Selecting the Best Fit Platform for New and Mixed Workloads

The future runs on System z

Jim Porell
IBM Distinguished Engineer
IBM System z Business Development

Topics

- Business drivers for server consolidation
- System z: wired for virtualization
- System z workload management: what every business needs
- Save money, grow your business

Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration

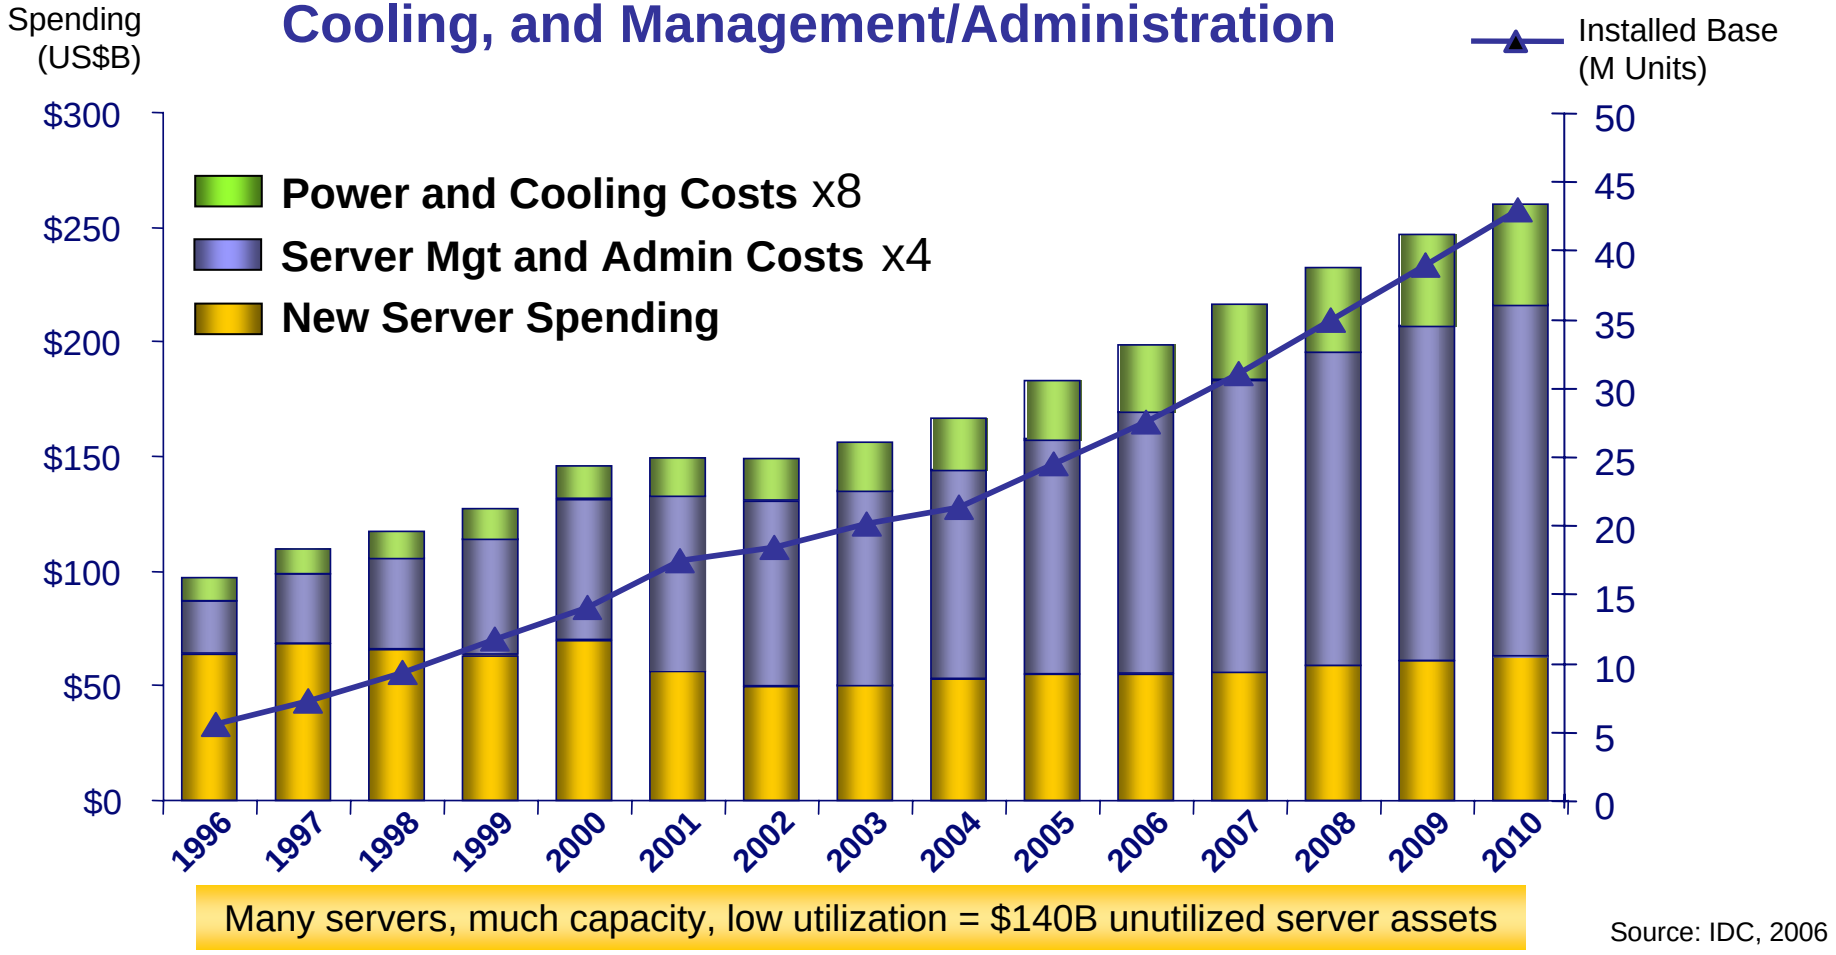

Many servers, much capacity, low utilization = $140B unutilized server assets

Source: IDC, 2006

As server volumes increase, so does complexity, making basic business requirements like availability, security, and disaster recovery more difficult to achieve.

Server Architecture Genetics

Consider the Heritage of Today's Server Platforms

- **x86 systems**
  - Key value proposition: *end-user autonomy*
  - "Ctl-Alt-Del" not a problem for a single-user system
- **UNIX systems**
  - Key value proposition: *processor speed*
  - Sweet spot: engineering/scientific computing
- **Mainframe systems**
  - Key value proposition: *mixed workloads*
  - Highest degrees of efficiency, availability, workload mgmt, security

Virtualization Essentials

Virtualization technology can be significantly constrained or compromised by the underlying system architecture.

Virtualization and Security
Should IT Managers Be Concerned?

Virtualization security risks being overlooked, Gartner warns

Gartner raises warning on virtualization and security.

Companies in a rush to deploy virtualization technologies for server consolidation efforts could wind up overlooking many security issues and exposing themselves to risks, warns research firm Gartner.

“Virtualization, as with any emerging technology, will be the target of new security threats,” said Neil MacDonald, a vice president at Gartner, in a published statement.

– NetworkWorld.com, April 6, 2007

Posted September 21, 2007 at http://blogs.iss.net/archive/virtblog.html

"It is clear that with the increase in popularity, relevance and deployment of virtualization starting in 2006, vulnerability discovery energies have increasingly focused on finding ways to exploit virtualization technologies."

"...in a virtual environment all your exploitation risks are now consolidated into one physical target where exploiting one system could potentially allow access and control of multiple systems on that server (or the server itself). In total, this adds up to a more complex and risky security environment."

Known vulnerabilities across all of VMware's products*

VMware Vulns by Year	Total Vulns	High Risk Vulns	Remote Vulns	Vulns in 1st Party Code	Vulns in 3rd Party Code
Vulns in 2003	9	5	5	5	4
Vulns in 2004	4	2	0	2	2
Vulns in 2005	10	5	5	4	6
Vulns in 2006	38	13	27	10	28
Vulns in 2007	34	18	19	22	12

z/VM Security Architecture

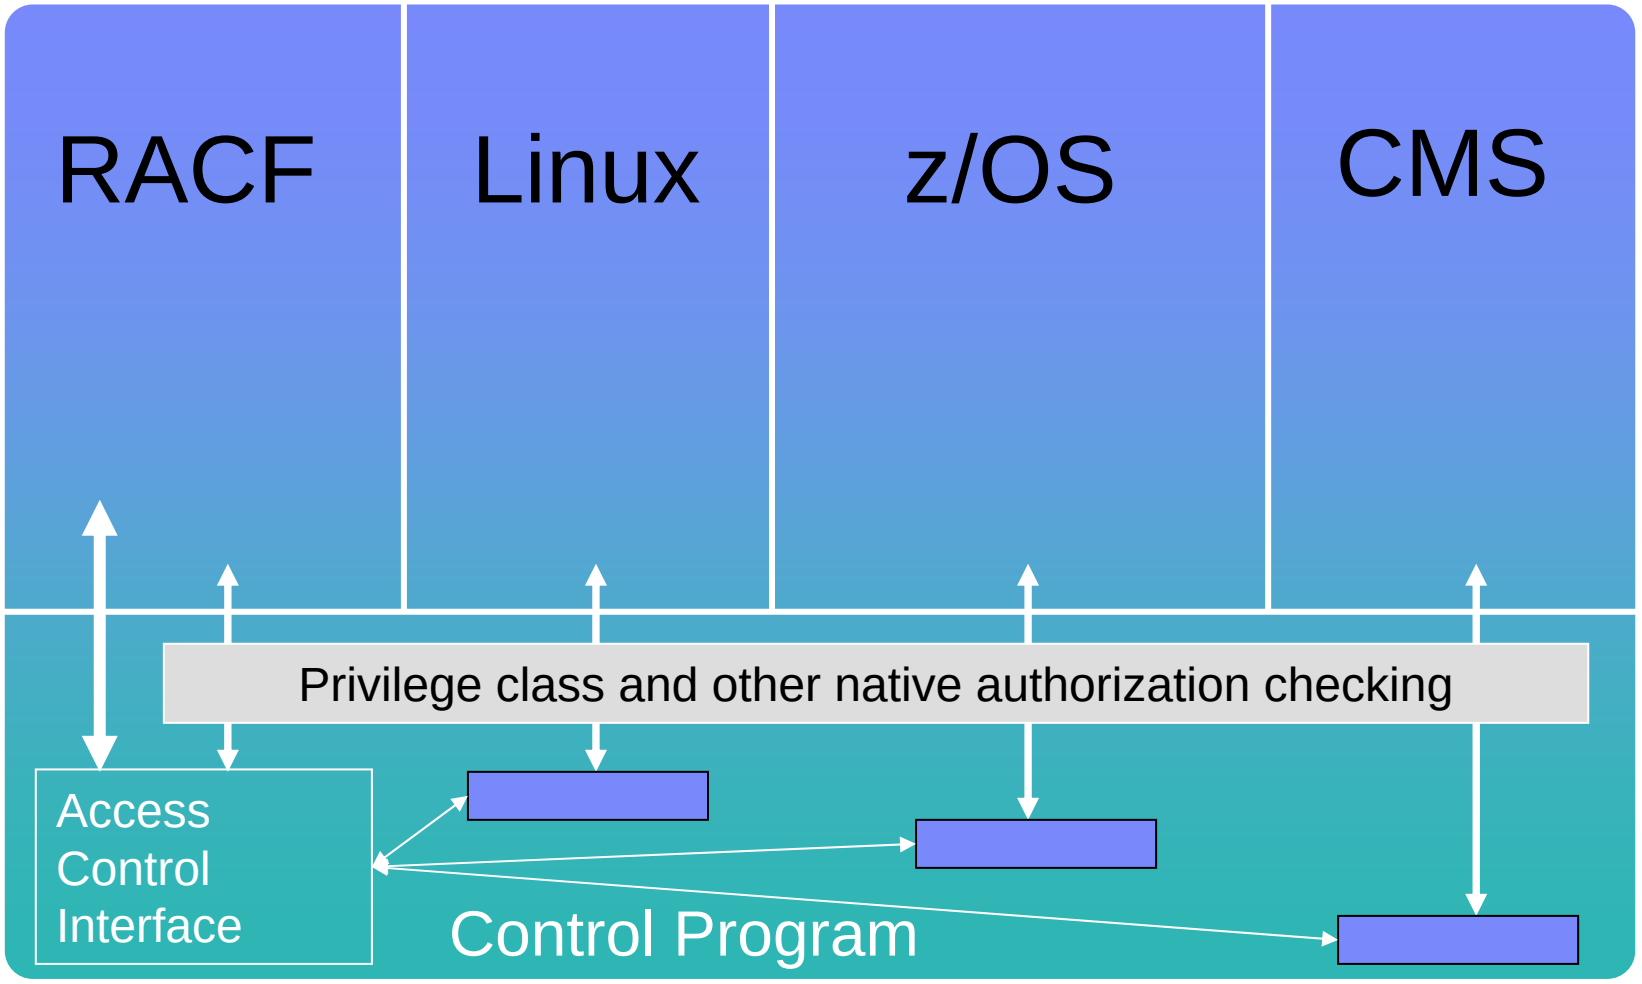

Topics

- Business drivers for server consolidation
- **System z: wired for virtualization**
- System z workload management: what every business needs
- Saving money, growing your business

IBM System z Virtualization Genetics

- **System z is *thoroughly* architected to host applications in a virtualized environment**
- **This is accomplished with a coordinated set of investments that permeate the technology stack of *hardware*, *firmware*, *hypervisors*, and *operating systems***
- **This means clients can maximize the utilization, scalability, and security of all system assets, including:**
  - *CPU*
  - *Memory*
  - *I/O*
  - *Networking*
  - *Cryptography*
- **All with exceptional levels of operational ease and cost efficiencies**

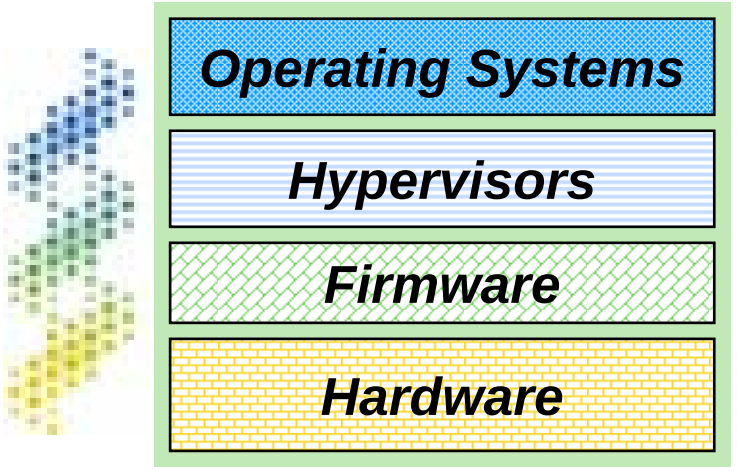

IBM System z Virtualization Leadership

Extreme Levels of CPU Sharing

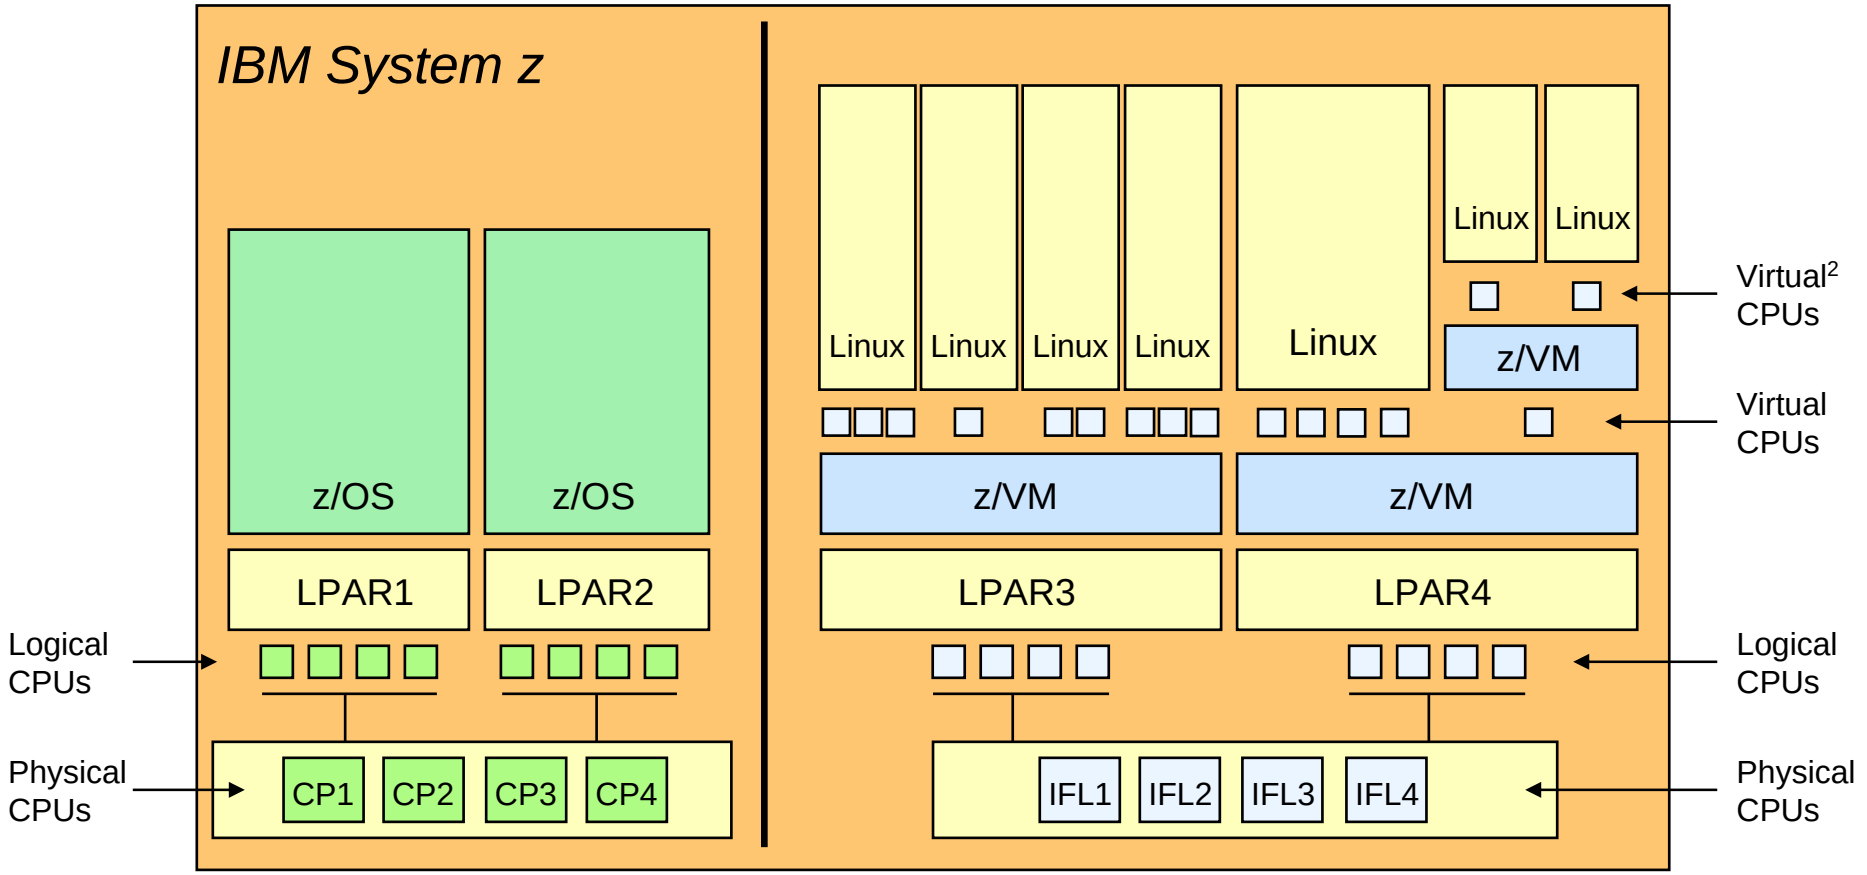

IBM System z CPU High Availability
Concurrent Processor Reassignment

- **Used to concurrently change the physical backing of one or more logical processors**
- **The state of source physical processor is captured and transplanted into the target physical processor**
- **Operation is transparent to operating systems**
- **Used for *processor sparing* and *book replacement***

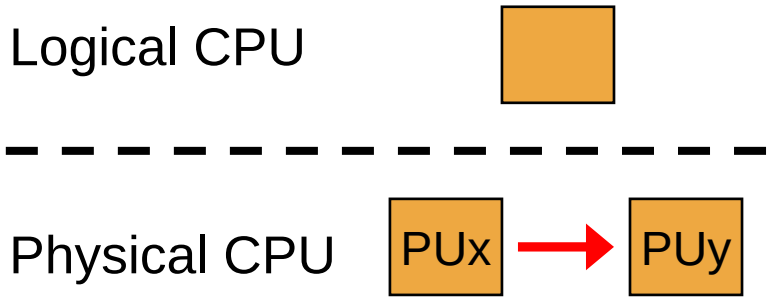

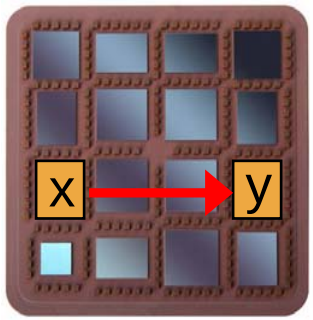

System Design Affects Virtualization Capabilities

System z packs a lot of compute power into a single box

- With TCO-friendly pricing

Up to 64-way SMP

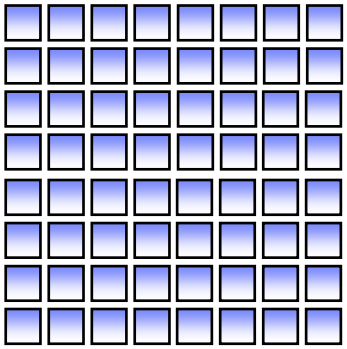

Share up to 64 processors with up to 60 LPARs

Configure these processors as CPs, IFLs, zAAPs*, zIIPs*, or ICFs*

* No software license fees

2 Standard Spare PUs

Up to 11 System Assist Processors

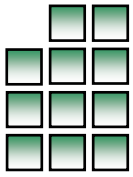

Offload system processing to dedicated CPUs (no impact to software license fees)

Up to 336 I/O Processors

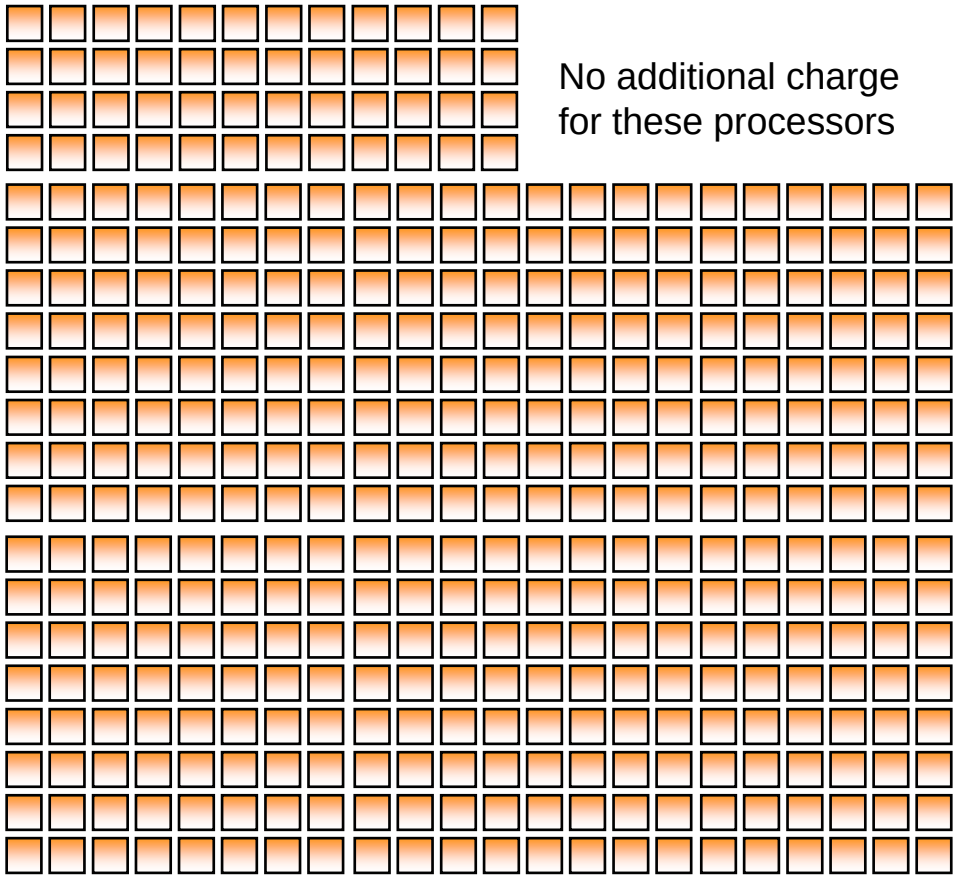

No additional charge for these processors

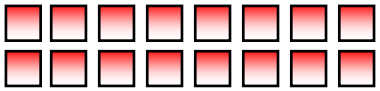

Up to 16 Crypto Express2 CPUs

High scale performance for SSL transactions

System Design Affects Virtualization Capabilities

Compare to typical UNIX system design...

Up to 128-way SMP configuration

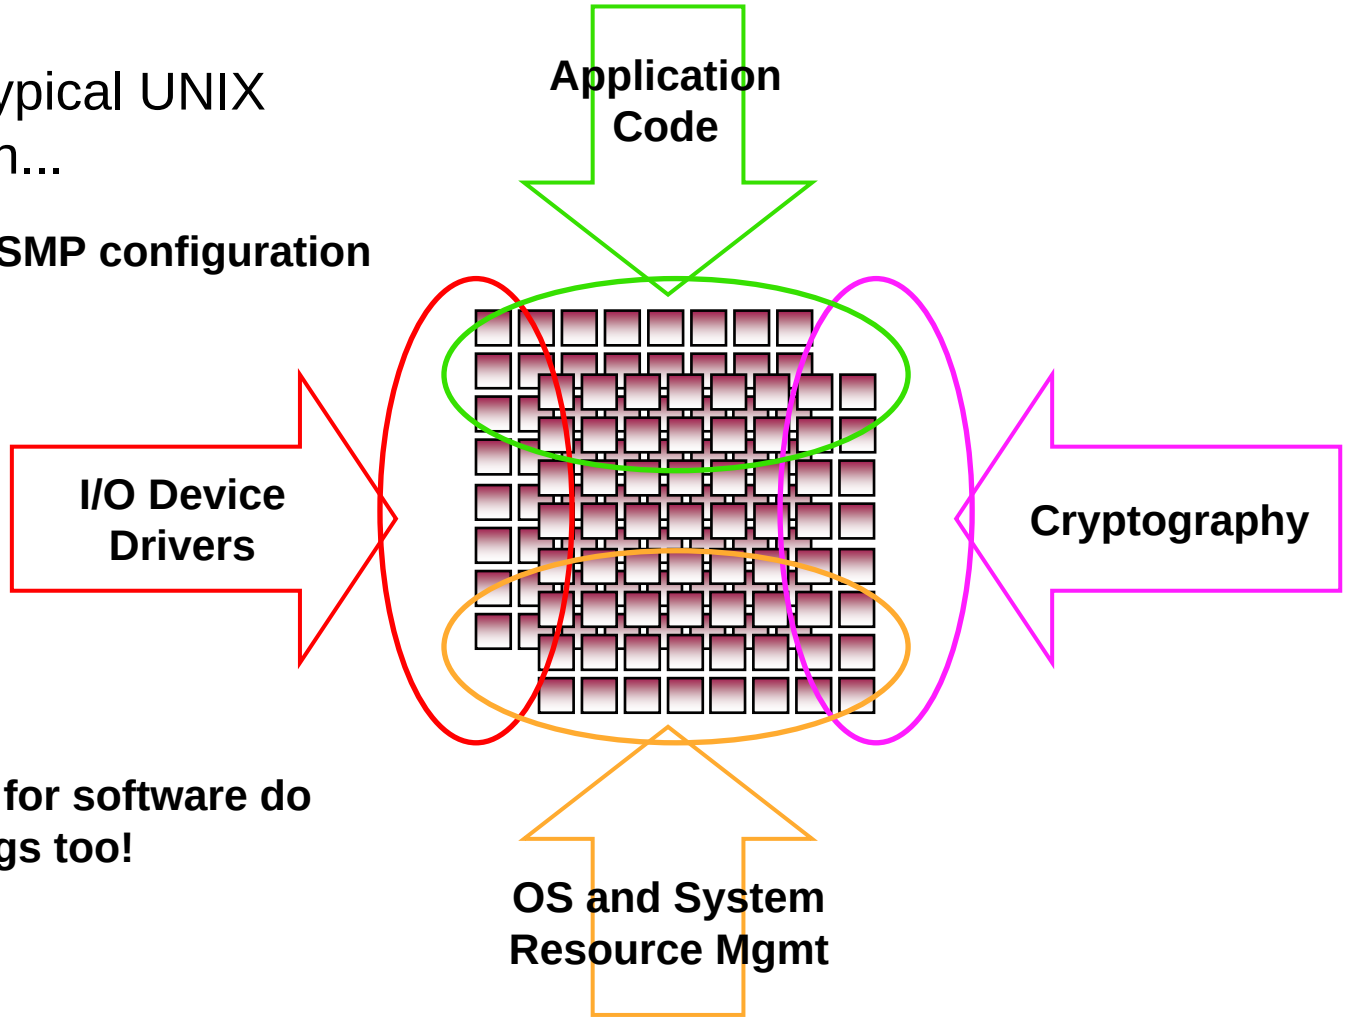

CPUs licensed for software do a lot other things too!

Chip Design Affects Virtualization Capabilities

Replicated Server Chip Design

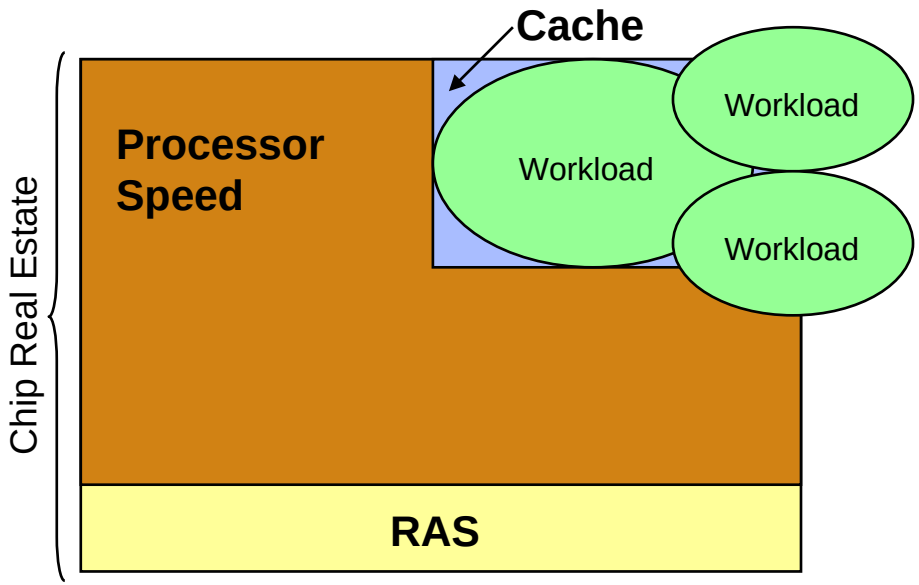

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- “Fast” processor speed is not fully realized due to cache misses

Consolidated Server Chip Design

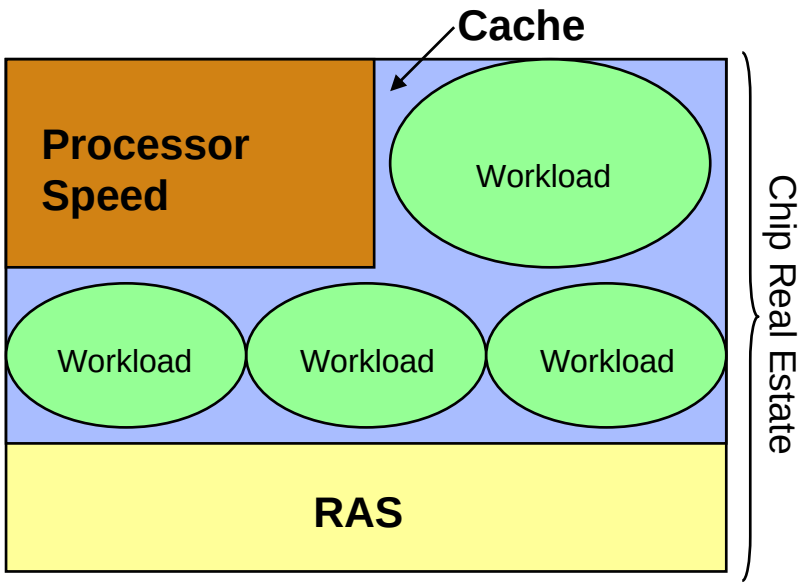

- System z cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

Note: System representations are not to scale, proportions may vary based on generation of chip and model

IBM System z Channel and I/O Architecture

The Big Picture

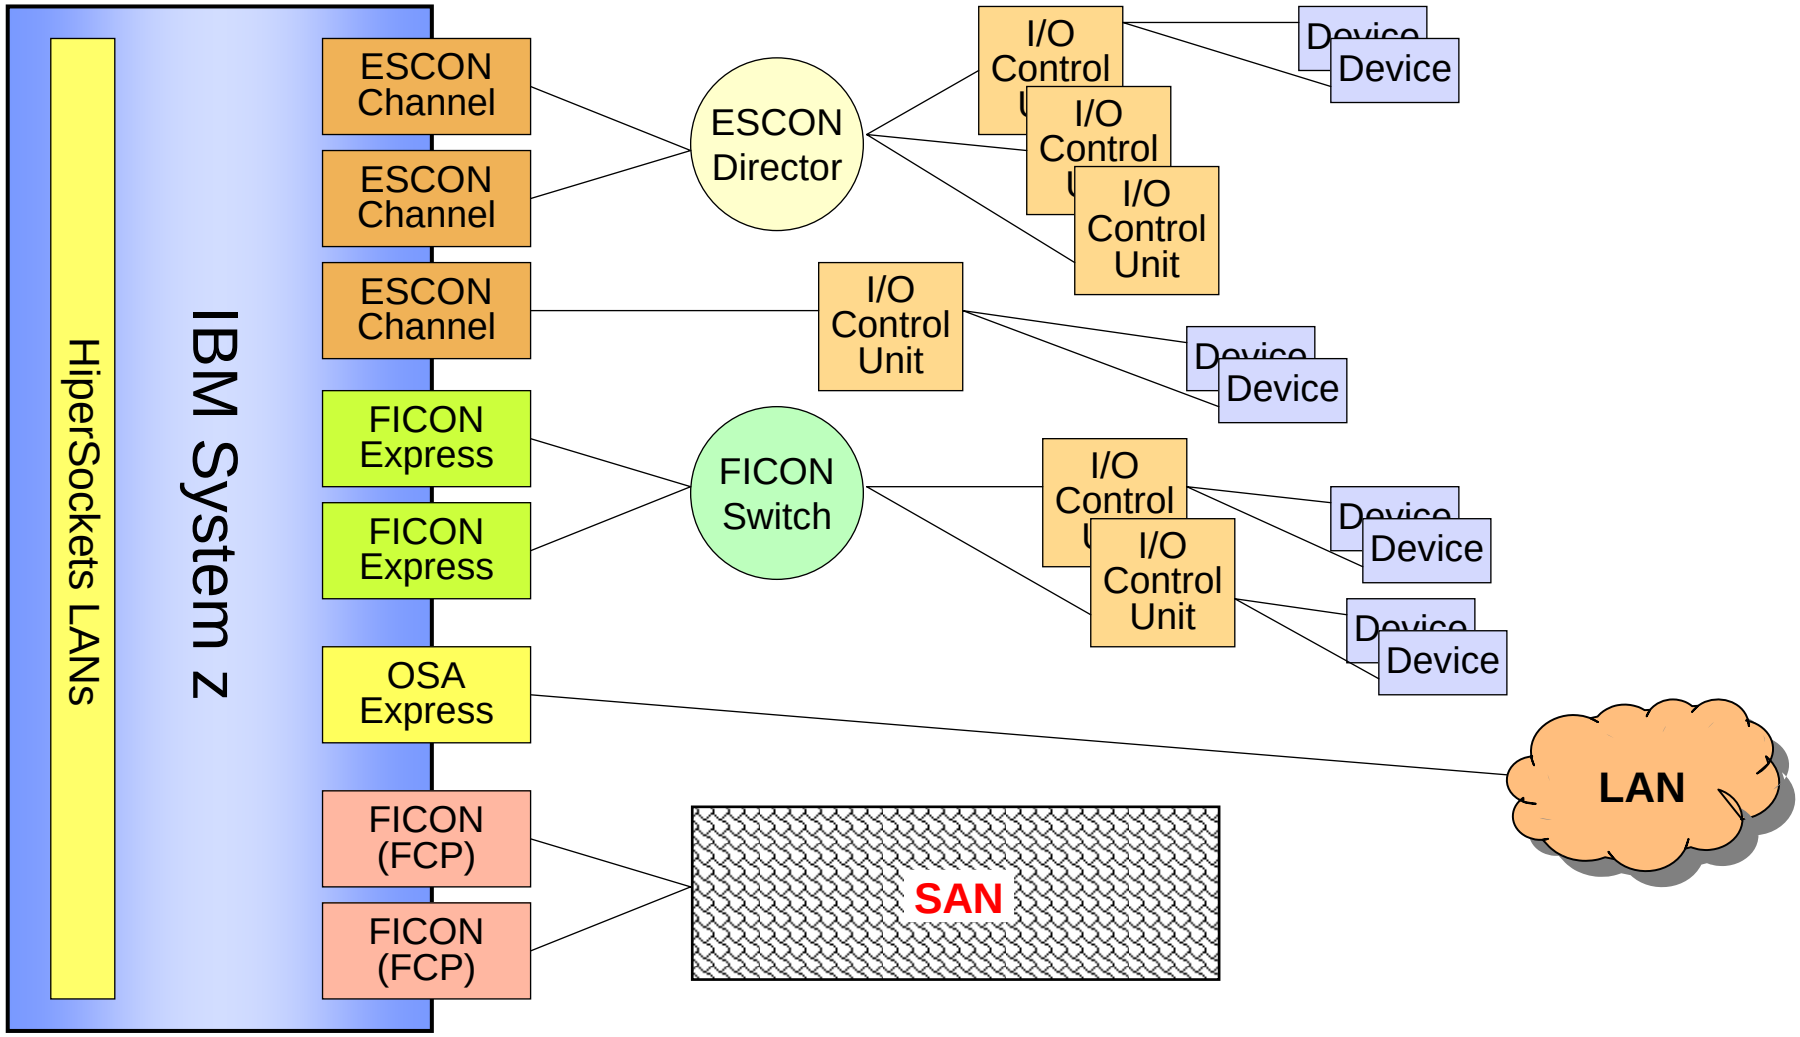

PR/SM High-Performance I/O Sharing (Multi-Image Facility)

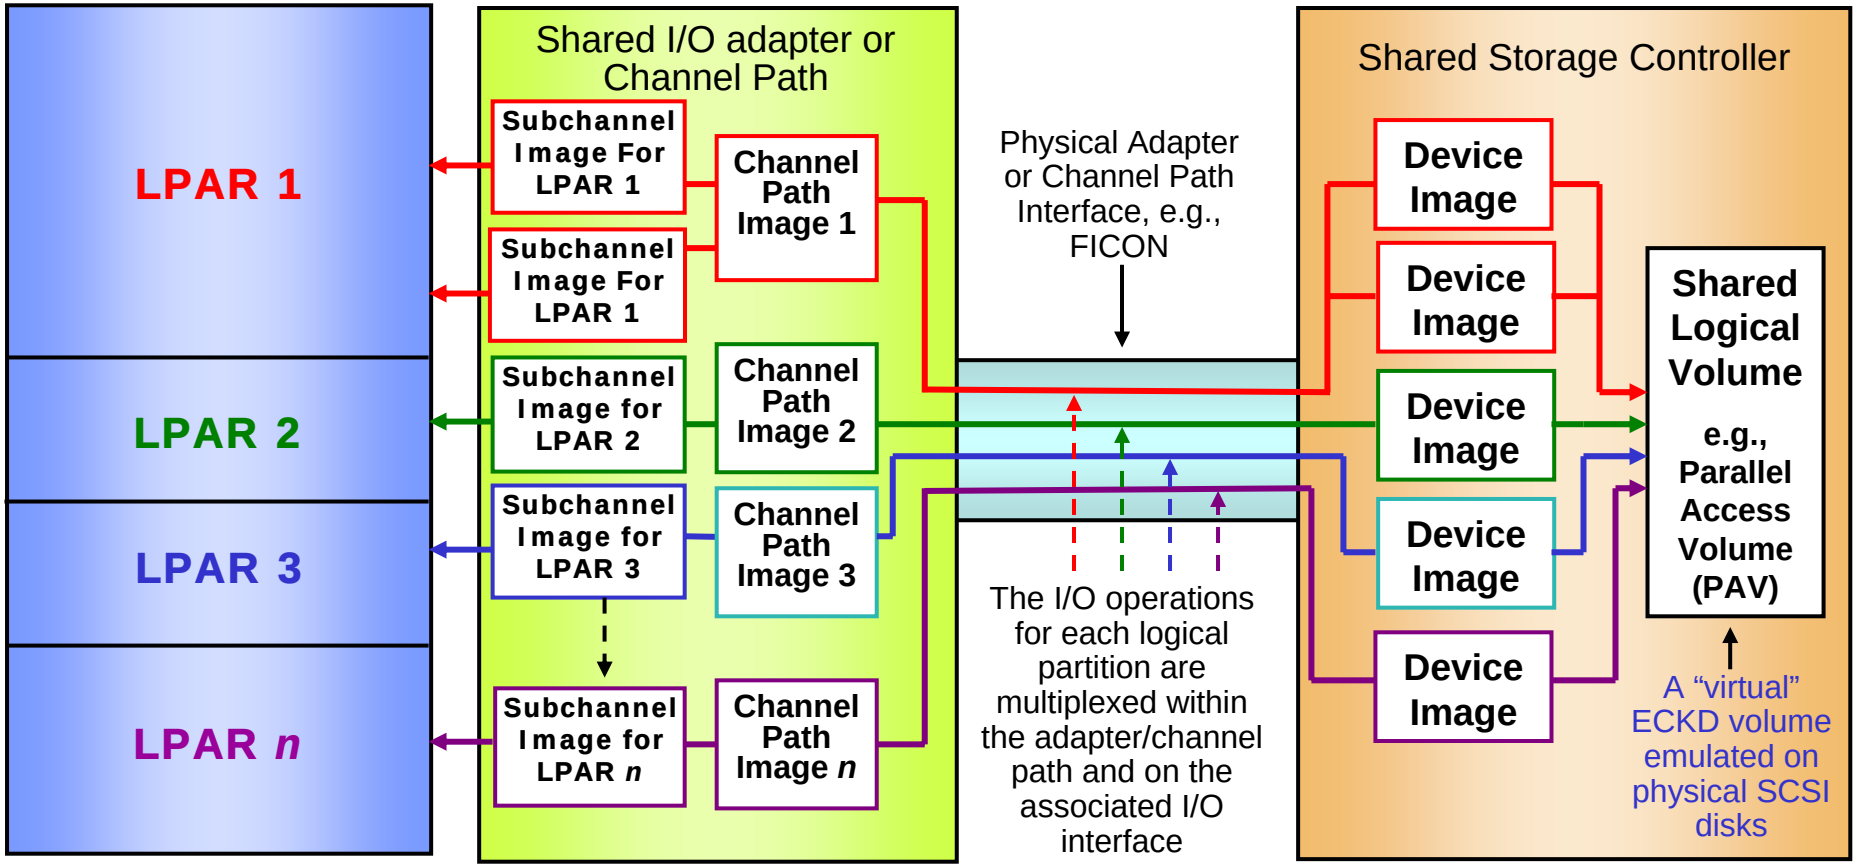

- The I/O infrastructure is shared by LPARs at native speeds, without hypervisor involvement
- Up to 8 physical channels process the I/O requests to the shared devices
  - This reduces the possibility of I/O queuing delays at the channels or at the shared storage controller

Extreme Virtualization with z/VM V5.3

z/VM can provision virtual machines with a mix of real and virtual resources with exceptional levels of scalability, availability and security

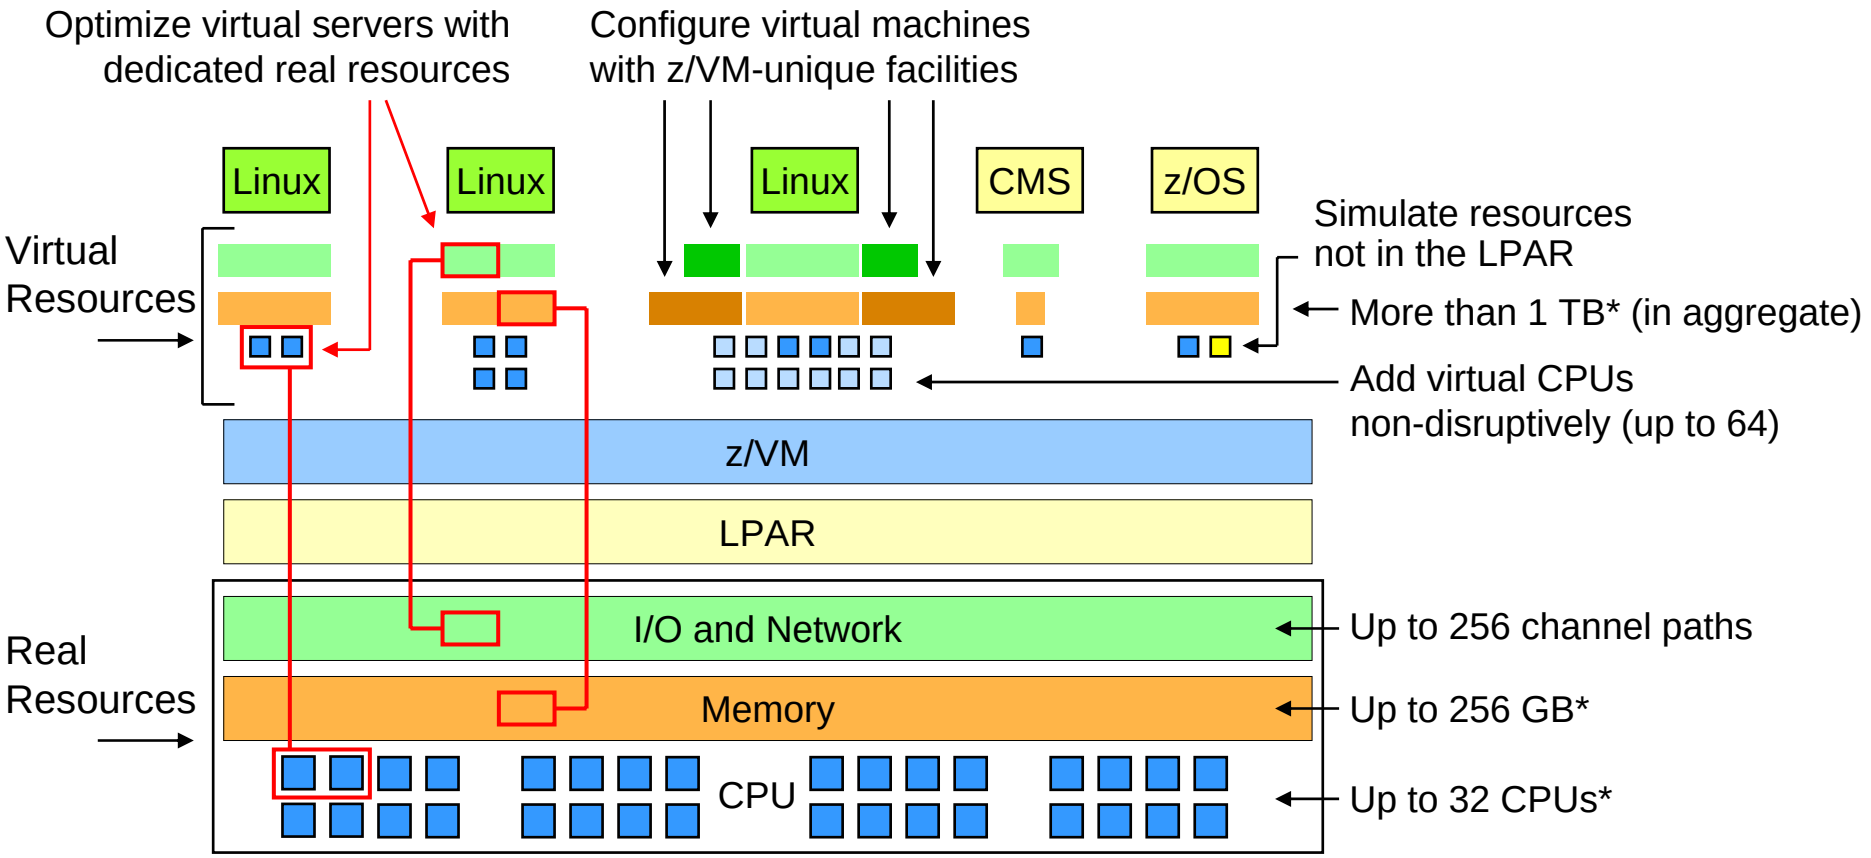

* z/VM V5.3 maximums

Extreme Virtualization with Linux on z/VM

VMRM Cooperative Memory Management (VMRM-CMM)

- Problem scenario: virtual memory utilization far exceeds real memory availability
- Solution: real memory constraint corrected by z/VM *Virtual Machine Resource Manager*
- Linux images signaled to reduce virtual memory consumption
- Demand on real memory and z/VM paging subsystem is reduced
- Helps improve overall system performance and guest image throughput

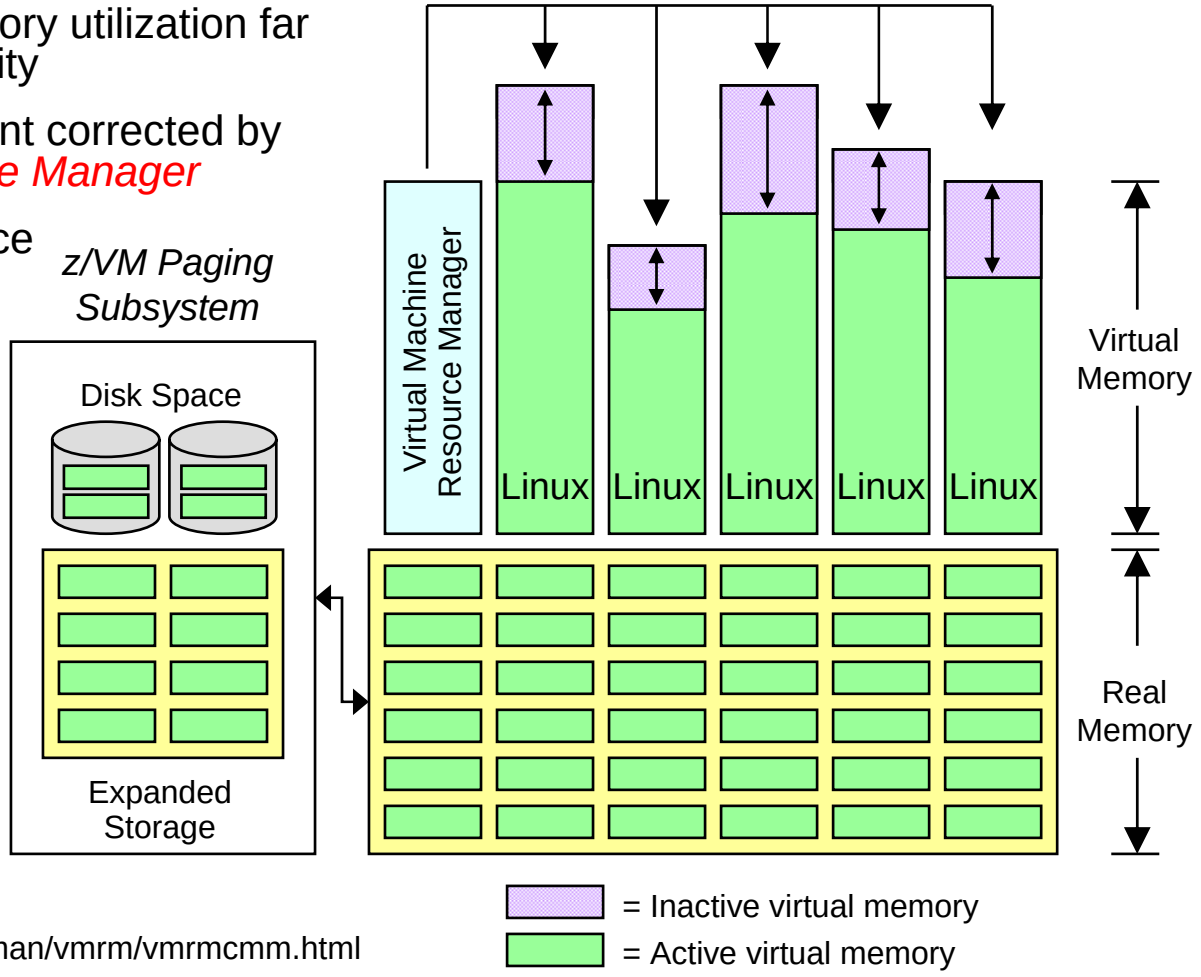

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

Extreme Virtualization with Linux on z/VM

VMRM Cooperative Memory Management (VMRM-CMM)

- Problem scenario: virtual memory utilization far exceeds real memory availability
- Solution: real memory constraint corrected by z/VM *Virtual Machine Resource Manager*
- Linux images signaled to reduce virtual memory consumption
- Demand on real memory and z/VM paging subsystem is reduced
- Helps improve overall system performance and guest image throughput

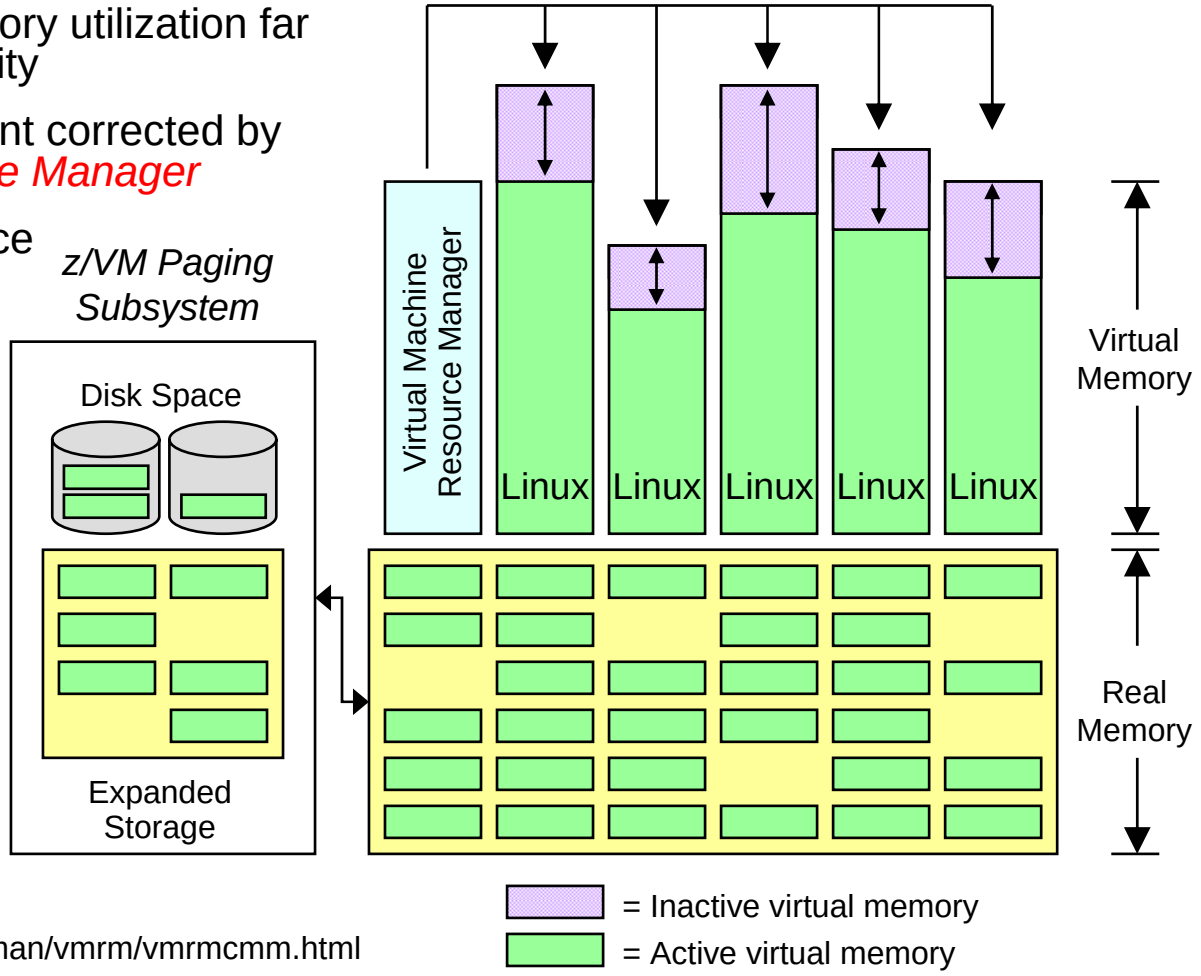

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

OLTP Database Environment with VMRM-CMM and CMMA

Excerpt from "z/VM Large Memory – Linux on System z" Whitepaper

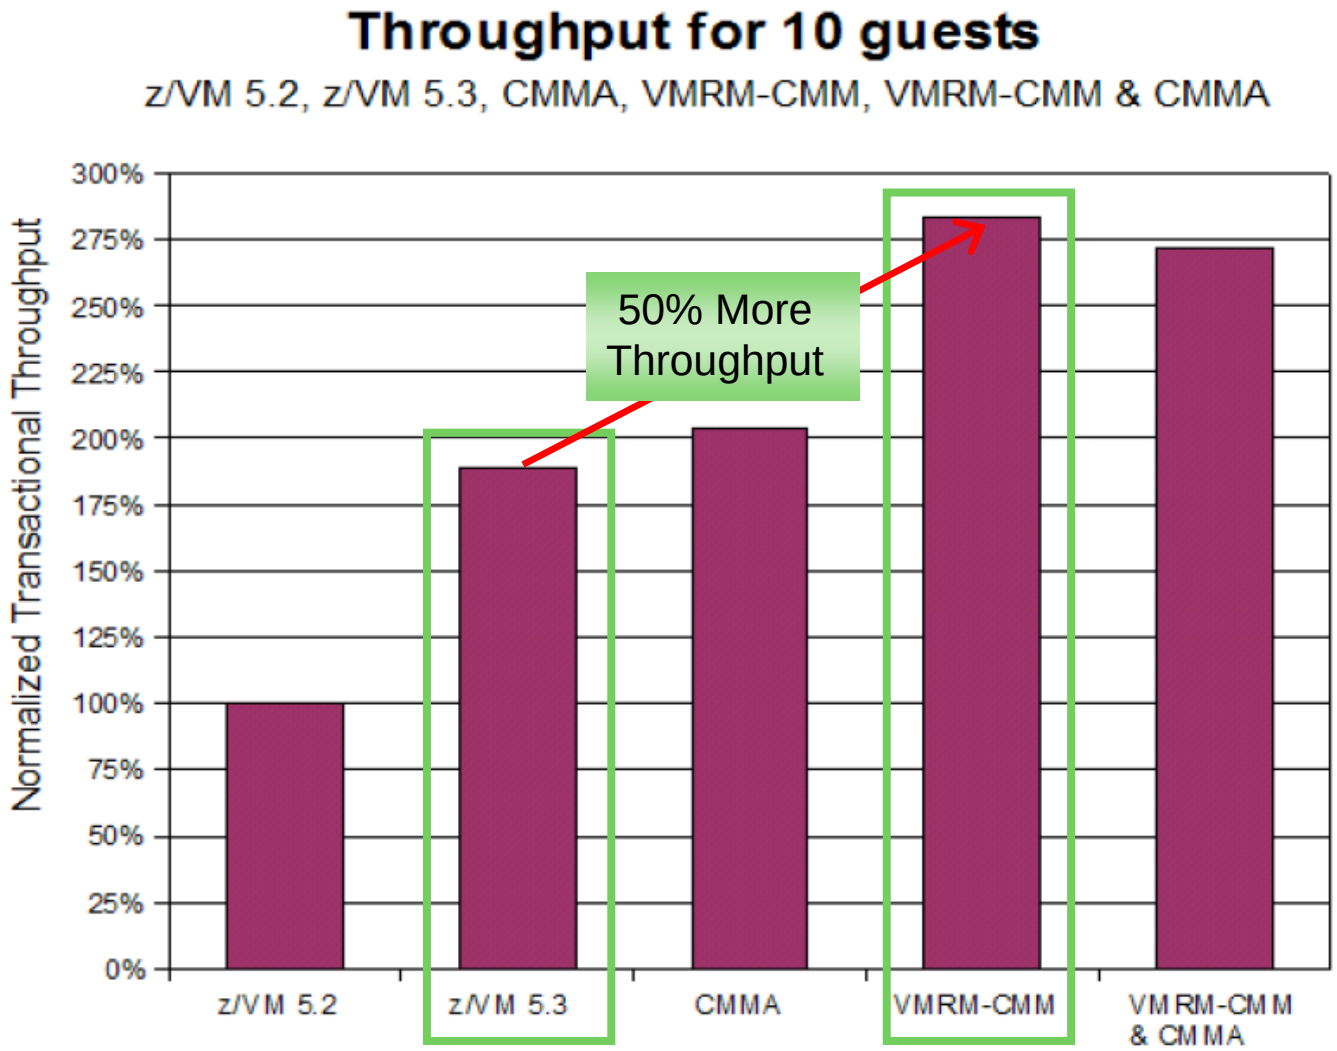

Extreme Virtualization with Linux on z/VM

Linux Exploitation of z/VM Discontiguous Saved Segments (DCSS)

- DCSS support is Data-in-Memory technology
  - Share a single, real memory location among multiple virtual machines
  - Can reduce real memory utilization
- Linux exploitation: shared program executables
  - Program executables are stored in an execute-in-place file system, then loaded into a DCSS
  - 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
  - Helps enhance overall system performance and scalability

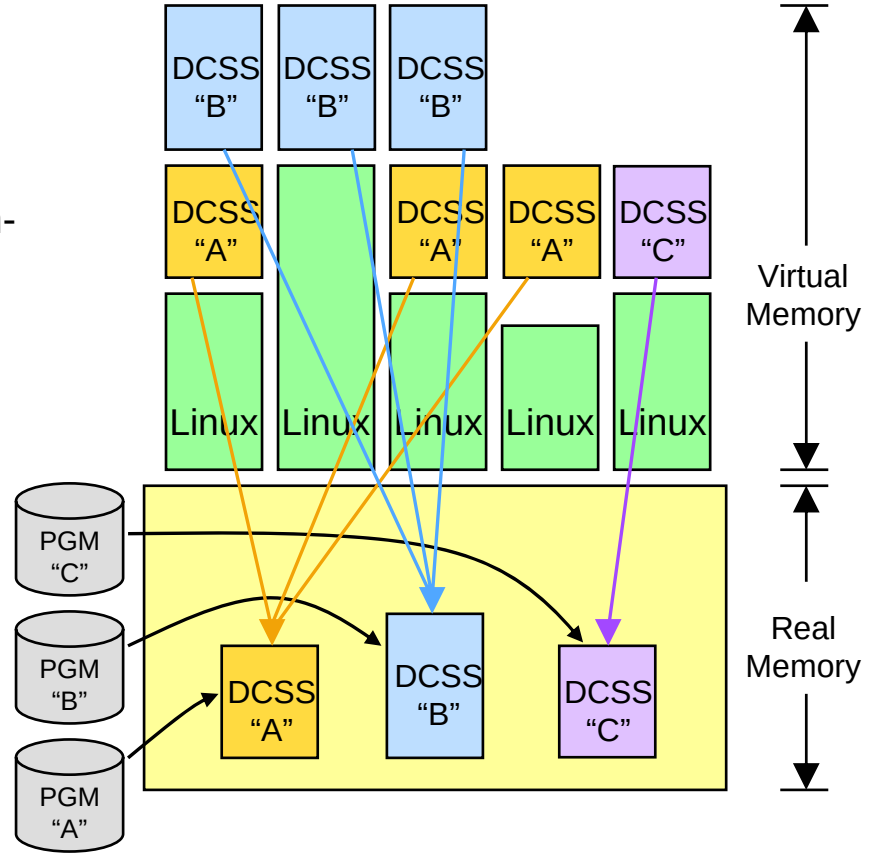

Learn more:
"Using DCSS/XIP with Oracle 10g on Linux for System z"
www.redbooks.ibm.com/redpieces/abstracts/sg247285.html

Extreme Virtualization with Linux on z/VM

Linux Exploitation of z/VM Virtual Disks in Storage (VDISK)

- VDISK support is Data-in-Memory technology
  - Simulate a disk device using real memory
  - Achieve memory speeds on disk I/O operations
  - VDISKs can be shared among virtual machines
- Linux exploitation: high-speed swap device
  - Use VDISKs for Linux swap devices instead of real disk volumes
  - Reduces demand on I/O subsystem
  - Helps reduce the performance penalty normally associated with swapping operations
  - An excellent configuration tool that helps clients minimize the memory footprint required for virtual Linux servers
  - Helps improve the efficiency of sharing real resources among virtual machines

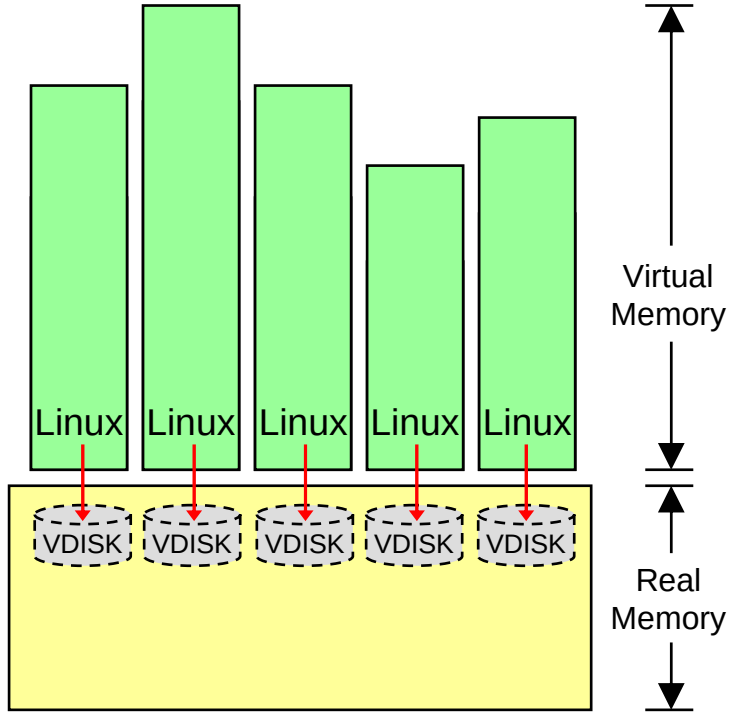

z/VM Virtual Switch Link Aggregation Support

Enhanced Networking Bandwidth and Business Continuance

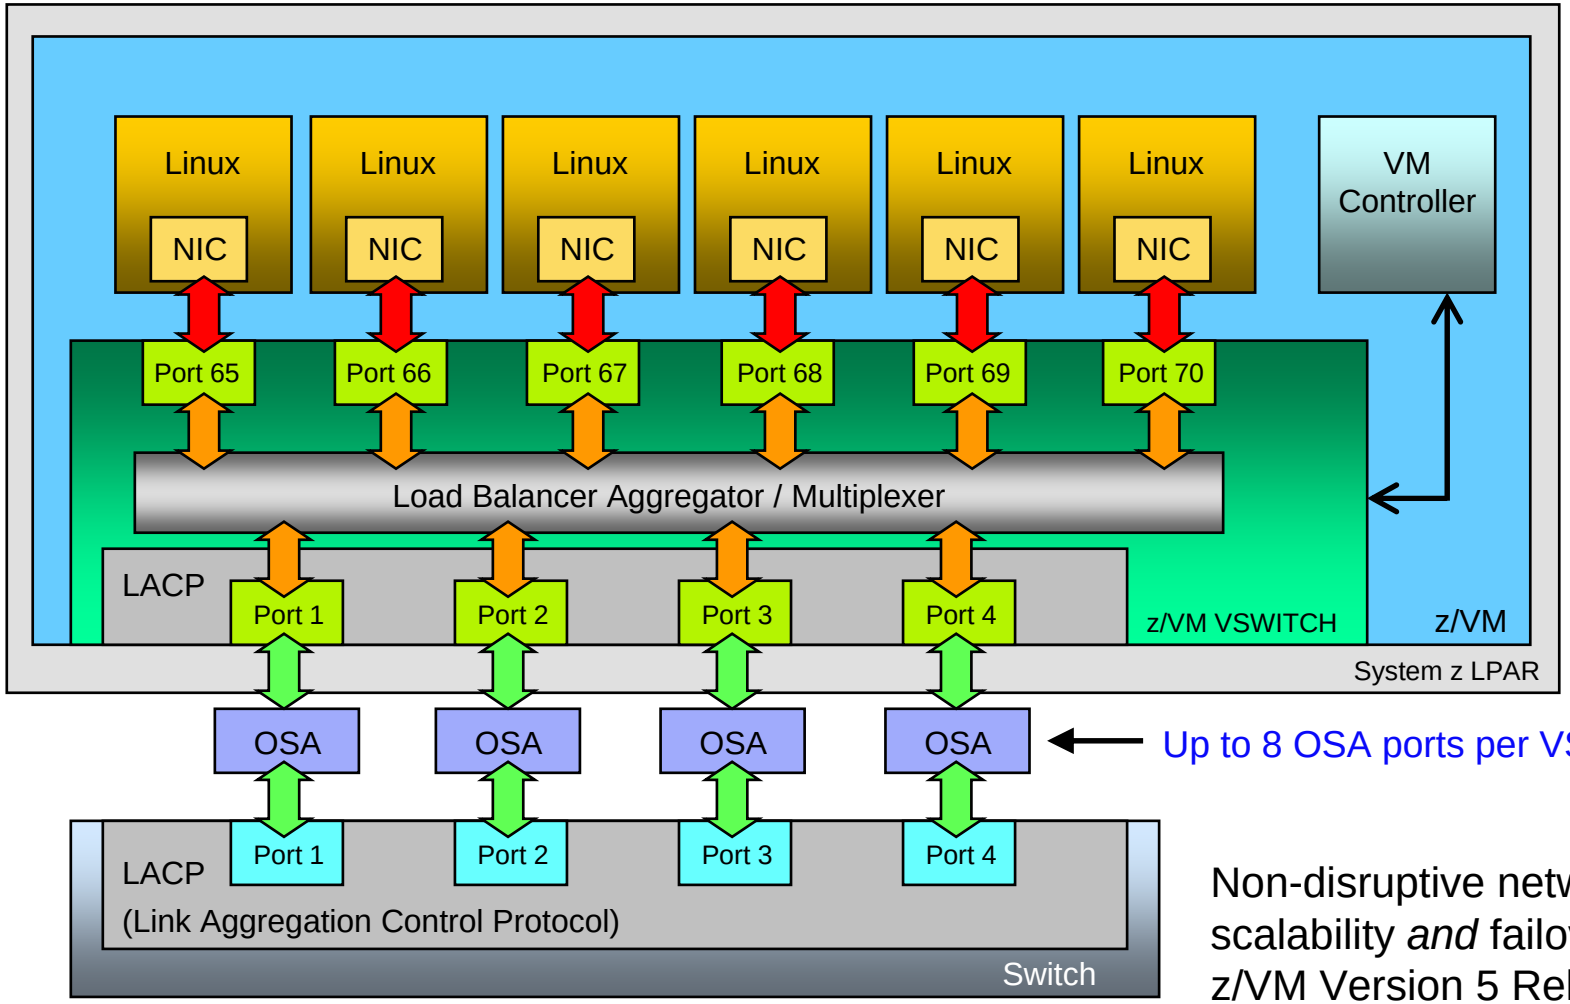

Non-disruptive networking scalability *and* failover with z/VM Version 5 Release 3

Note: Requires OSA-Express2 support available with IBM System z9 and later servers

z/VM Product Futures

- **Support for "VM-mode" LPARs**
  - Enabling clients to mix standard CPUs and IFLs in a single z/VM LPAR
  - Add IFLs to an existing z/VM standard-engine LPAR to host Linux workloads
- **Simplify end-users tasks for managing a z/VM-hosted virtual Linux environment**
  - Install Linux into a virtual machine using the HMC
  - Enhanced HMC GUI functionality for managing z/VM hypervisor-configuration tasks
- **Dynamic memory upgrade for z/VM LPARs**
  - Delivers added flexibility and virtual server availability to meet workload growth requirements
  - Add CPU, I/O, networking, *and* memory to a z/VM LPAR without disruption
- **Rapid data provisioning for test systems**
  - Generate test data quickly and affordably with space-efficient FlashCopy
- **Multi-system virtualization support**
  - Enhanced workload management and scalability across a cluster of z/VM images – improves workload availability and virtual server provisioning

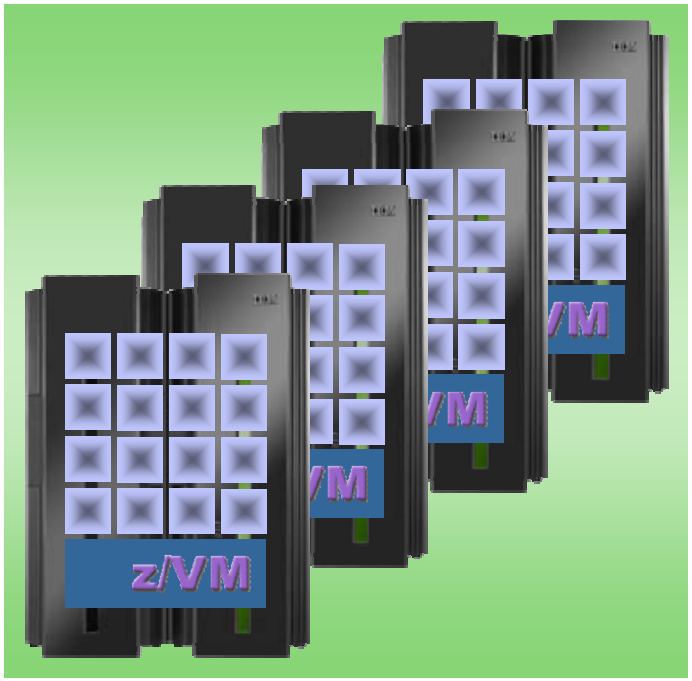

Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

What Kind of Workload Runs Best on Mainframes?

- **Consider *system* performance and capacity in selecting "best fit" for your applications**

  Data intensive workloads like large databases, transaction processing, object oriented code and context switching often run better on System z, due to its processor caching and I/O architecture.

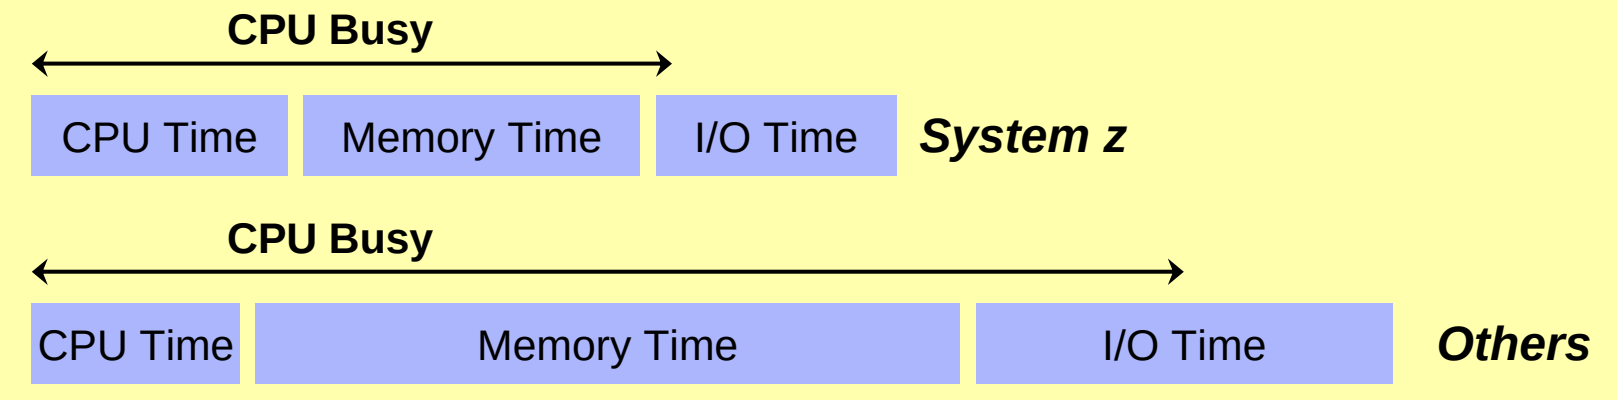

- **Also factor in the value derived from the co-residency of applications and data servers on a single mainframe**

Topics

- Business drivers for server consolidation
- System z: wired for virtualization
- **System z workload management: what every business needs**
- Saving money, growing your business

System z Workload Management

- **IBM System z workload management is expertly crafted to:**
  - Maximize the virtualization capabilities of an infrastructure already designed to share hardware assets
  - Align IT resources with business needs with a degree of certainty unmatched by other platforms
  - Minimize the time-consuming, low-value busy work associated with a server sprawl solution (e.g., “Security Tuesdays”)
  - Deliver tangible cost savings to the bottom line
    - Software license fees
    - Staff productivity
    - Cost-efficient support for business continuance

IBM System z Virtualization Workload Management

The Big Picture

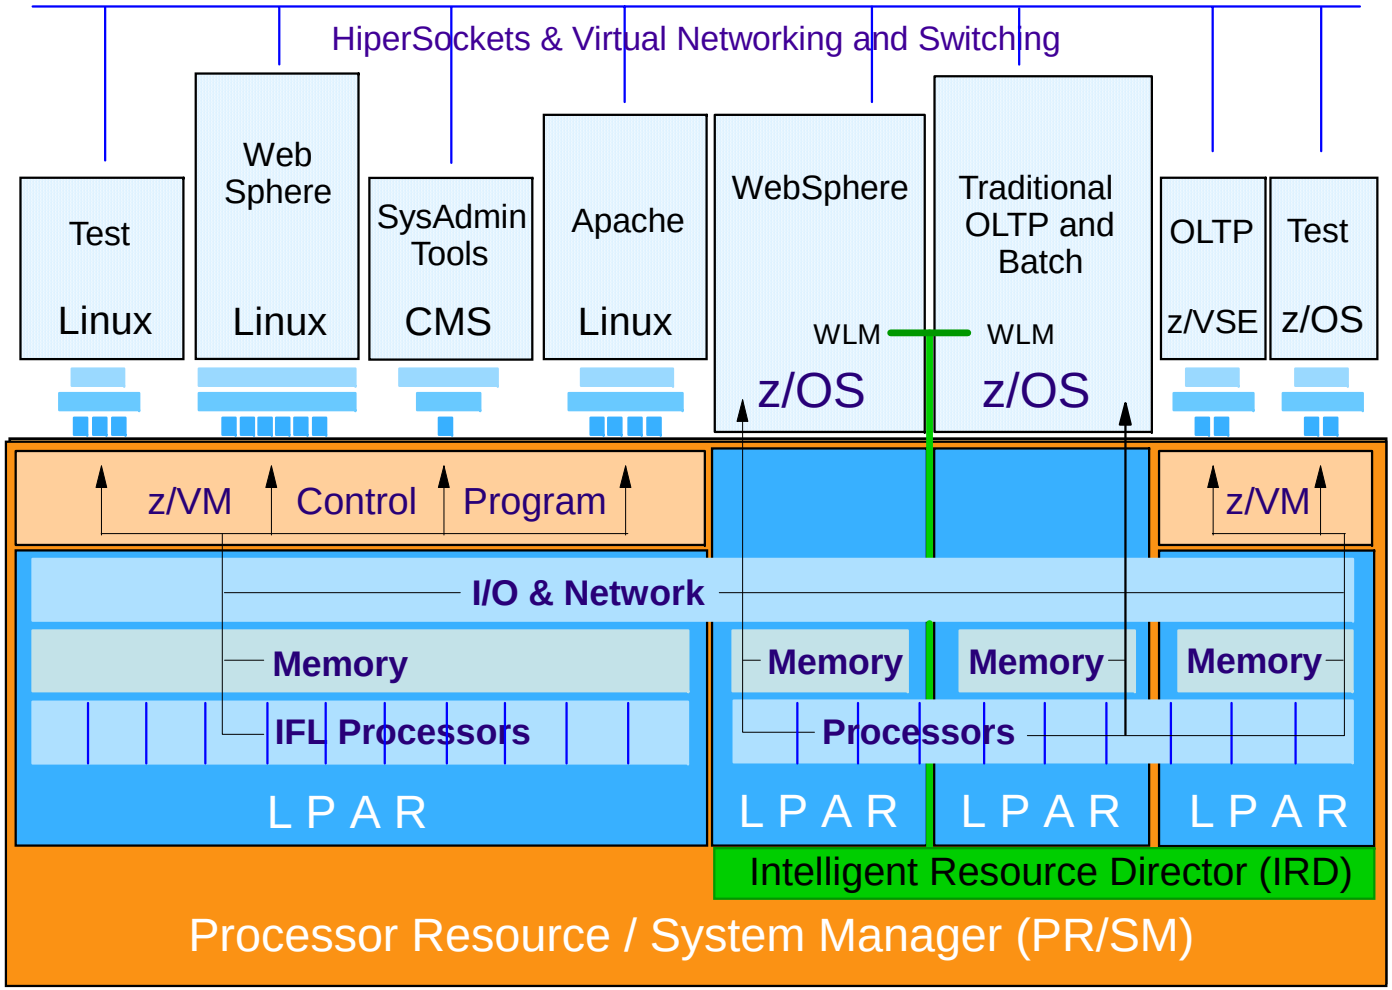

Multi-dimensional virtualization technology

- 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*
- Ideally suited for large-scale hosting of mixed workloads

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

z/OS Workload Management

Prioritizing work in a single z/OS image, across LPARs, across a Parallel Sysplex, to your network

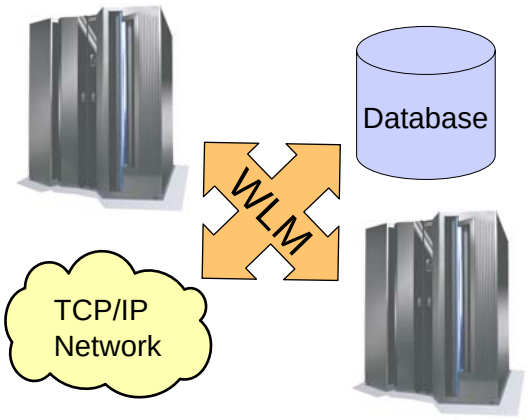

- Define performance goals in business terms and assign importance to each goal: manage to service level agreements
- The system decides how much resource is needed to meet the goals
- Resources are shared dynamically across mixed applications
- End-to-end prioritization capability, from the network to the data
- Workload Manager will monitor the system and adjust processing to meet the goals

With z/OS 1.8: more integration with EWLM (Enterprise Workload Manager), facilitating the end-to-end management of workloads.

A History of Advanced Technology	
Workload Manager (WLM)	1994
Parallel Sysplex	1994
Sysplex Data Sharing	1994
TCP/IP Sysplex Distributor	2000
Intelligent Resource Director (IRD)	2001
Transactional VSAM (DFSMStvs)	2002
TCP/IP Sysplex Health Monitoring	2004
Cross platform monitoring (EWLM)	2004
WebSphere Application Server	2004
DB2 Stored Procedures / Latches	2005
Sysplex Distributor Coordination	2005
Load Balancing Advisor	2005
Support for zAAP	2005
Group Capacity Limits	2006
Support for zIIP	2007

z/VM CPU Resource Controls

Highly Granular Sharing of System Resources

- Allocate system resources per guest image using 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“
  - Also use VMRM to prioritize I/O operations among guest images via “I/O Priority Queuing”

```
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
```

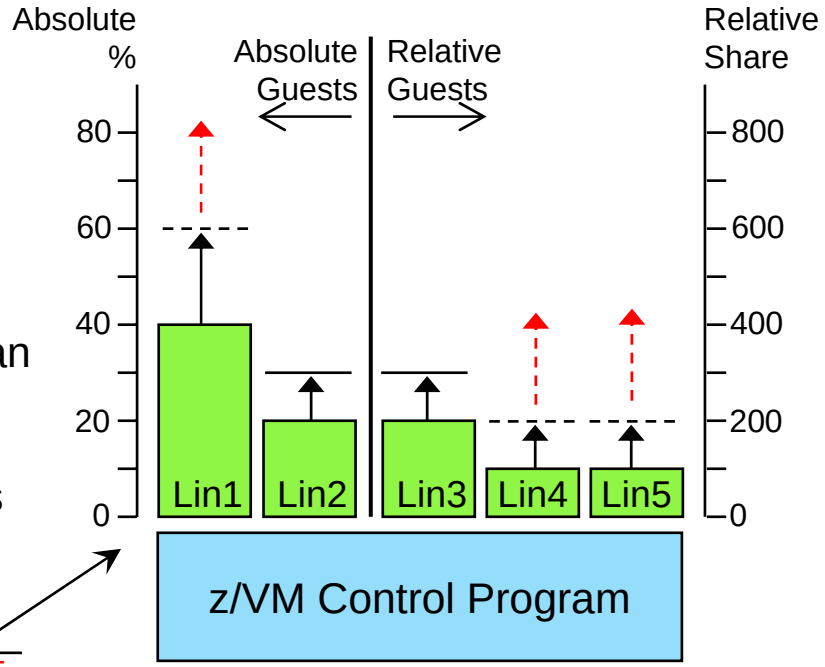

Notes:

----- = limit can be exceeded if unused capacity is available (LIMITSOFT)

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

IBM System z Application Hosting

Offering Clients Industry-Leading Virtualization Options

- z/OS: large-scale application integration with transaction and data services
- Linux-on-z/VM: large-scale virtual server hosting side-by-side z/OS systems
- z/OS and Linux: mix and match as business demands dictate
- Consider WebSphere Application Server for z/OS and Linux
  - Both adhere to the WebSphere development model and tools
  - Both adhere to WebSphere systems management and admin model and tools
  - Each provides unique value for WebSphere applications

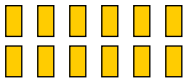

- ***The Linux value proposition***: simplify and optimize existing infrastructure for end-to-end WebSphere applications with the goal of reducing costs and complexities

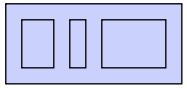

- ***The z/OS value proposition***: provide highest possible qualities of service (QoS) in an efficient, cost-effective manner via integration with z/OS WLM

Topics

- Business drivers for server consolidation
- System z: wired for virtualization
- System z workload management: what every business needs
- **Saving money, growing your business**

IBM System z: The Ultimate Virtualization Platform

- ***Virtualize*** **everything with up to 100% utilization rates**
  - CPU, memory, network, I/O, cryptographic features, coupling facility, ...

Consolidate all types of workloads

- ***Massively scale*** **your workload on a single System z mainframe**
  - The Linux-on-z/VM record is 97,943 virtual machines
  - Each virtual machine on z/VM can access up to 24,576 devices

Smart economics: start small and grow big in the same box

- ***Non-disruptively add*** **anything**
  - Up to 64x CPU scalability per mainframe, 32x scalability per z/VM LPAR
  - z/VM is designed to support more than 1 TB of active virtual memory

Rapidly respond to workload spikes

- ***Security*** **for everything**
  - Highest security classification for general purpose servers in the world
  - System z LPAR technology is EAL 5 certified

Secure your virtual servers and reduce business risk

- ***Optimize and integrate*** **it all with the IBM software portfolio**

Increase staff productivity and virtualize the enterprise

IBM Tivoli Virtualization Management for System z

Helping Clients Manage and Control Their Virtualized IT Infrastructure

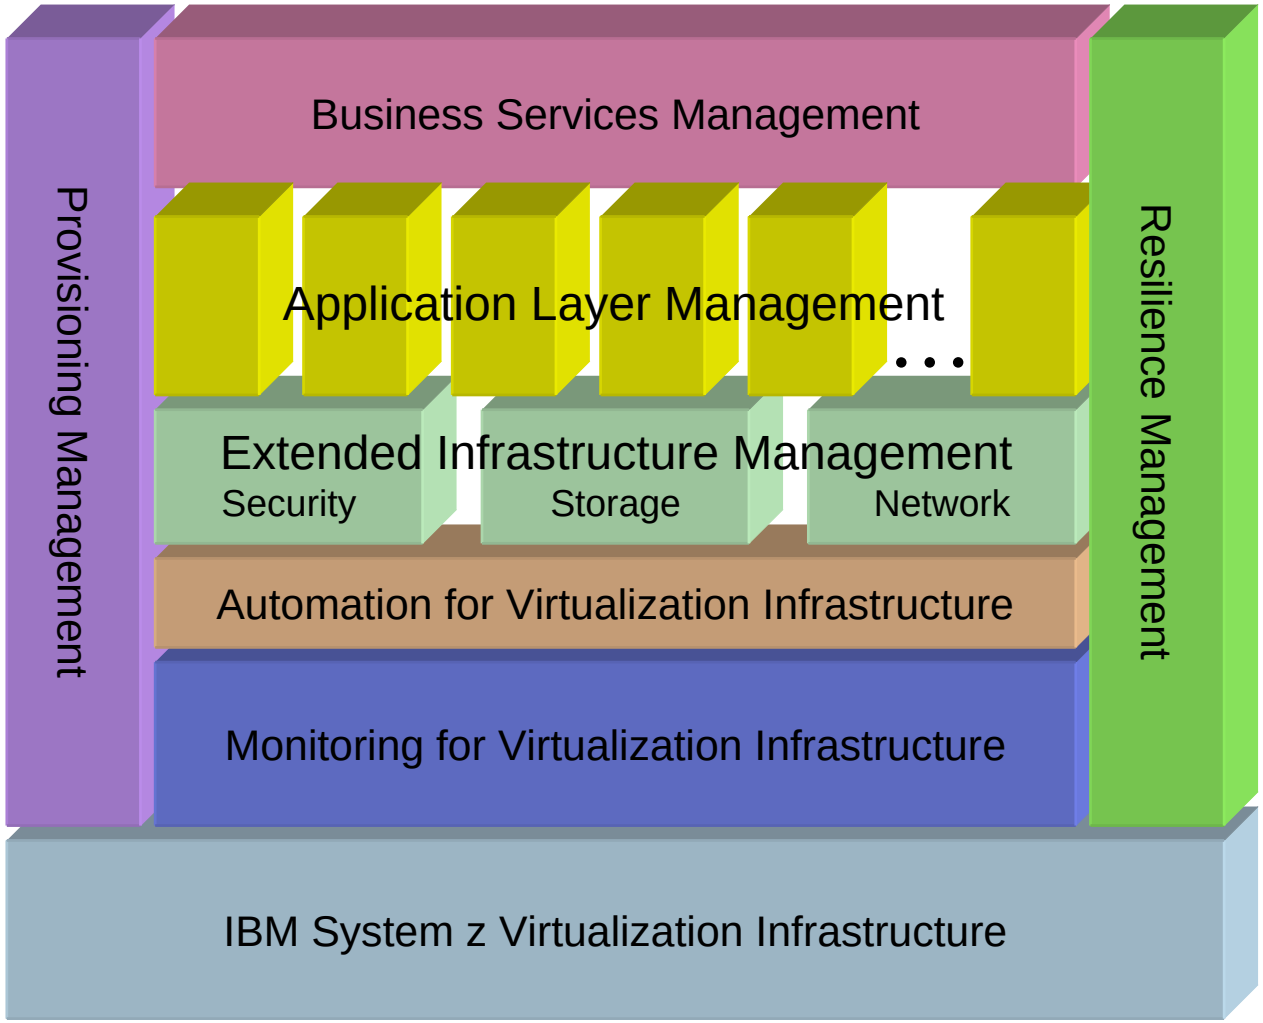

Why Virtualize on System z Now?

New HW / SW spending

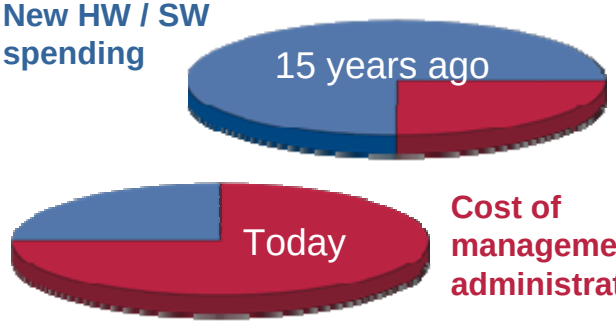

Cost of management & administration

Source: Tony Picardi, IDC
Economist.com: Make it simple. October 28th, 2004
From The Economist print edition

System z9 Managing Growth and Complexity

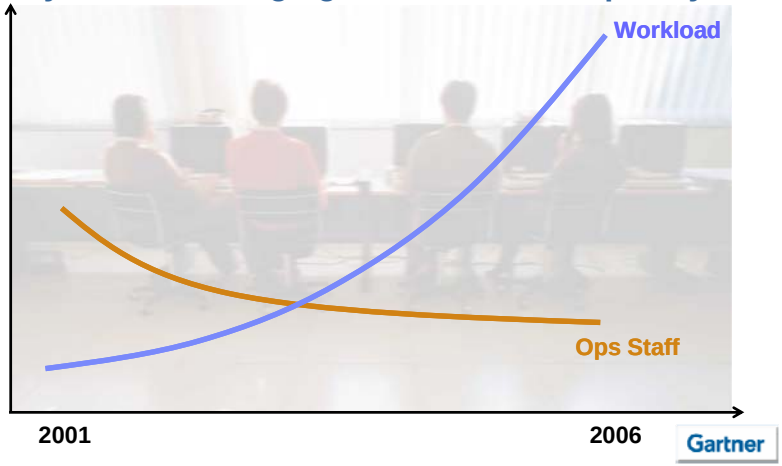

Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration

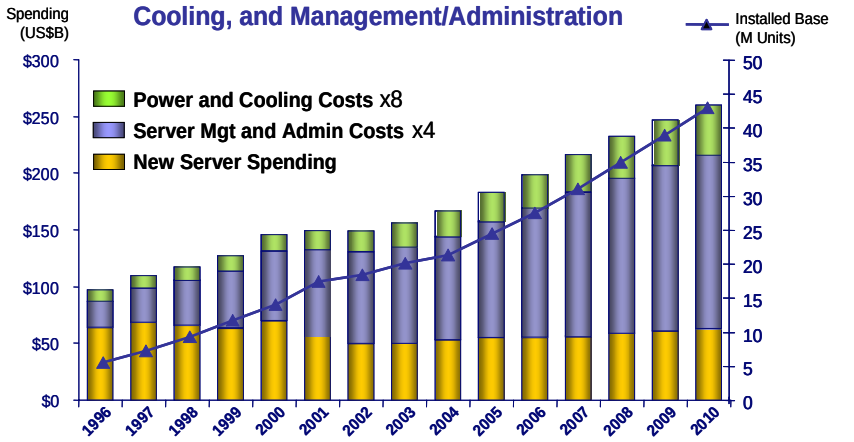

Many servers, much capacity, low utilization = $140B unutilized server assets

Source: IDC, 2006

POWER:
System z9 vs. Linux on Intel

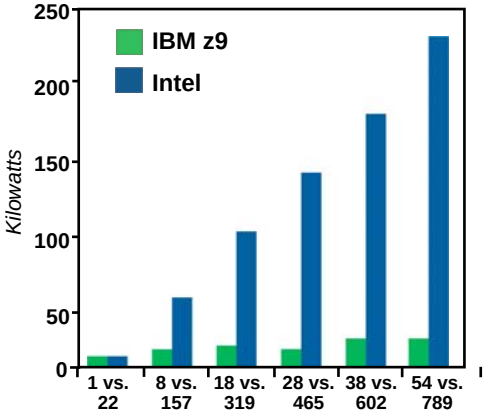

SPACE:
System z9 vs. Linux on Intel

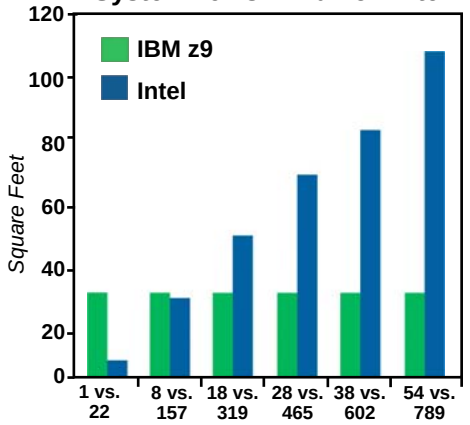

The Linux on Intel servers selected in this example are functionally eligible servers considered for consolidation to a System z running at low utilization such that the composite utilization is approximately 5%. The utilization rate assumed for System z EC is 90%. This is for illustration only actual power and space reductions, if any, will vary according to the actual servers selected for consolidation.

Secure Virtualization Changes Operational Model

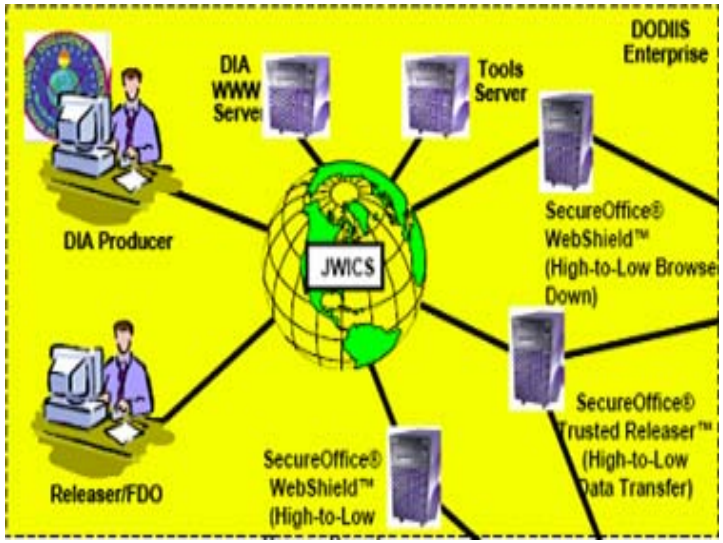

- **Cross Domain Database:**
  - Provide real time access to data avoid batch delays
  - collaboration across communities
- **Cross Domain Presentation Client**
  - Reduces desktop clutter
  - Reduce power consumption
  - Reduces leak potential with central mgt

Near-linear scalability	up to 900,000+ concurrent users; TBs of data
"Mean Time Between Failure"	measured in decades versus months
¼ network equipment costs	virtual and physical connectivity
1/25th floor space	400 sq. ft. versus 10,000 sq. ft
1/20 energy requirement	$32/day versus $600/day
1/5 the administration	< 5 people versus > 25 people
Highest average resource utilization	Up to 100% versus < 15%
Capacity Management & upgrades	On demand; in hours, not weeks/months
Security intrusion points	Reduced by z architecture and # of access pts.
Higher concurrent workload	hundreds of applications versus few

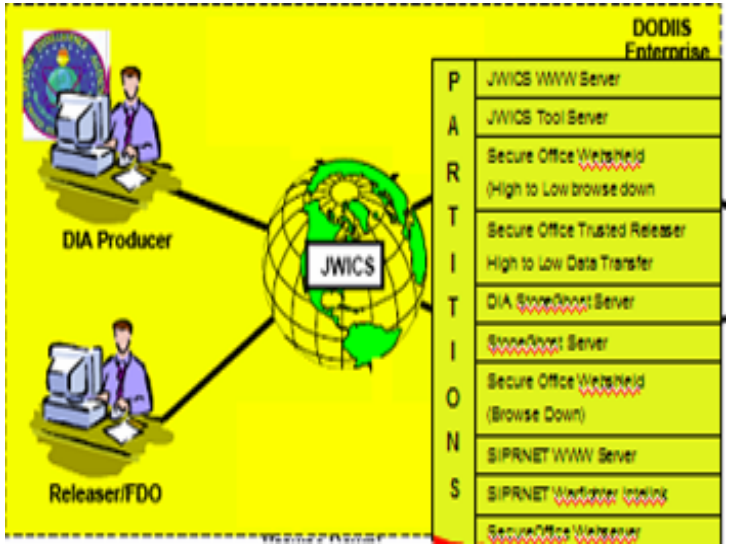

Opportunities for Cost Savings

- **Overcommitment of CPU resources can reduce software license fees**
- **Large-scale virtual server deployment on a single z/VM hypervisor can greatly enhance staff productivity**
- **Reliability and redundancy of System z infrastructure helps lessen application outages**
- **Flexible configuration options for business continuance (e.g., Capacity Backup on Demand)**
- **Cost-attractive economic model for technology refreshes (e.g., specialty engines carry forward to next generation)**

Extreme Virtualization with System z
Understanding the Value Proposition

- **Business pain points addressed by server virtualization:**
  - Underutilized IT assets
  - Environmental costs
  - Linear software costs per server image
  - Staff inefficiencies managing multiple real servers
  - Spiraling people costs
- **x86 virtualization pain points addressed by System z**
  - Virtual server workload management
  - Reliable high-bandwidth I/O virtualization
  - Virtual server and total system performance reporting and planning
  - Virtual server reconfiguration outages
  - Virtual machine security and integrity
  - Server sprawl with added complexity

Clients need to develop an enterprise-wide virtualization strategy that leverages the strengths of mainframe virtualization

The Power and Flexibility of System z Virtualization

- Over 40 years of continuous innovation in virtualization technologies
- Multiple images concurrently share all physical resources
- Resources delivered as required, automatically, based on business-oriented goals
- New OS images can be started without affecting ongoing work
- Hardware assists used to accelerate virtualization operations (e.g., SIE)

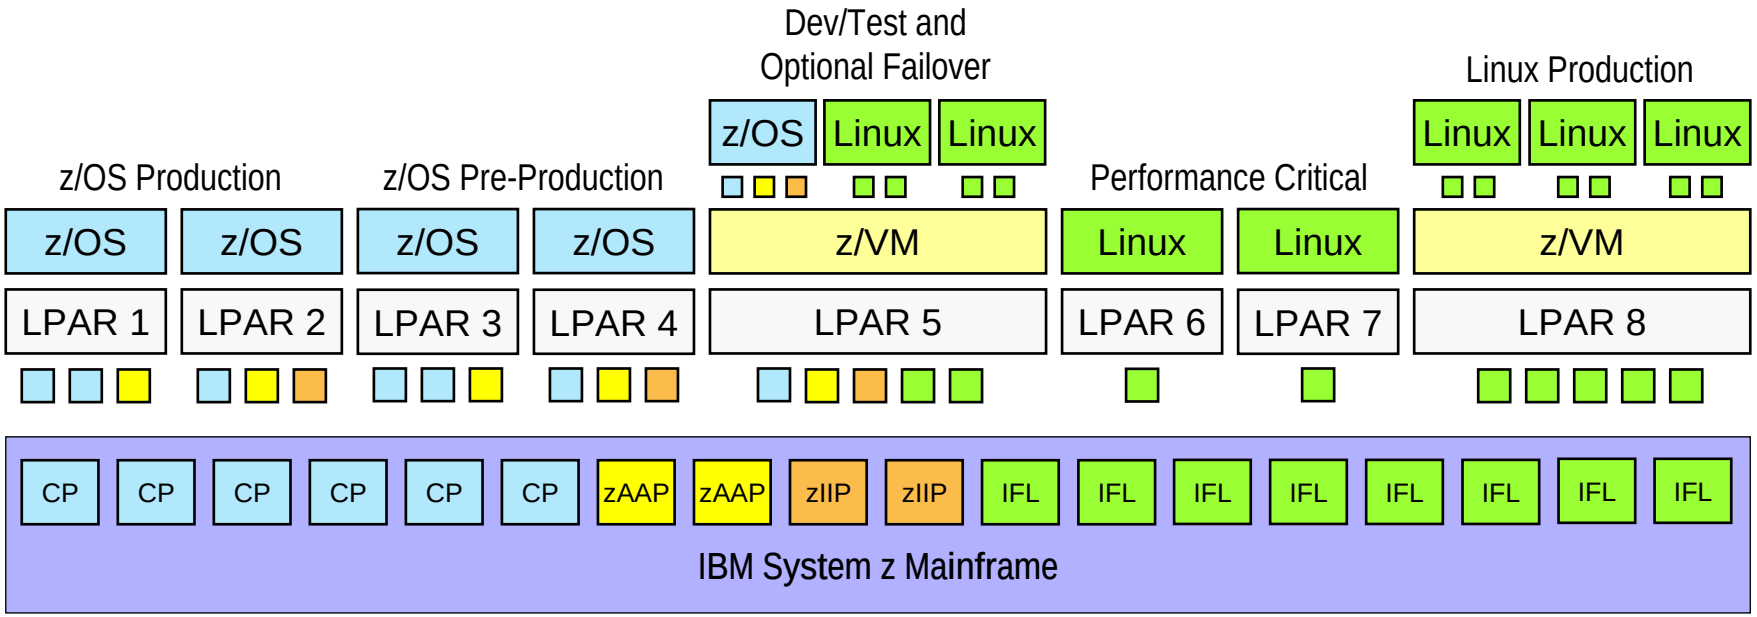

Questions