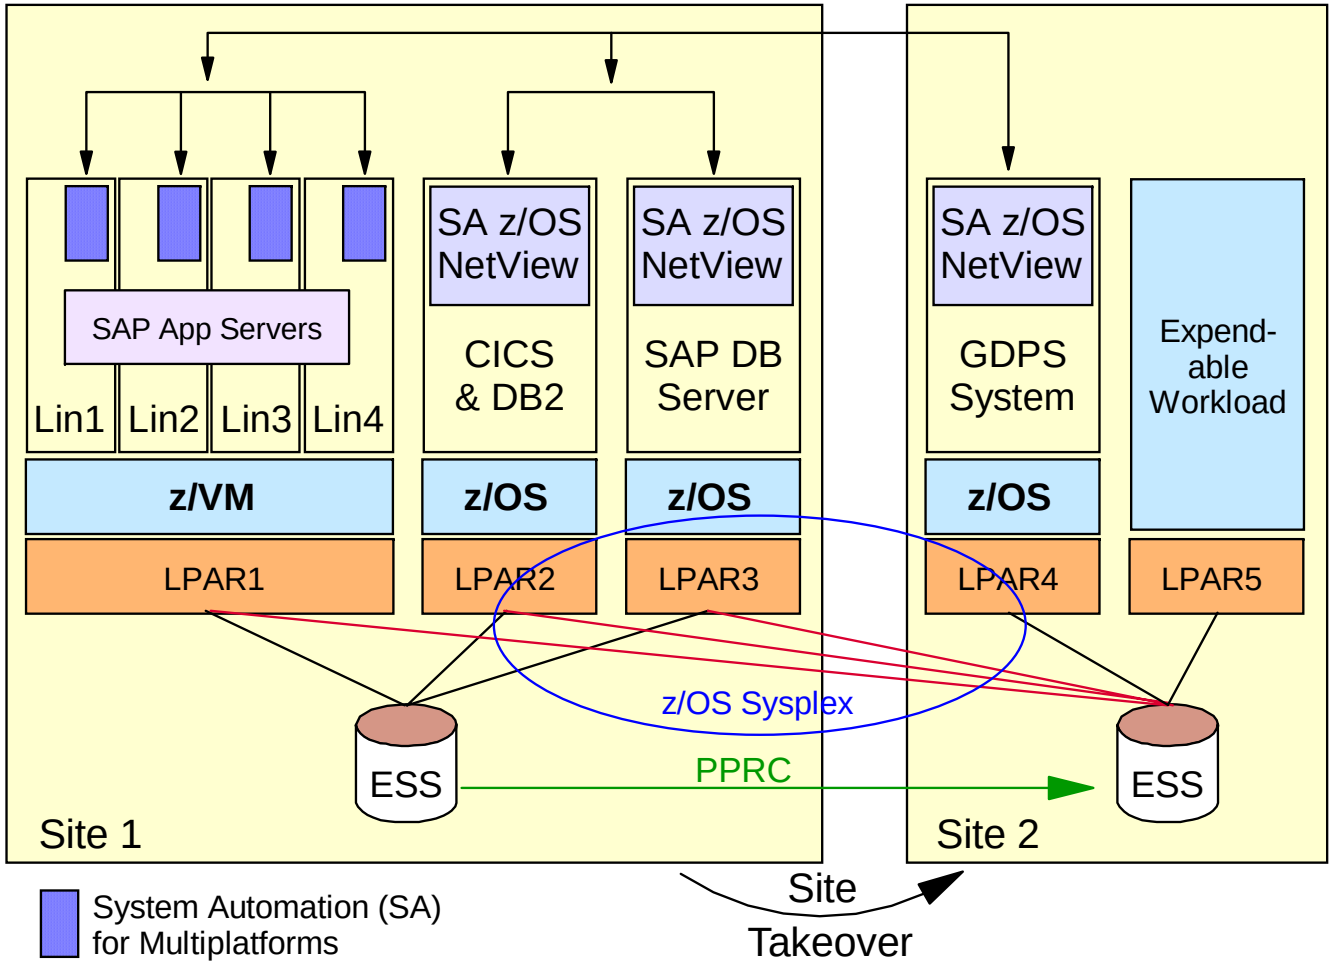

- Designed for customers with distributed applications
- SAP application server running on Linux for zSeries
- SAP DB server running on z/OS
- Coordinated near-continuous availability and DR solution for z/OS, Linux guests, and z/VM
- Uses z/VM HyperSwap function to switch to secondary disks
- Sysplex support allows for site recovery

IBM Systems and Technology Group University 2005

zSeries Partitioning Differentiation

A “Genetic” View of Virtualization Technologies

- **Today’s server virtualization capabilities are influenced by the heritage of the platform**
  - zSeries
    - Mainframes have hosted mixed workloads and have optimized resource sharing ever since their introduction in 1964
    - Mainframe investments have been driven by the business community, where high qualities of service are essential
  - pSeries
    - Unix systems achieved early success in the scientific and engineering community, where high-performance number crunching was a priority
    - Rebooting a failed server was not considered a serious problem
  - xSeries
    - Intel servers gained popularity as an inexpensive platform offering “application freedom” from the IT organization
    - One application per Intel server was the norm; “Ctl-Alt-Del” was the standard method of solving problems

Implications of Virtualization Genetics

- **Virtualization technology and "awareness" is built into mainframe architectures, hardware, firmware, hypervisors, and operating systems**
  - The technology is very responsive and stable, "hardened" by 40 years of customer usage and feedback
- **Virtualization technology is largely an "add-on" for pSeries and Intel system designs**
  - Longer feedback paths reduce responsiveness and stability
  - The technology does not enjoy years of high-stress deployment in a wide-range of customer environments
- **Nonetheless...**
  - One server platform does not satisfy all customer requirements
  - A multi-platform deployment of virtualization technology offers maximum responsiveness and efficient resource utilization for on demand computing
  - Exploiting zSeries in this environment is good for our customers and good for IBM
  - Tivoli solutions and the Virtualization Engine can help customers manage an environment that has a mix of server technologies

Chip Design Affects Virtualization Capabilities

Replicated Server Chip Design

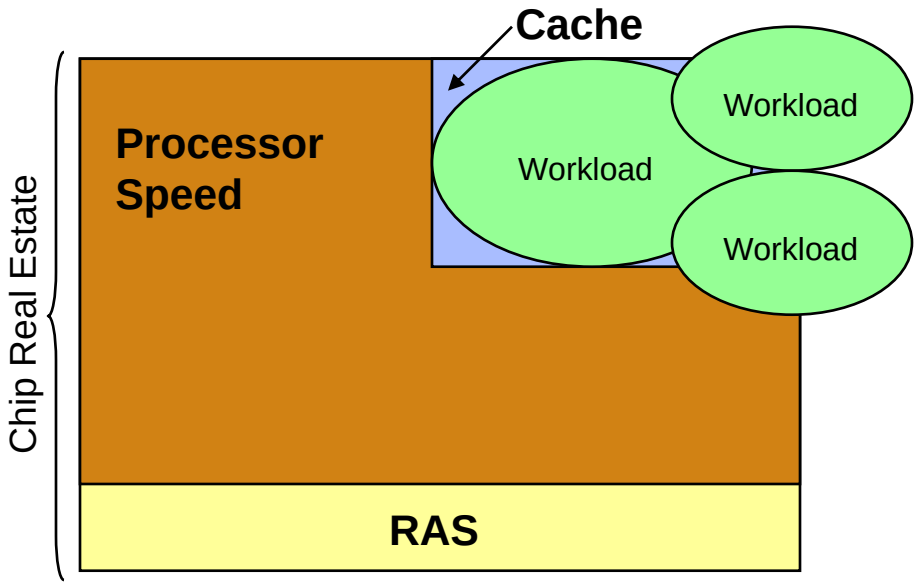

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- "Fast" processor speed is not fully realized due to cache misses

Consolidated Server Chip Design

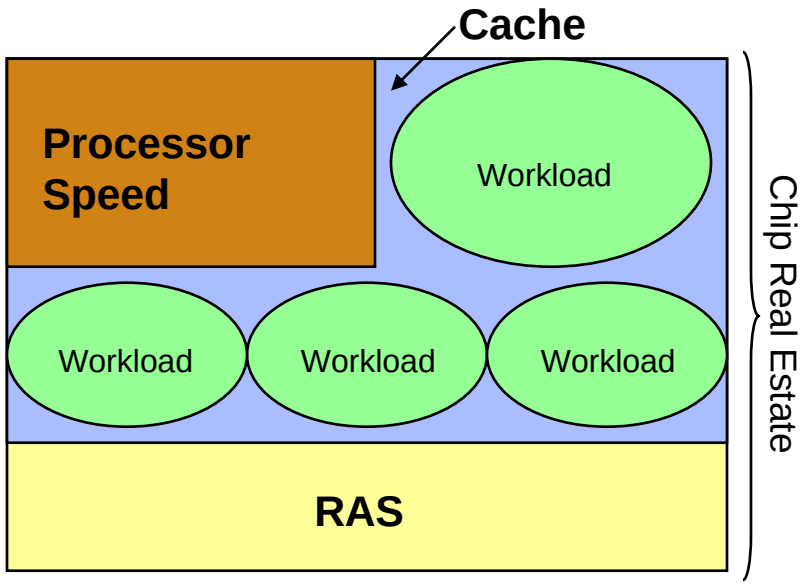

- zSeries cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

Scalability Considerations: Context Switching

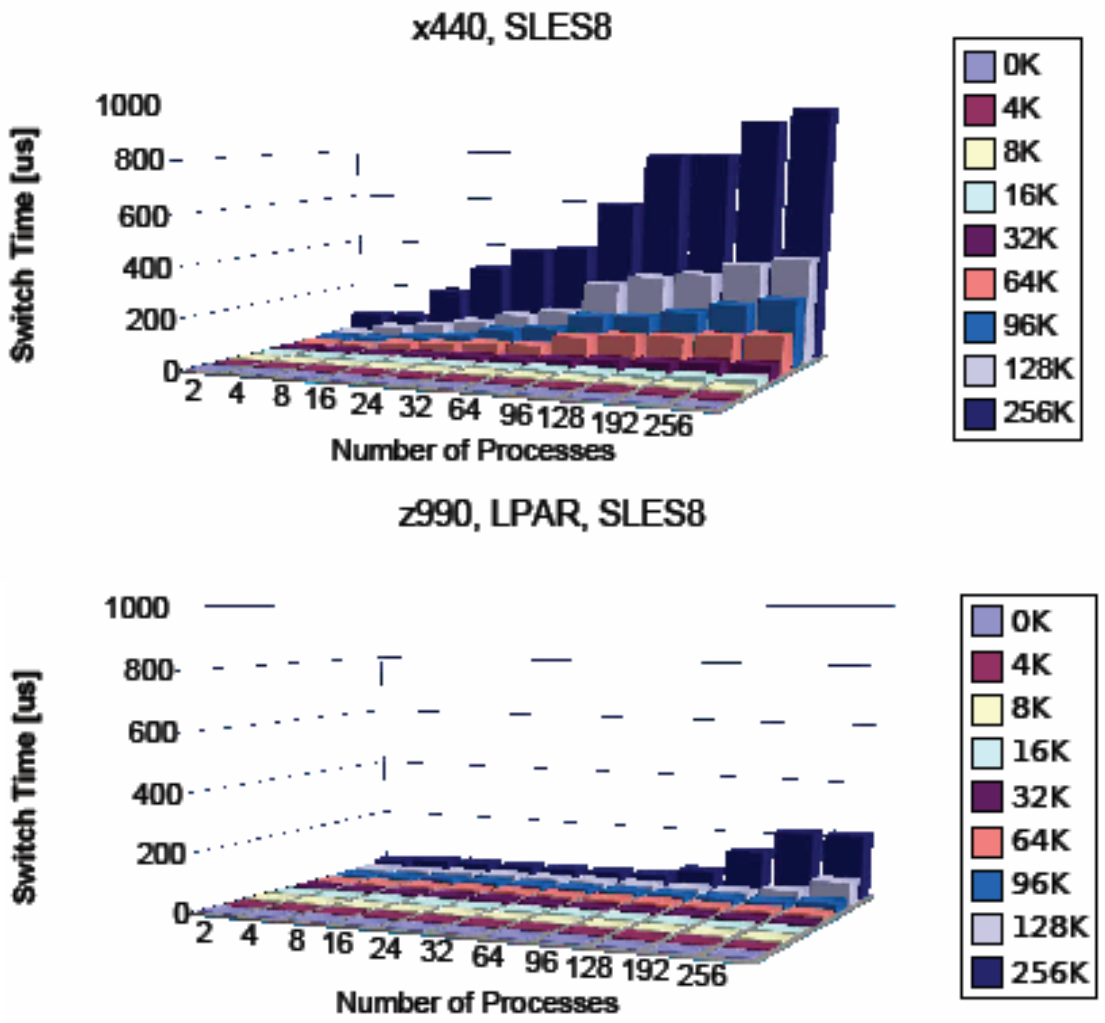

- Time required to perform context switching is an indication of memory time
- Virtualization, by definition, involves context switching

pSeries measurements not available yet

zSeries Partitioning Leadership Support

- **Highest levels of RAS built into the hardware**
- **Non-disruptive On/Off Capacity on Demand**
- **Linux and z/OS application integration**
- **Highly granular allocation of hardware assets**
  - Add “small” server images to existing config with minimal impact to other server images
- **Large-scale server hosting**
  - Can run more than a thousand server images
- **Resource consumption recording / reporting**
  - Capture data at hypervisor level (CP Monitor)
  - Useful for charge-back, capacity planning, problem determination, and fix verification
- **Hot stand-by without the hardware expense**
  - Idle backup images ready to run (or be booted) if primary servers fail
- **Autonomic, non-disruptive disk failover to secondary storage subsystem**
- **Architecture simulation**
  - Satisfy server configuration requirements without suffering expense of real hardware
- **In-memory application sharing**
  - Share program executables among multiple server images
- **Server-memory-cached disk I/O**
  - High-speed read access to files on disk
- **Virtual Disks in Storage**
  - High-speed read and write access to files in memory (excellent swap devices for Linux)
- **Built-in console message routing**
  - Route messages from all virtual servers to a single virtual machine (system automation)
- **Virtual Machine Resource Manager**
- **“Hands free” auto-logon of server images**
  - Using z/VM “Autolog” support
- **Initiate operating system shutdown from “outside” the server image**
  - Without requiring agent running on guest operating system
- **Up to 256 Linux servers can share a single Cryptographic card using z/VM**

zSeries Versus pSeries Considerations

- **More granular allocation of CPU resources with zSeries**
- **z/VM shared memory model allows sharing and dynamic allocation of real memory to virtual server images**
- **Faster virtual server creation / provisioning with z/VM**
- **p5 has LAN-like latency for shared I/O and networking**
- **Compute-intensive workload is a better fit for p5**
- **Customer acceptance**
  - Linux on zSeries has a 5-year history of success
  - Linux on Power is relatively new

zSeries Versus VMware Considerations

- **zSeries enjoys significant scalability advantages over VMware**
  - CPU utilization and granularity
  - Sophisticated paging subsystem
  - Firmware-assisted I/O and networking
- **z/VM allows up to 64 virtual CPUs per virtual machine**
  - VMware restriction of 2 CPUs per virtual machine may constrain some Linux workloads
- **VMware only supports 32-bit x86 images, even on 64-bit hardware**
  - z/VM can support a mix of 32-bit and 64-bit virtual servers
- **Hosting multiple z/VM images on a single server (via LPAR) can enhance failover options and workload distribution**
- **Application integration of Linux and Windows favors VMware**
- **Application integration of Linux and z/OS favors z/VM**
- **EMC ownership of VMware may influence customer selection of storage and enterprise systems management solutions (SOSWOS issues)**
- **VMware support for non-IBM Intel systems gives customers negotiating leverage**

When Do You Need More than "Good Enough"?

Making the Case for zSeries Virtualization

- **When workload growth and decline is difficult to predict (be it production, development, or test/assurance systems)**
- **When customer demand does not match your IT resources and business results suffer**
- **When your IT staff wants to optimize their productivity for deploying and managing virtual servers**
- **When innovative is stifled because your staff cannot experiment or develop new solutions using existing resources**
- **When speed to market affects your business results**
- **When your server applications need fast and flexible access to z/OS data and applications**
- **When business resiliency is a high priority**
- **When you want more control over your environmental expenses (e.g., floor space, cooling)**

Quotes

"The advantages of virtualization in both management and efficiency are major contributors to the renaissance of the IBM zSeries, which has a richer history and greater intellectual property associated with virtualization than does any other platform."

- Gordon Haff, Illuminata, "IBM Fits Linux to Power", 9/13/2004

"When considering ESX Server, it's vital not to lose sight of one inescapable reality: PC servers are not designed for virtualization or hardware partitioning."

- Tom Yager, "VMware Delivers a Datacenter in a Box", Computerworld.com, 11/09/2004

Additional Resources

- **The Campus for more education:**
  - BPs: http://www.ibm.com/partnerworld/sales/systems/education
    - Includes link to IBM PartnerWorld University (Web lectures for key topics)
    - http://www.ibmweblectureservices.ihost.com/pwu
  - IBMers: w3.ibm.com/sales/systems/education
    - Includes links to the Online Universities for Cross-Brand and each Brand (Web lectures for key topics)
  - Customers: www-1.ibm.com/servers/eserver/education

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: 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

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
LINUX is a registered trademark of Linux Torvalds
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.

IBM Systems and Technology Group University 2005

Please Complete the Brief Survey on This Session

Your Feedback Helps Us Provide You Right Education