Siegfried Langer – Business Development Manager z/VSE & Linux on System z
April 16, 2011

Extreme Virtualization
with z/VM and Linux on System z

Abstract

Extreme Virtualization

Virtualization technologies are becoming a common technology to increase flexibility and resource utilization. However, there are big differences between implementations. The session will focus on the advanced virtualization capabilities of z/VM as the hypervisor for Linux on System z. Differences to other virtualization technologies like VMware or Hyper-V will be covered in general.

Resulting Sprawl Drives Unsustainable Costs

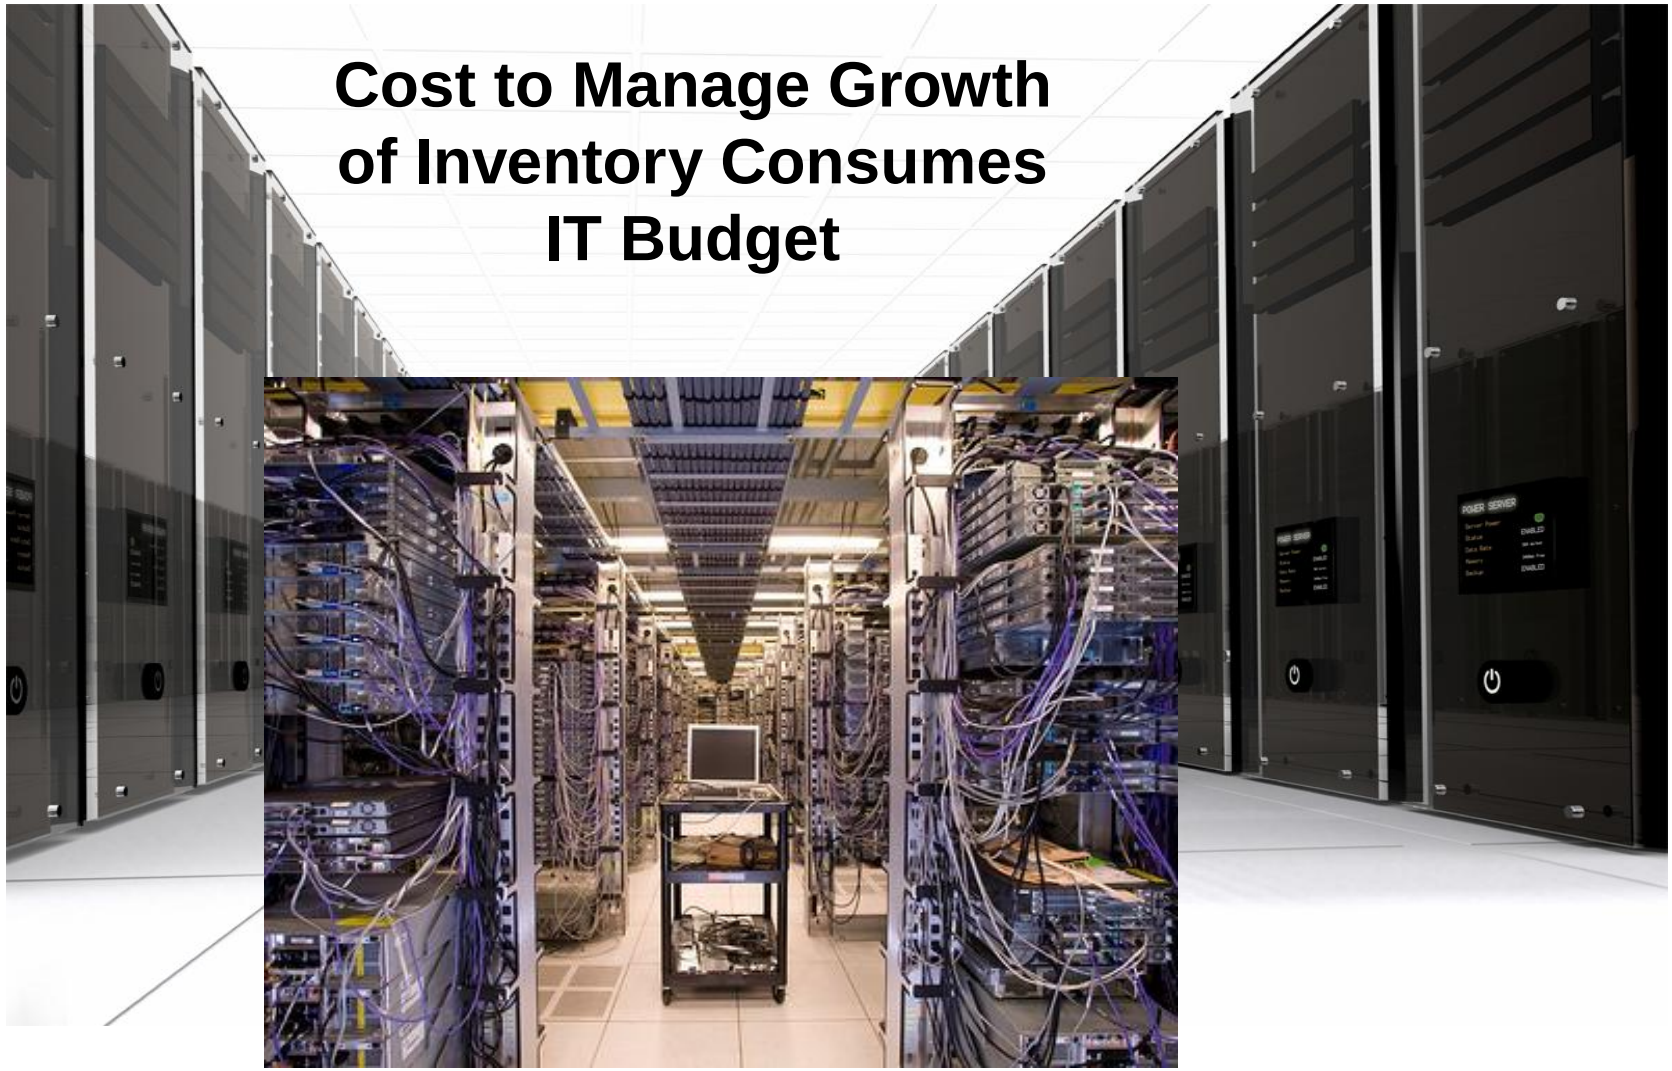

IT Operating Costs Are Out of Control

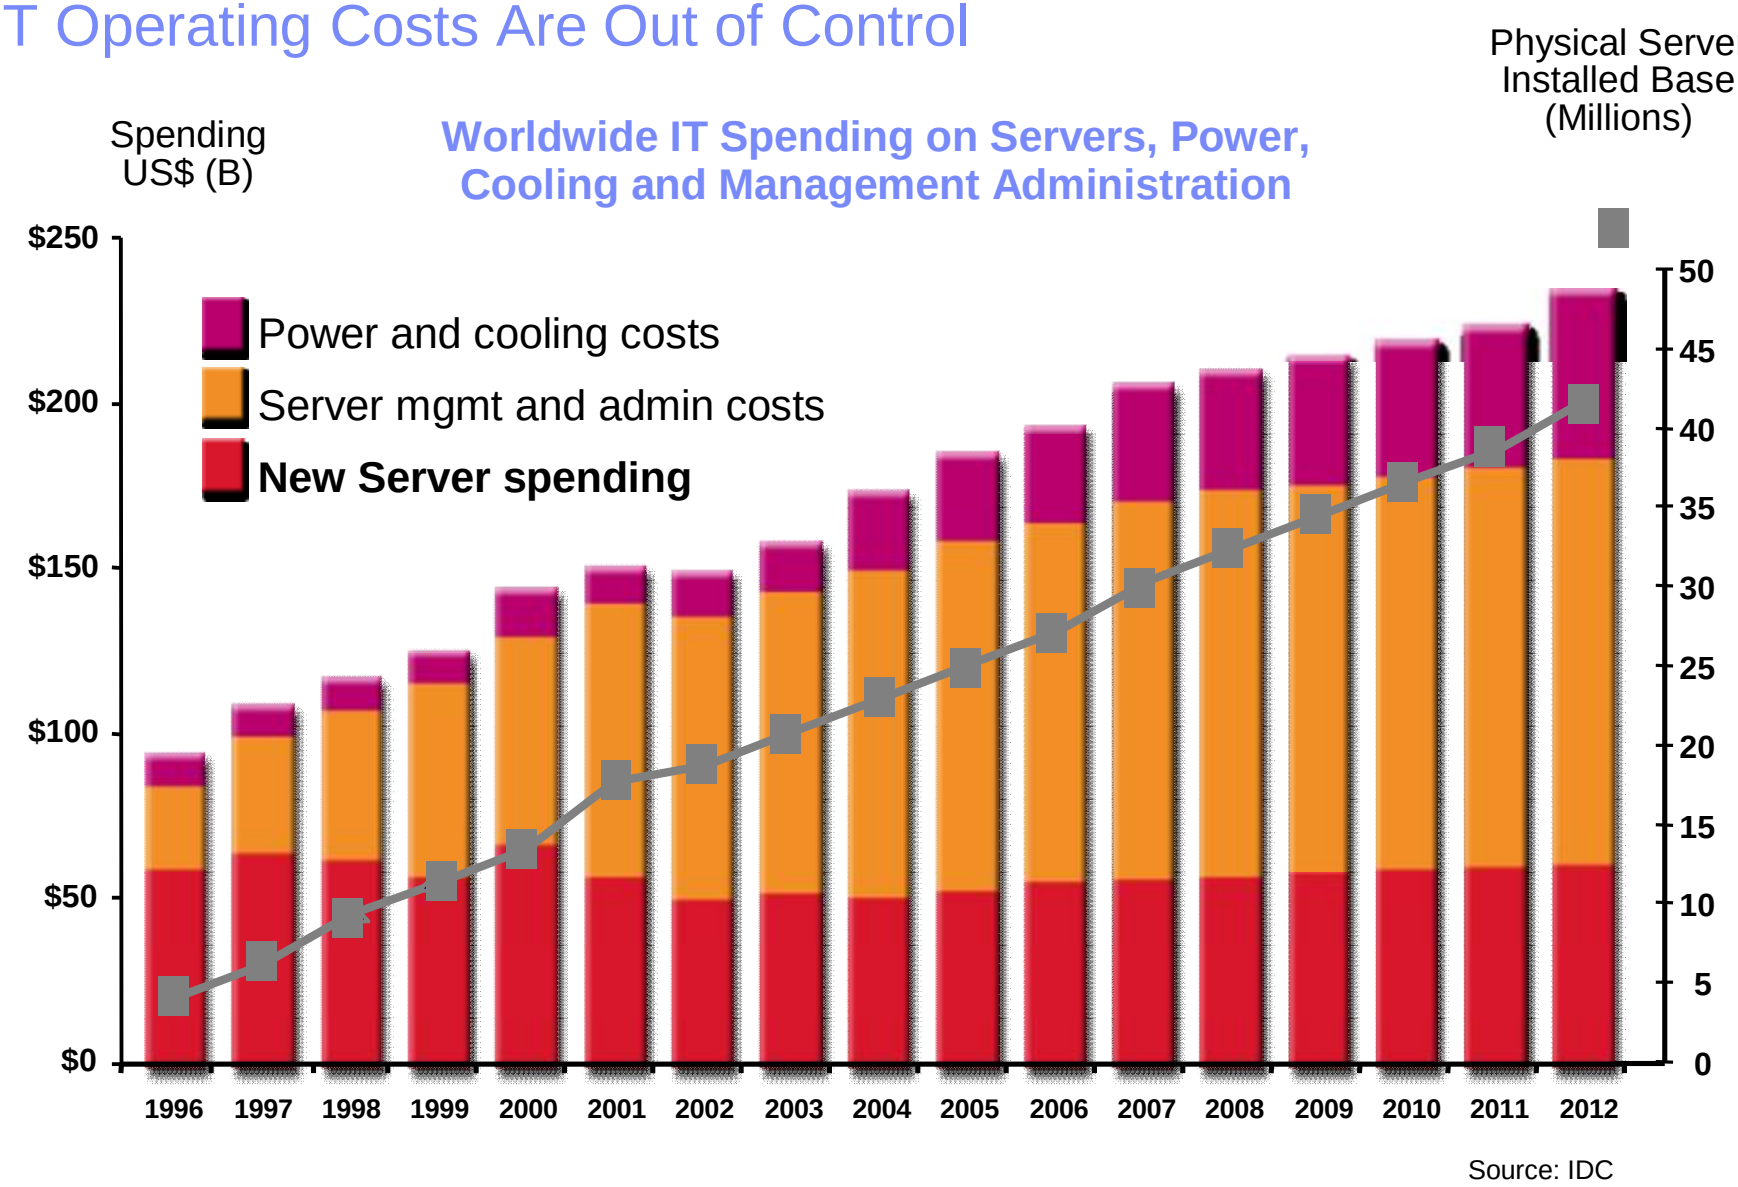

Source: IDC

IT Optimization and Consolidation

IBM 2009 CIO survey results

CIOs select their ten most important visionary plan elements

76%

of CIOs cited *"implementing a virtualized computing environment"* as part of their visionary plans to enhance competitiveness.

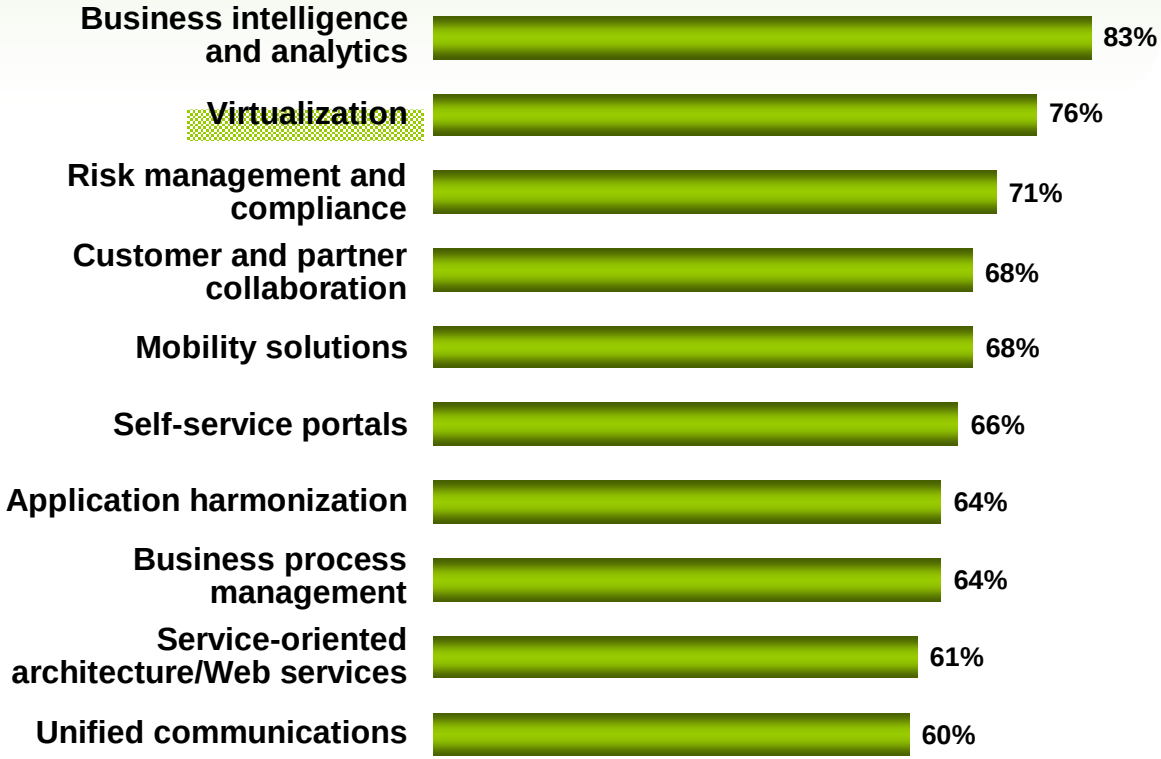

Note: CIOs were asked to select all applicable answers to the question, *"What kind of visionary plans do you have for enhanced competitiveness?"*

Strategies to Reduce Costs and Improve Value

Optimize the Overall IT Environment

Consolidate Hardware Infrastructure

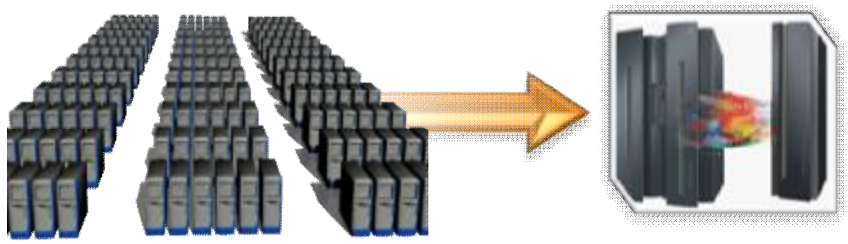

Eliminate Redundant Software and Data

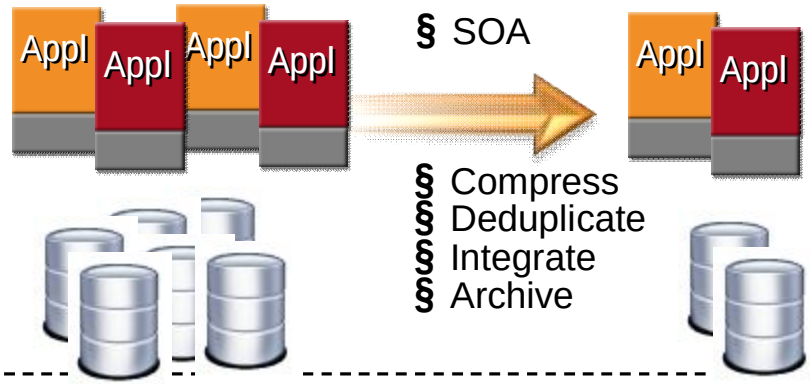

Improve Service Delivery

Integrated Service Management

Visibility Control Automation Cloud Computing

What is Virtualization?

Logical representation of resources not constrained by physical limitations

- Enables user flexibility
- Centrally manage many resources as one
- Dynamically change and adjust across the infrastructure
- Create many virtual resources within single physical device
- Eliminates trapped capacities

IBM Virtualization

A comprehensive platform to help virtualize the infrastructure

Utilization of Distributed Servers

Provision for expected growth

Provision capacity for peak workload

Server dedicated to one application

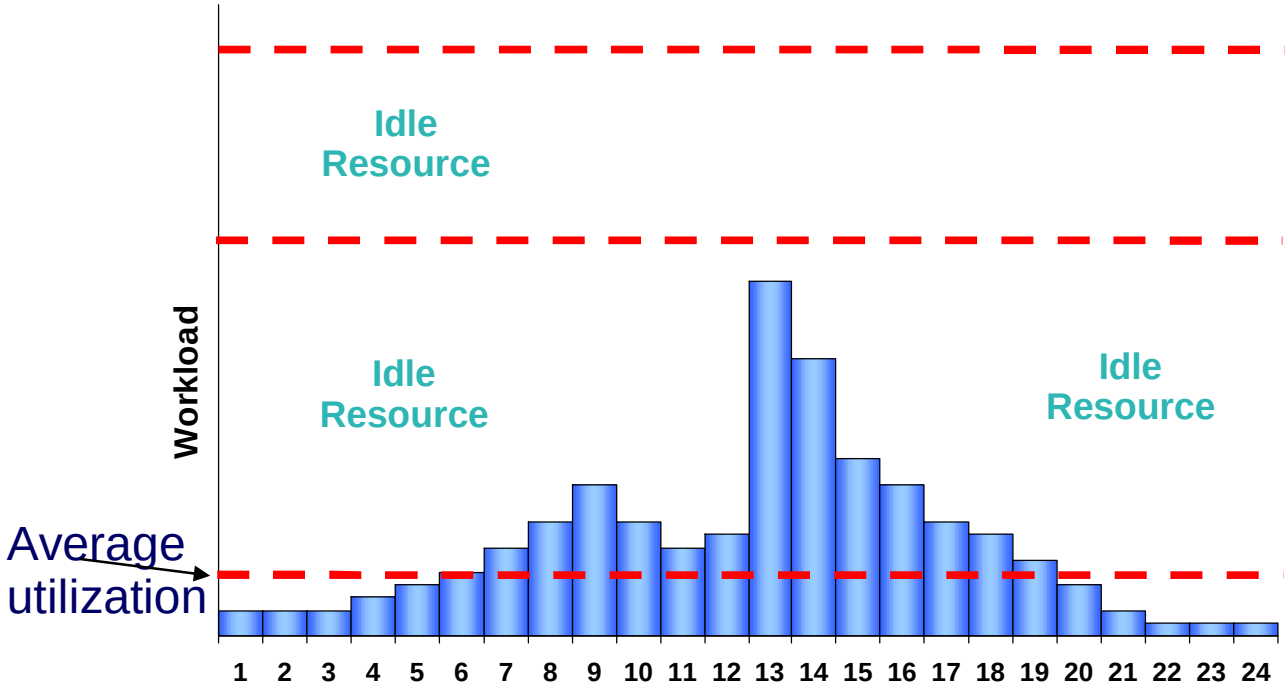

- **Typical utilization of Windows Servers 5 – 10%**
- **Typical utilization of UNIX Servers 10 – 20%**
- **Typical utilization of System z Servers 85 – 100%**

RPE2

(from Ideas International)

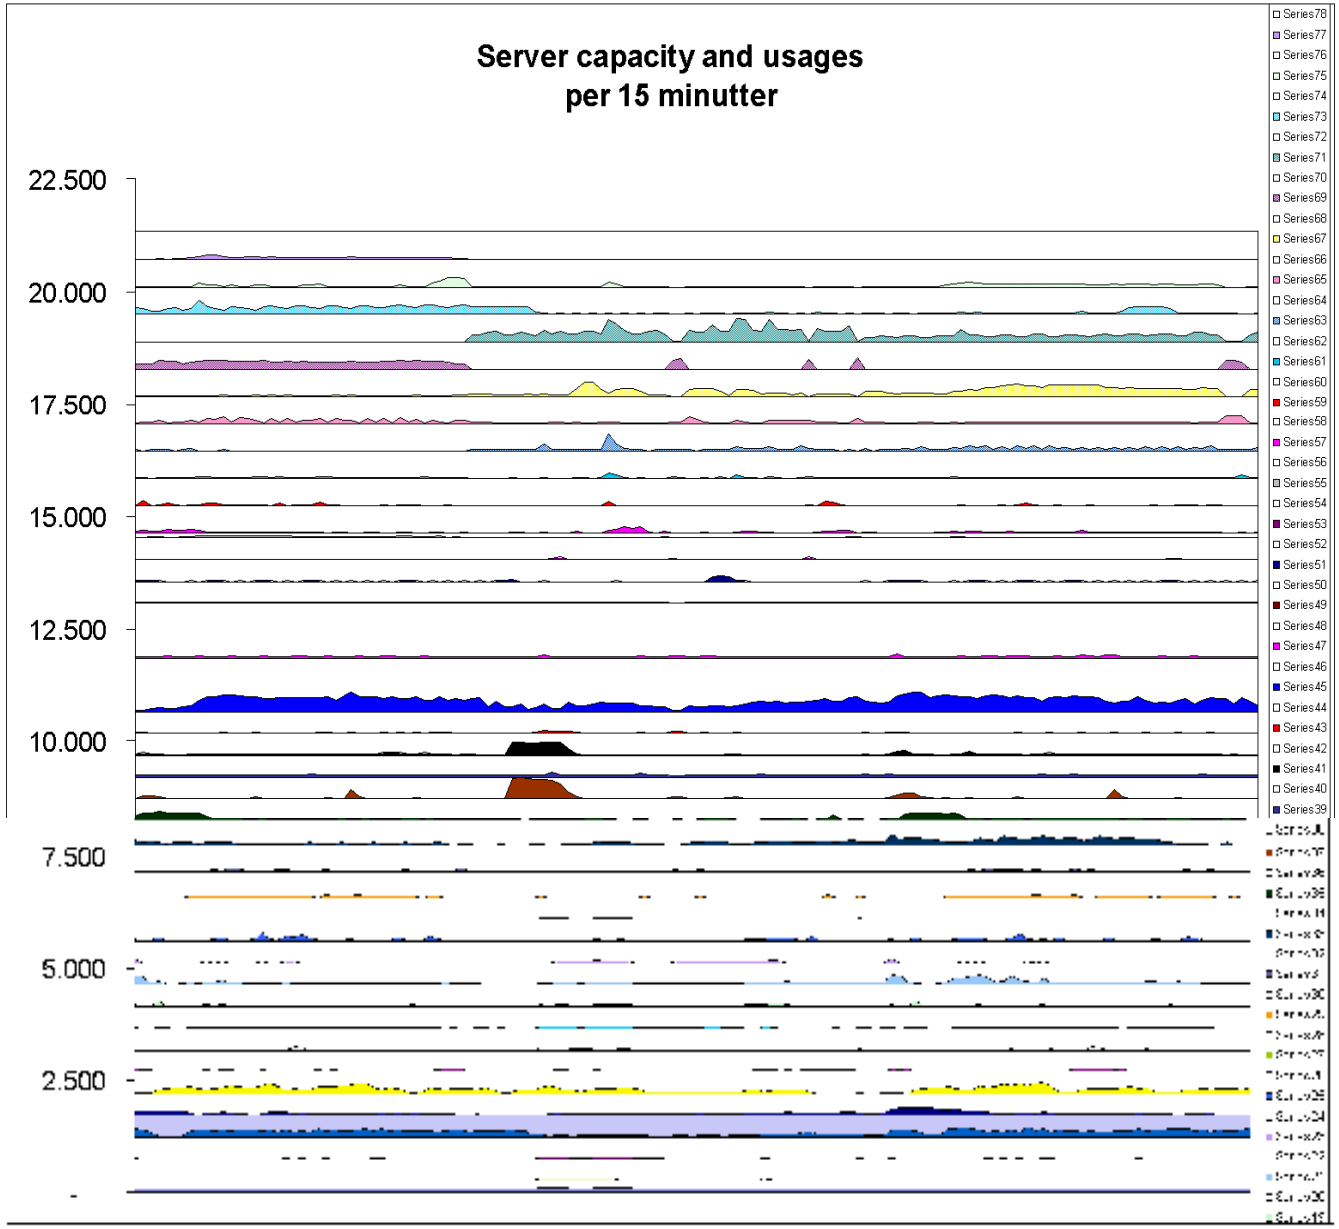

RPE2

(from Ideas International)

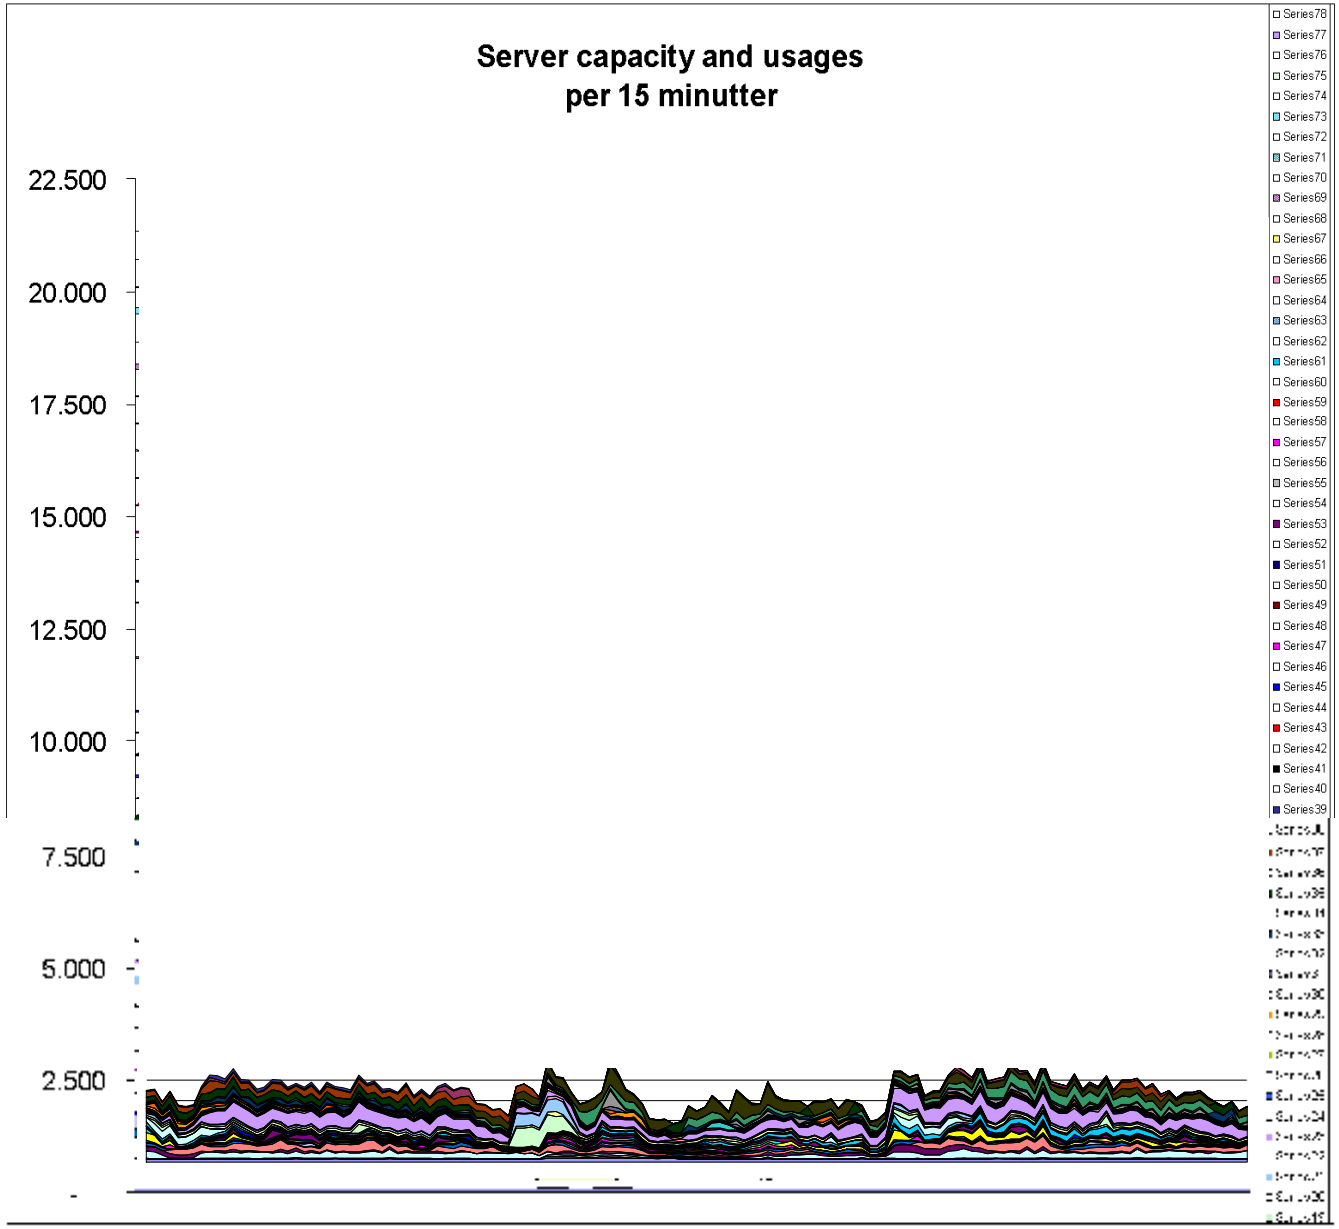

System z – Extreme **Virtualisation**

Build-in and Shared Everything Architecture

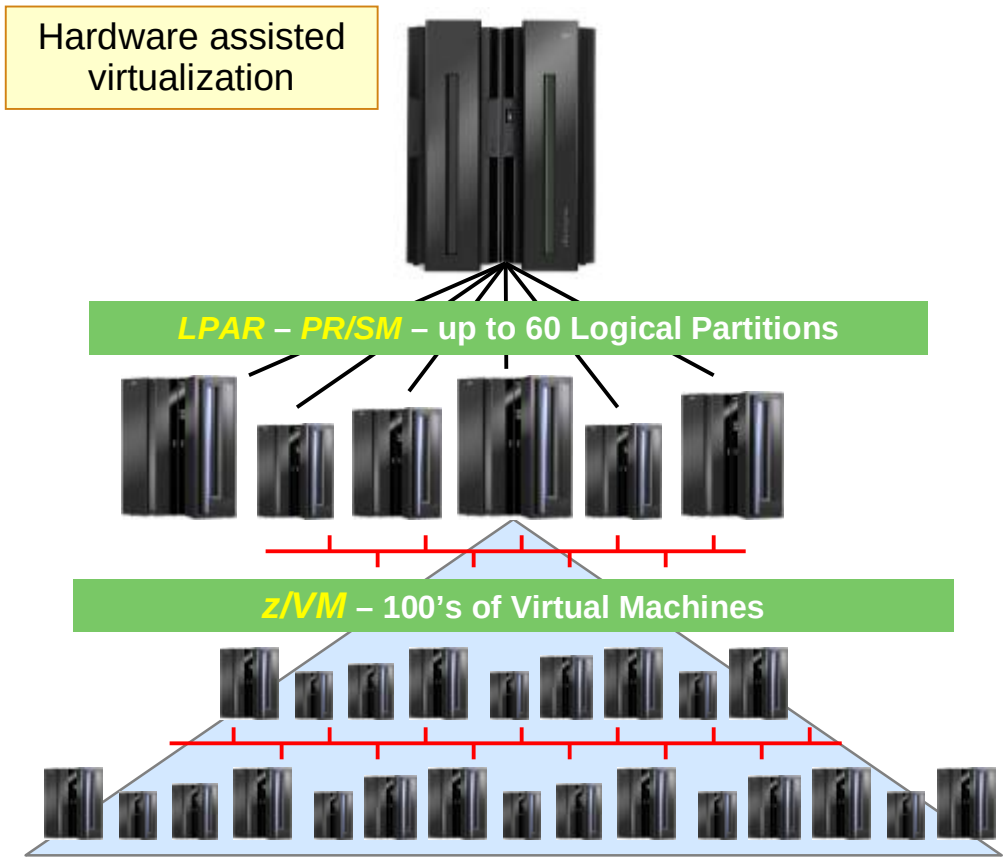

System z

- Provisioning of virtual servers in seconds
- High granularity of resource sharing (<1%)
- Upgrade of physical resources without taking the system down
- Scalability of up to 1000's of virtual servers
- More with less: more virtual servers per core, sharing of physical resources
- Extensive life-cycle management
- HW-supported isolation, highly secure (EAL5 or EAL4+ certified)

Distributed platforms

- Limited virtual server scalability per core
- Scaling requires additional physical servers
- Operational complexity increases with growth of virtual server images
- VMware, Xen, Hyper-V focus on x86, no HW management across multiple platforms

Scheidt & Bachmann: Introducing SAP on System z

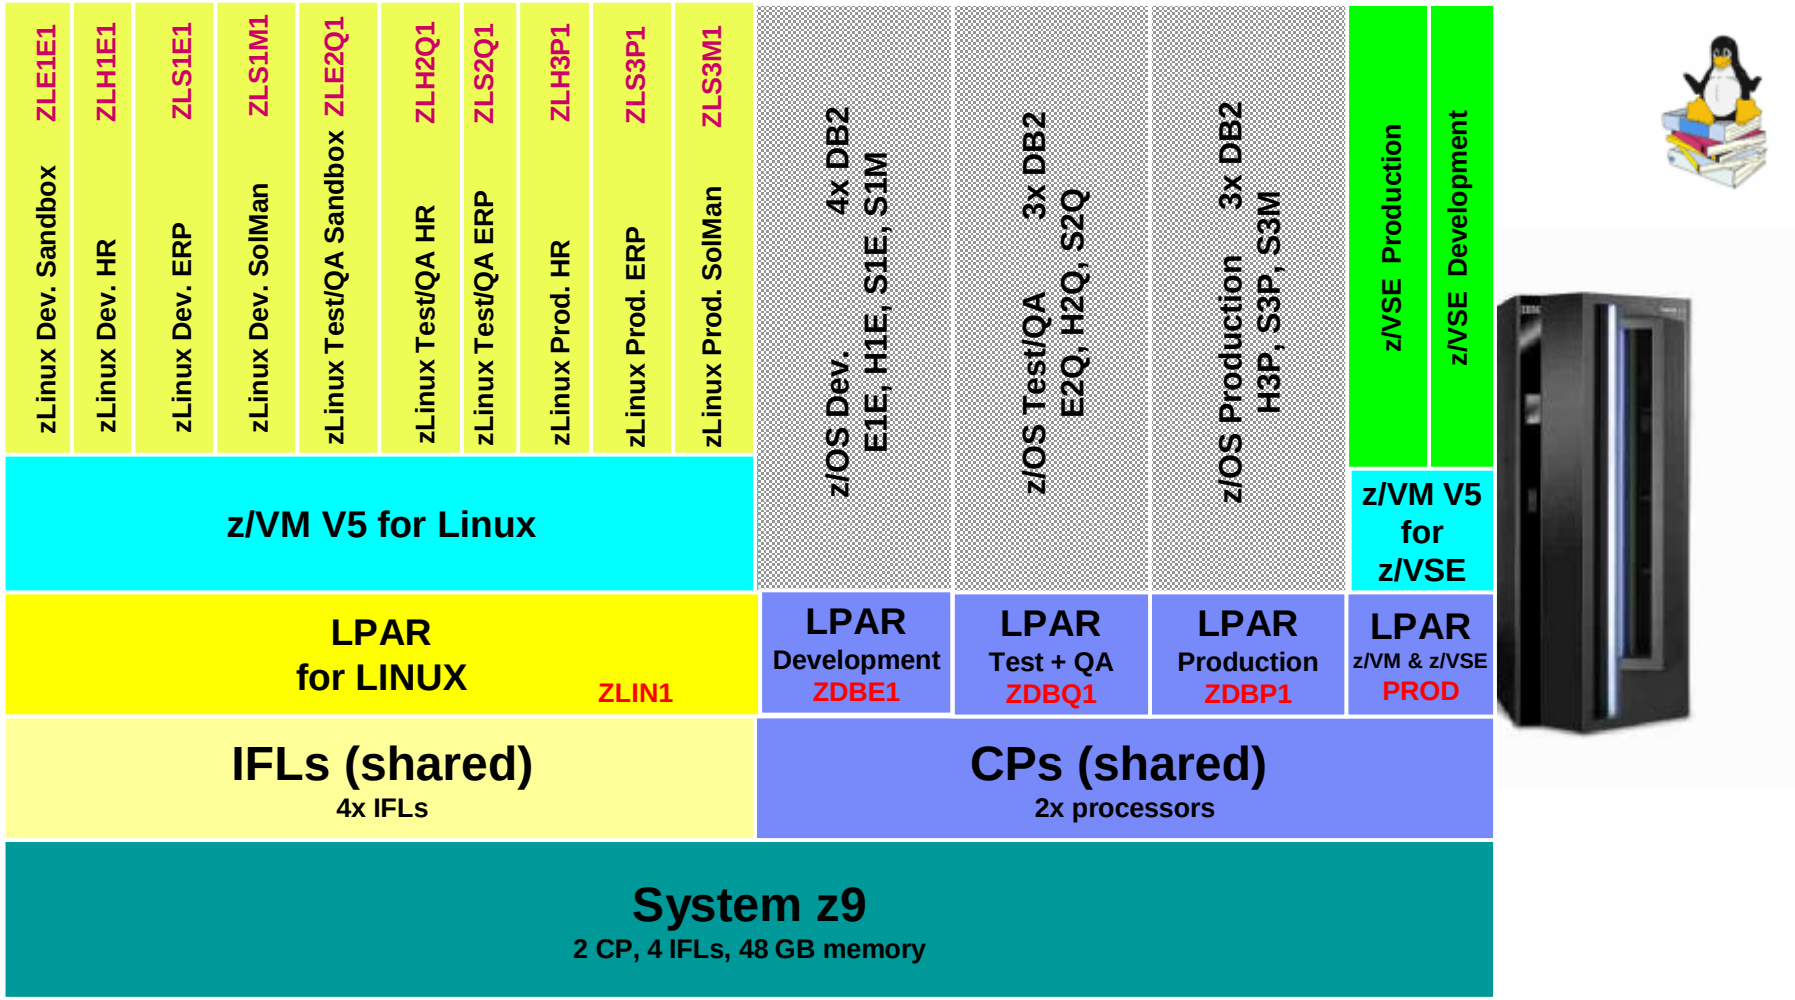

Extreme Virtualization with z/VM

z/VM can massively scale a virtual server environment with a mix of virtual and real resources for each virtual machine

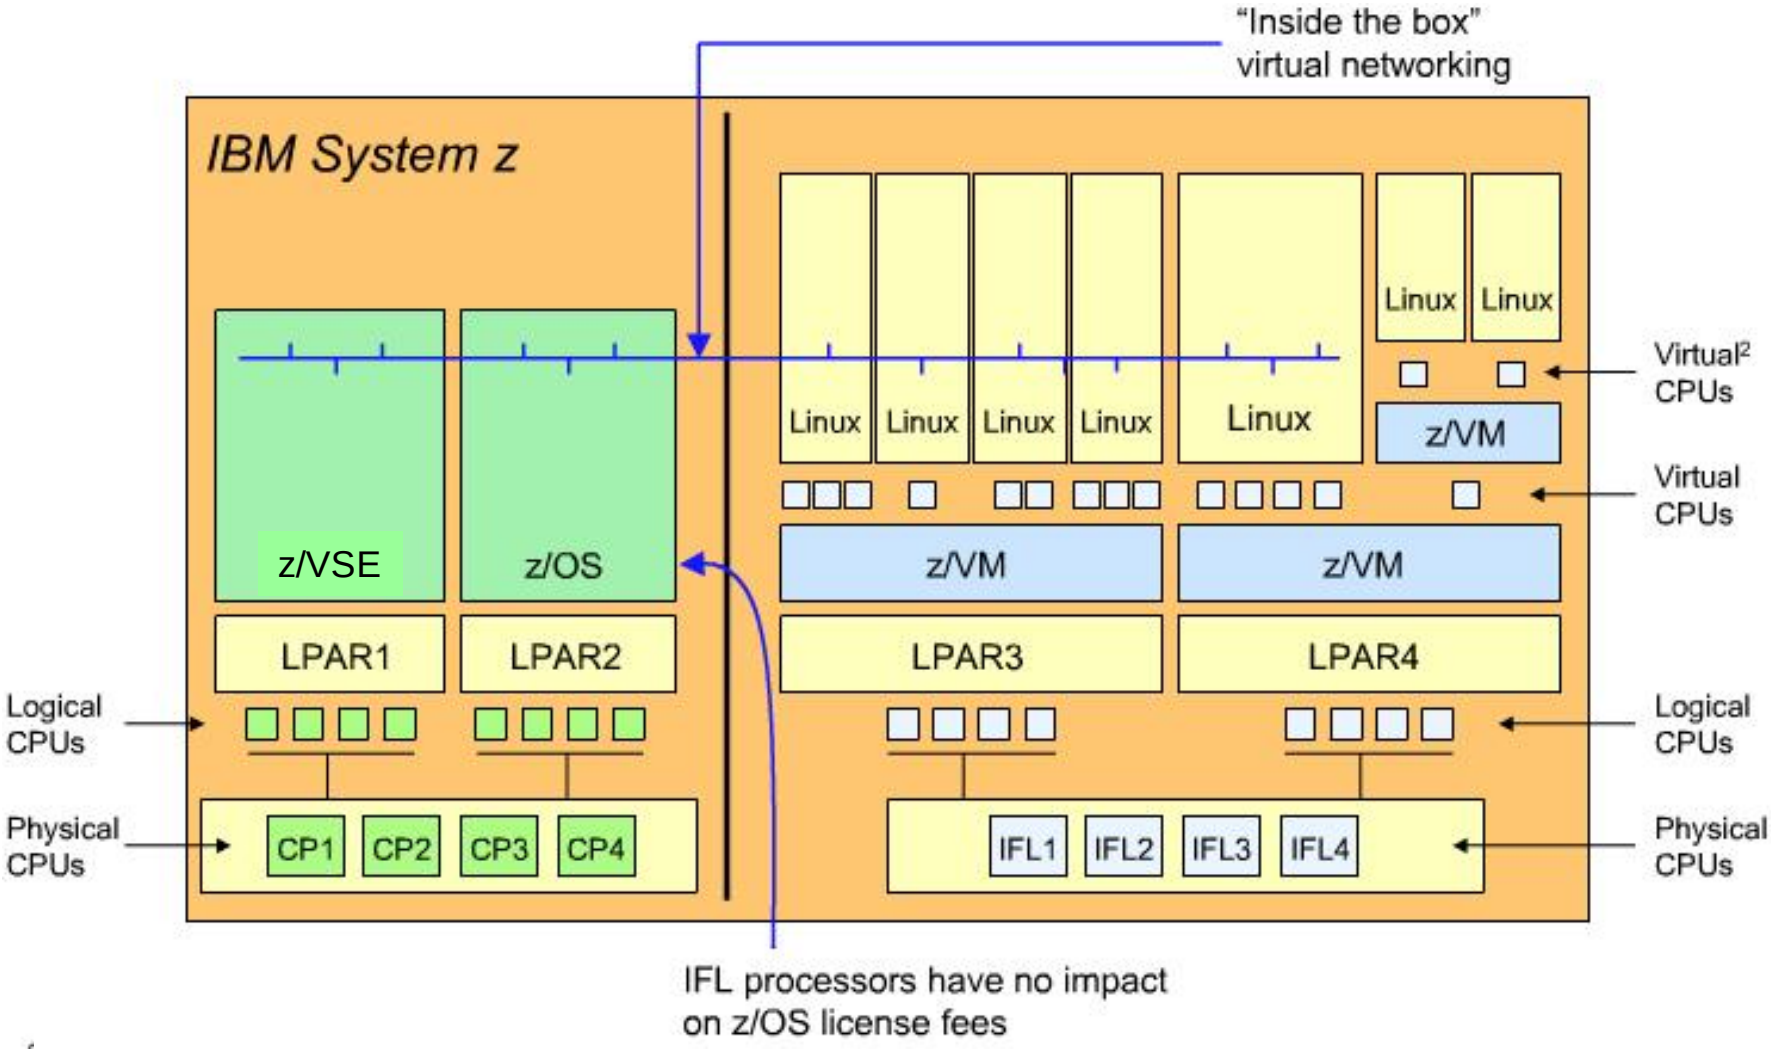

Extreme Virtualization with Linux on z/VM

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
  - Execute-in-place (xip2) file system
  - 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

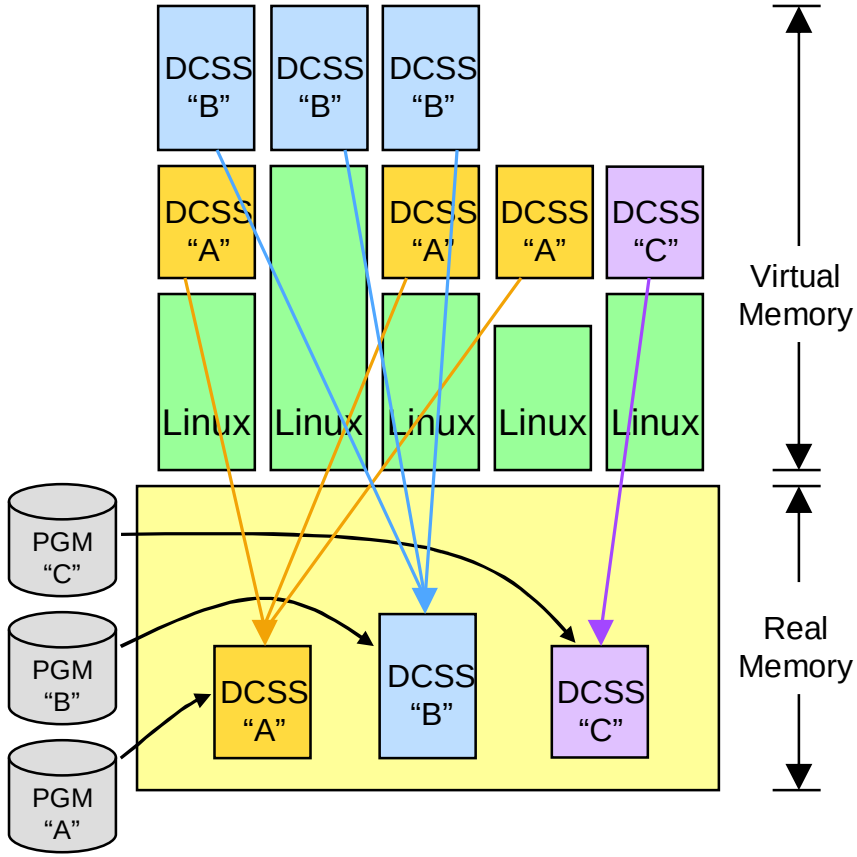

Learn more:
"Using DCSS/XIP with Oracle 10g on Linux for System z"
www.redbooks.ibm.com/redpieces/abstracts/sg247285.html

Extreme Virtualization with Linux on z/VM

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

Lab tests have shown up to 50% more throughput using CMM with z/VM 5.3

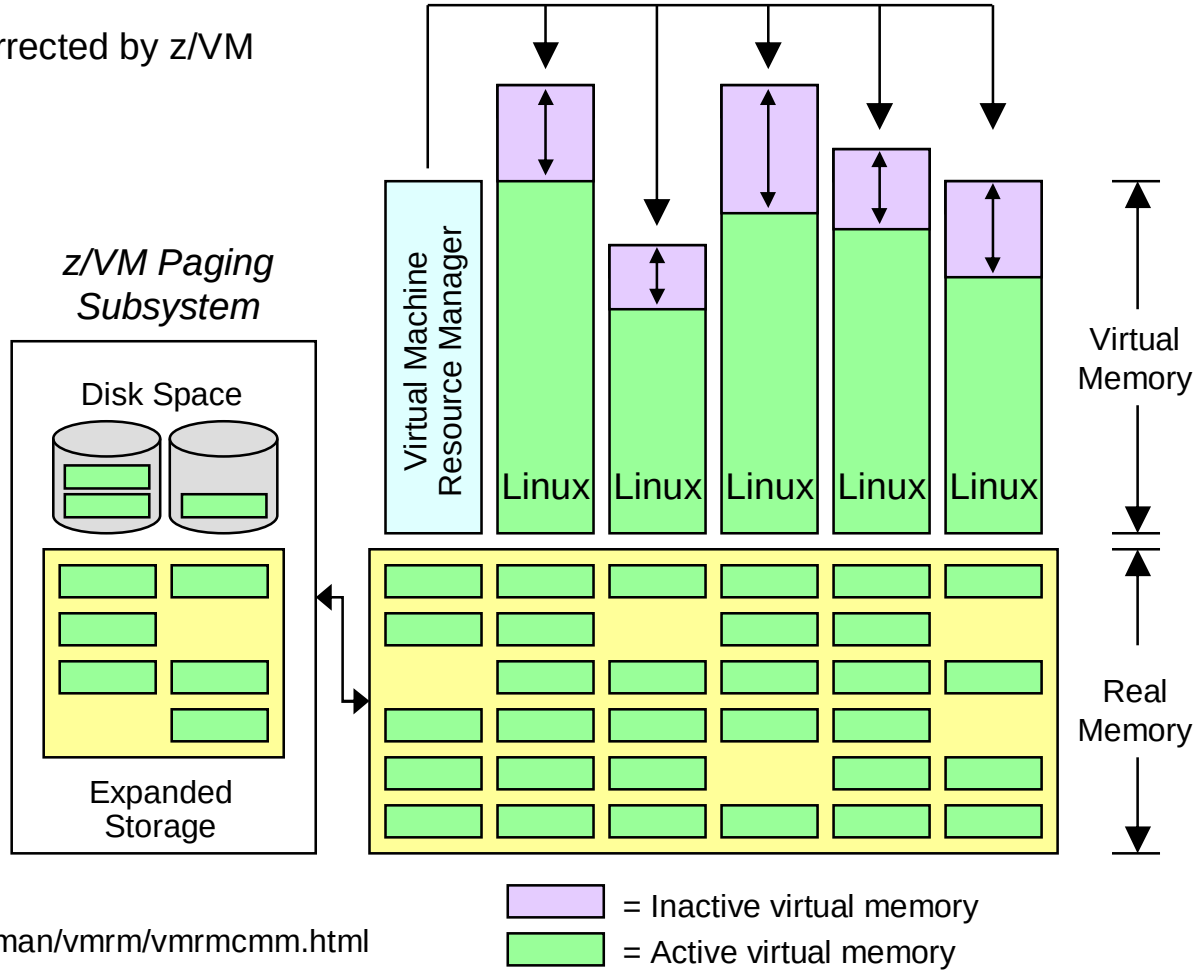

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

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

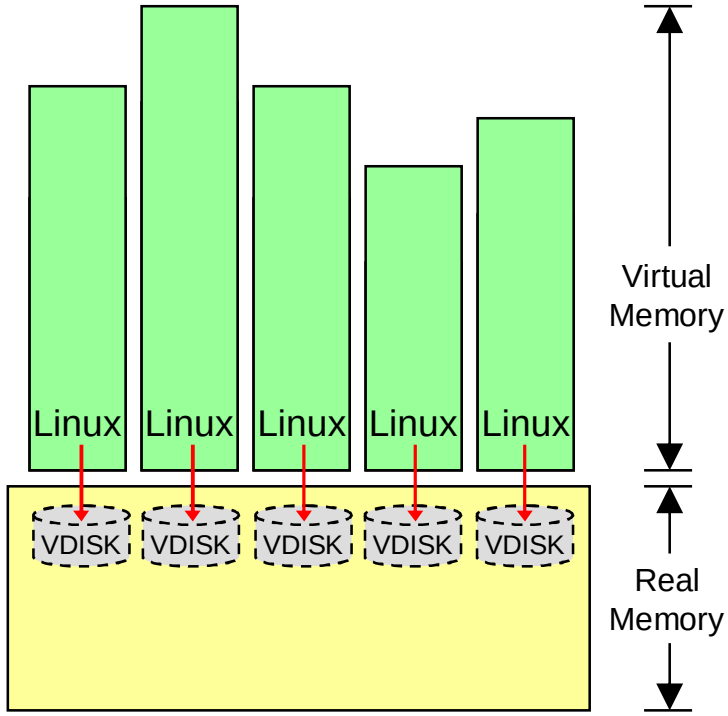

Bank of New Zealand

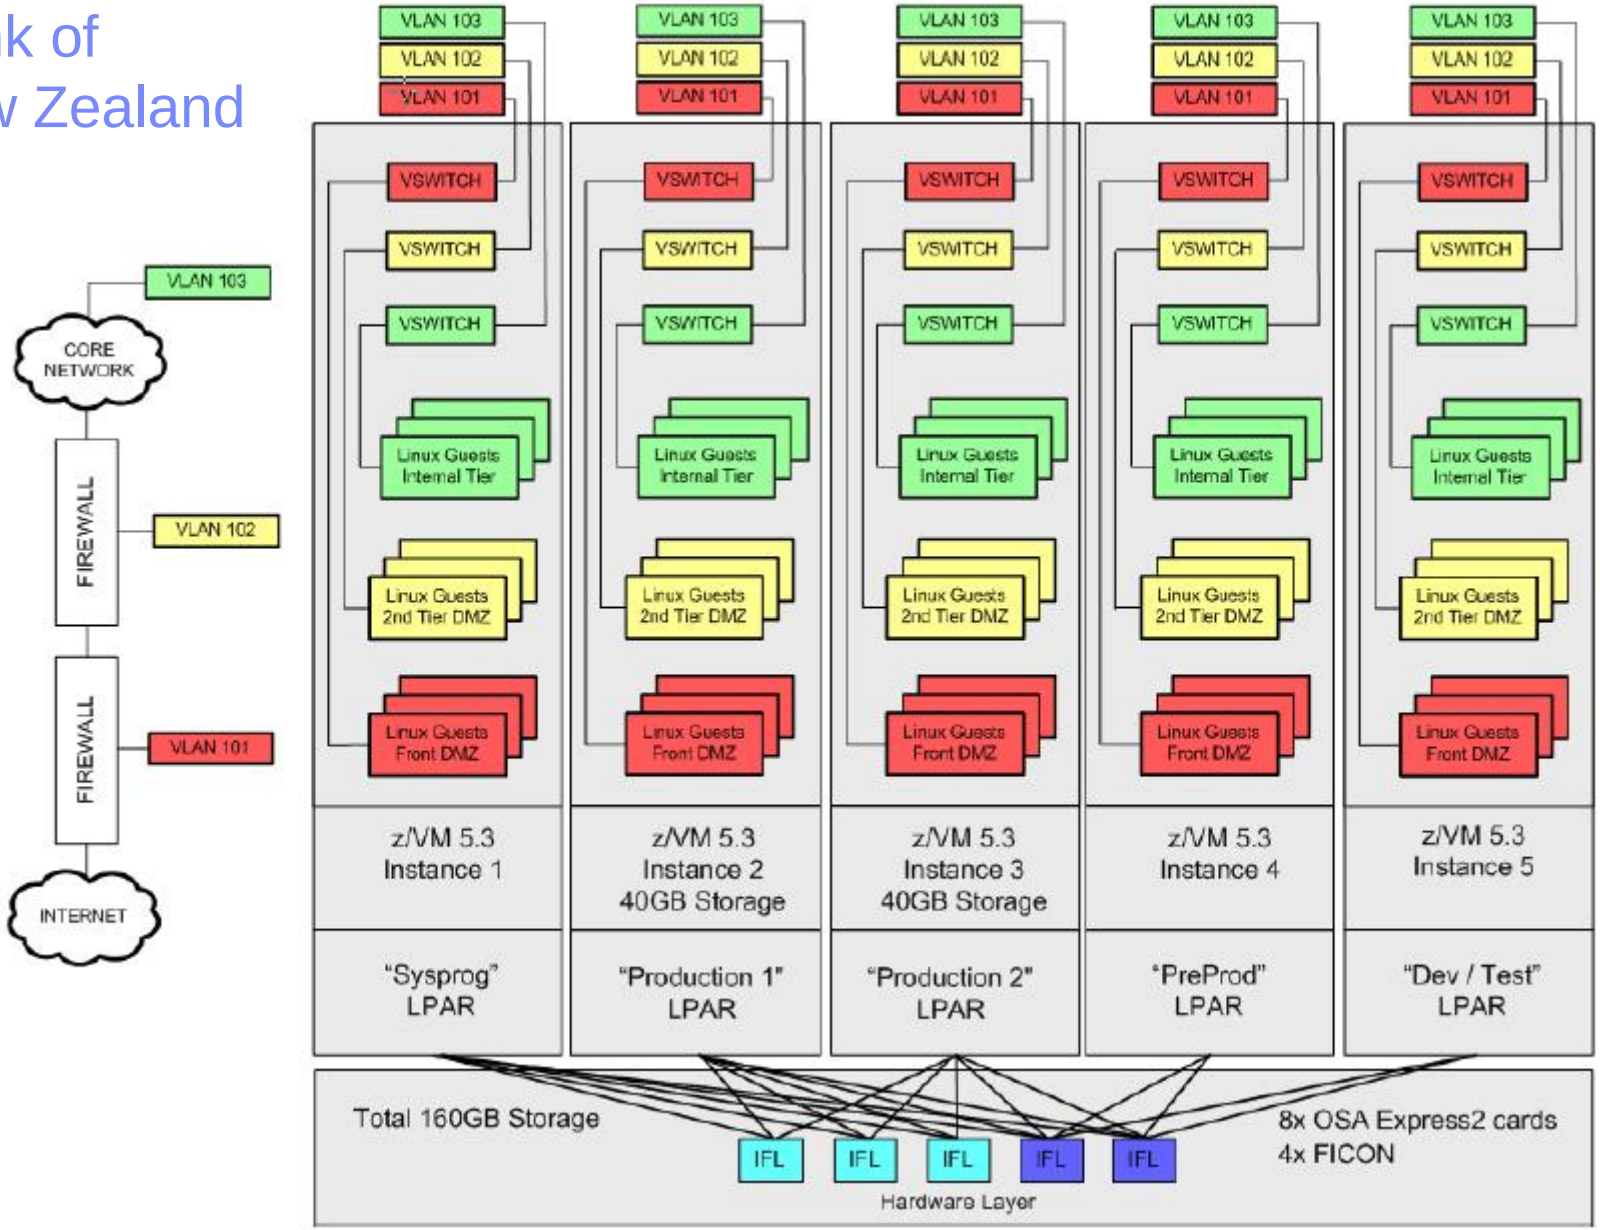

Throughput comparison

IBM x3950 (8 cores @ 3.5 Ghz) compared to z10 EC

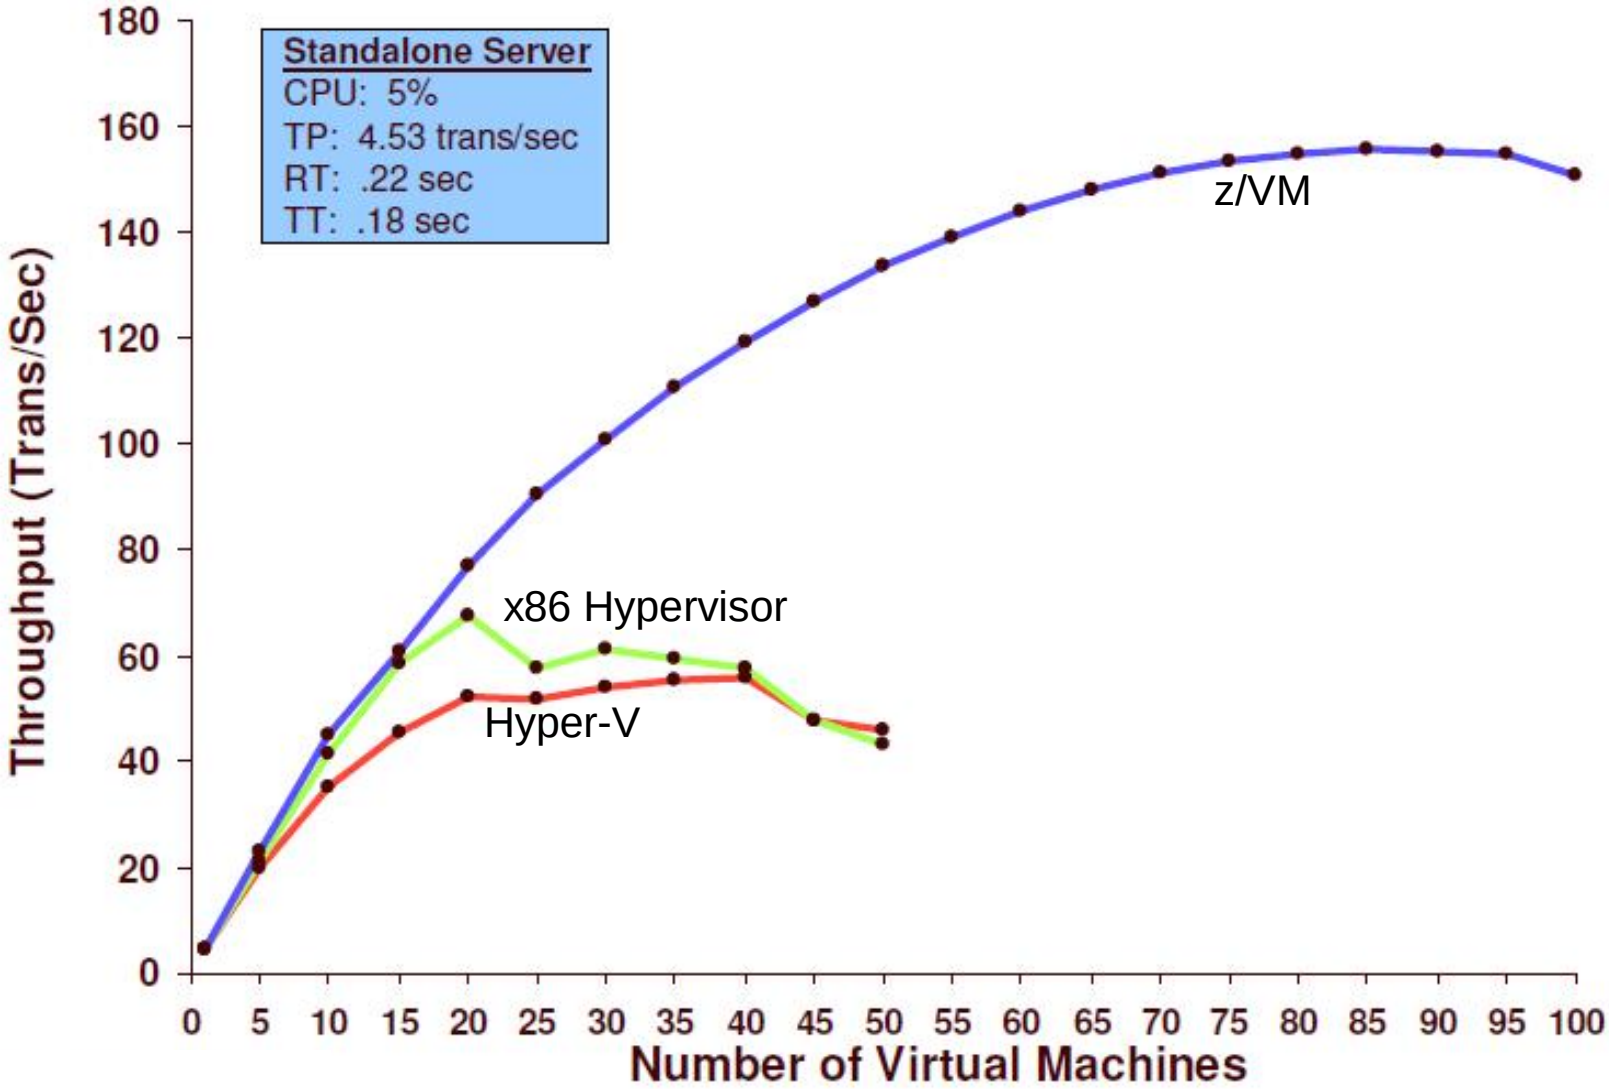

18

Response time comparison

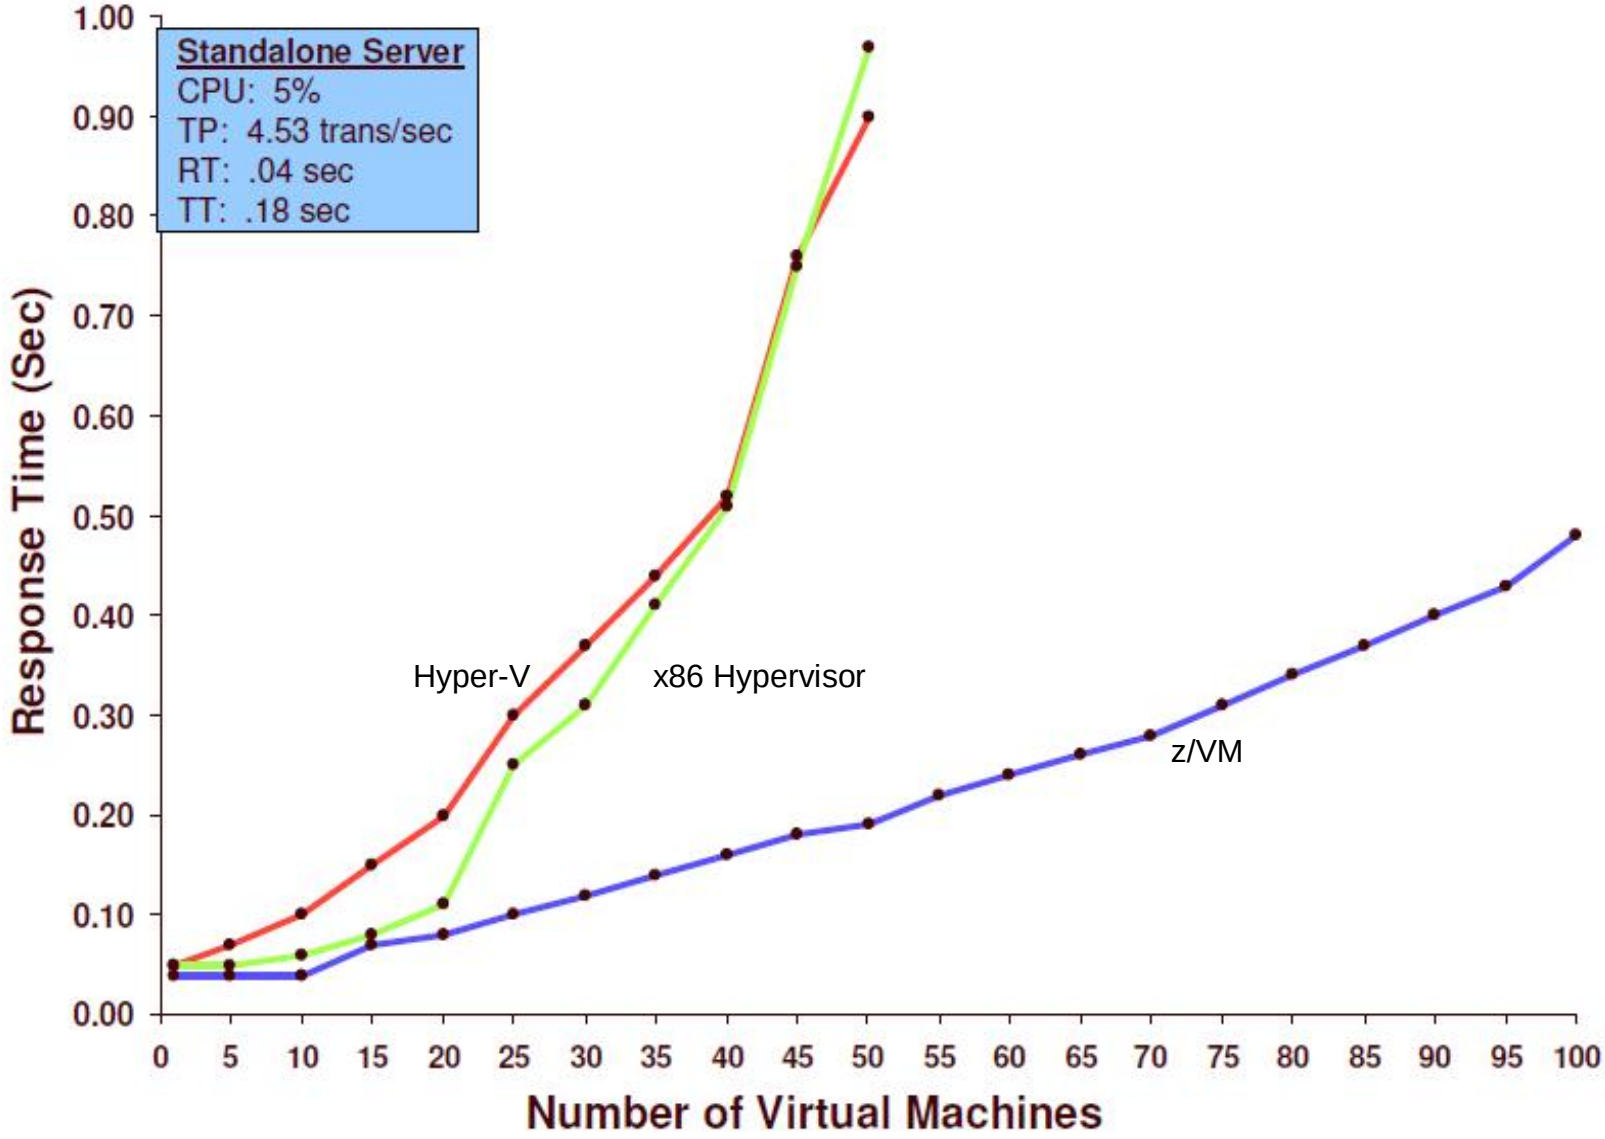

The effect of memory constraints on response time

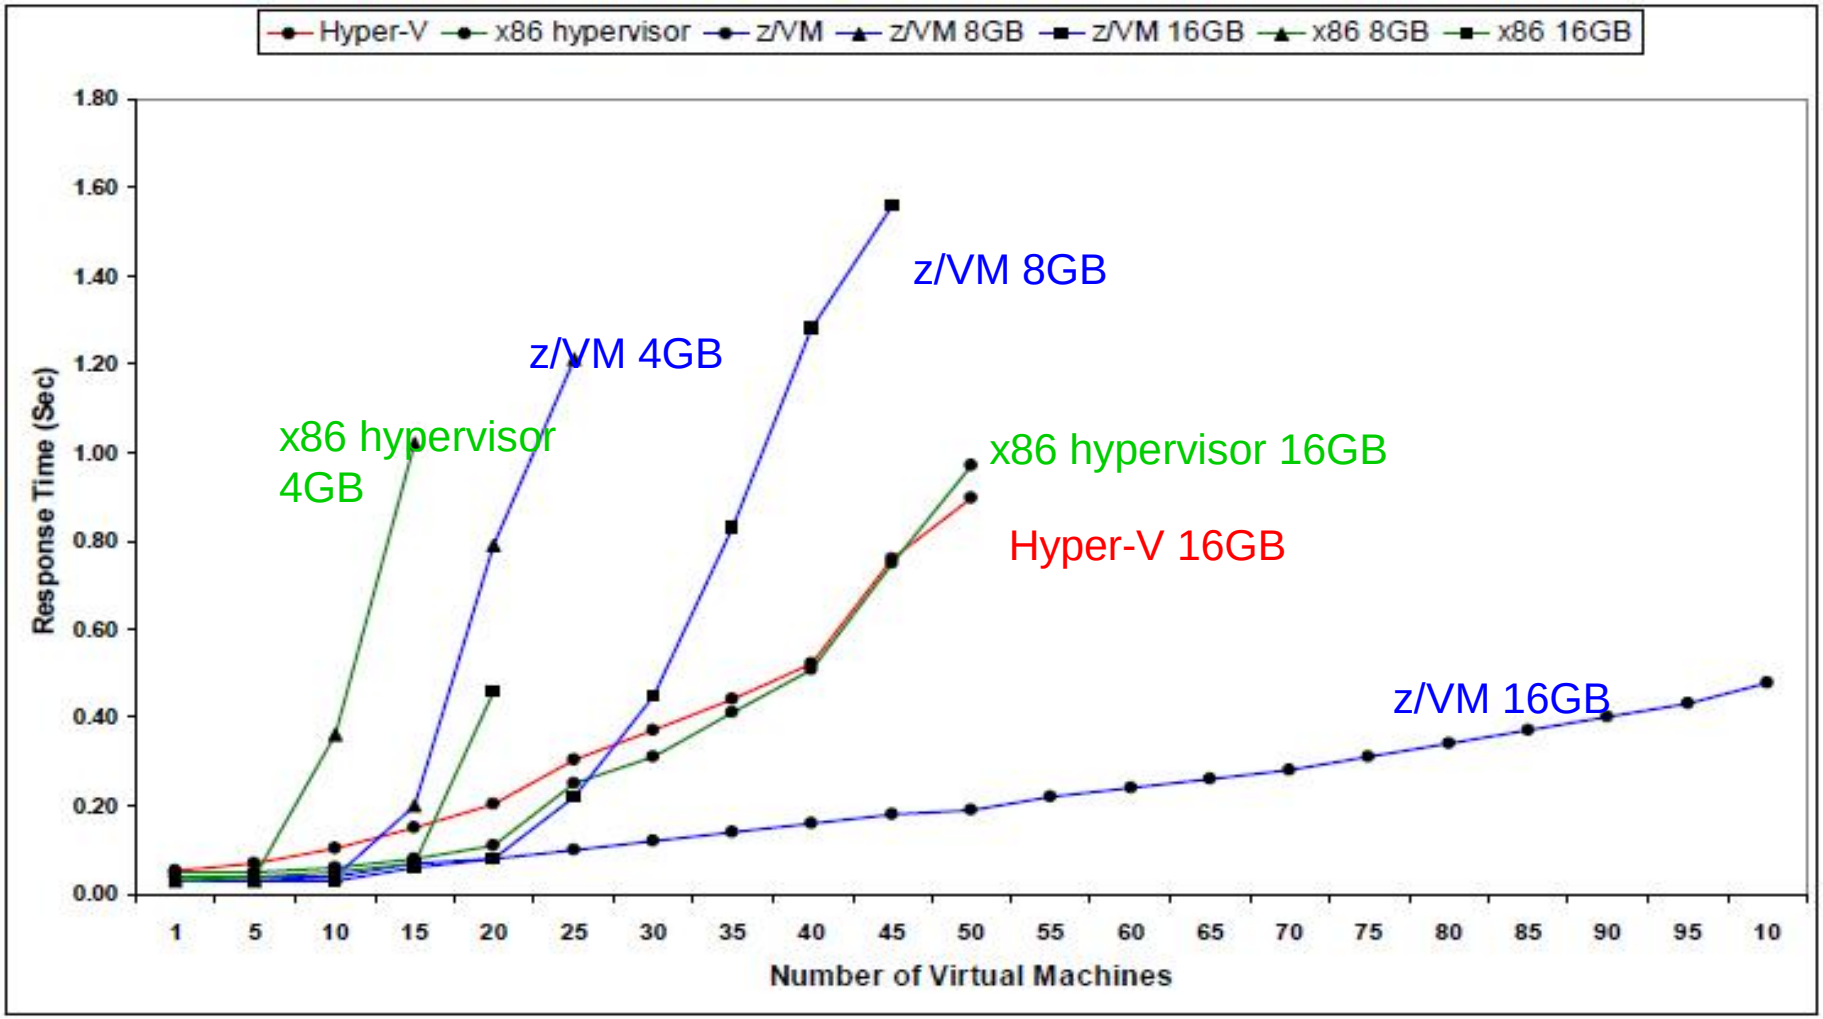

Variability In Workload Demand

- Variations in workload demand measured by time interval
- The mean **M** is the average utilization over all the time intervals
- **Sigma** (the Standard Deviation) is a measure of how widely the intervals differ from the mean (greater sigma means more variation in workload demand)
  - 2 standard deviations from the mean: approximately 95% of all values are less than this

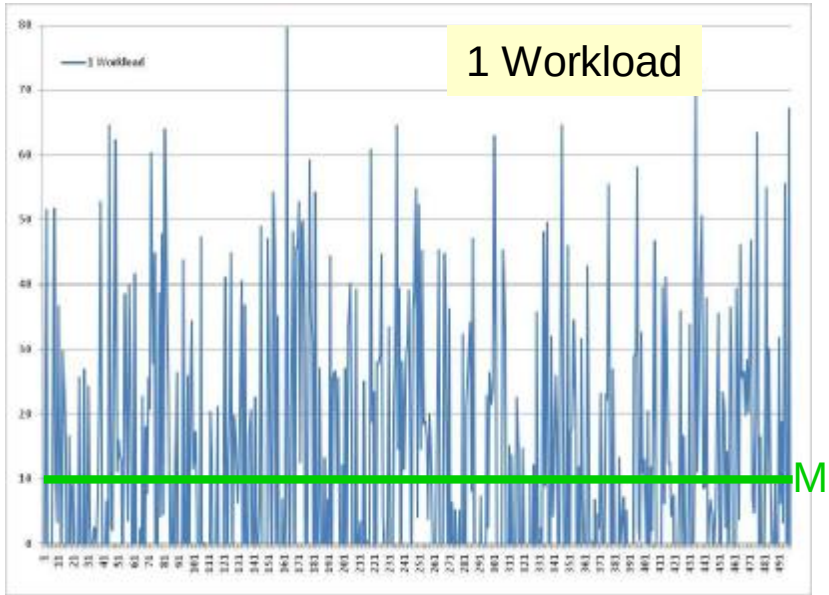

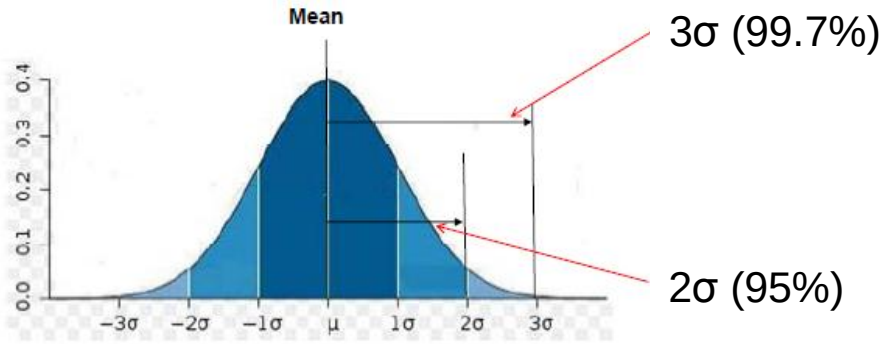

Service Level Agreements

- Service Level Agreement stipulates that only a small fraction of intervals may exceed the capacity of the machine
- For example*
  - Only 5% of intervals will exceed M + 2*Sigma

Only 5% of intervals will exceed M + 2*Sigma

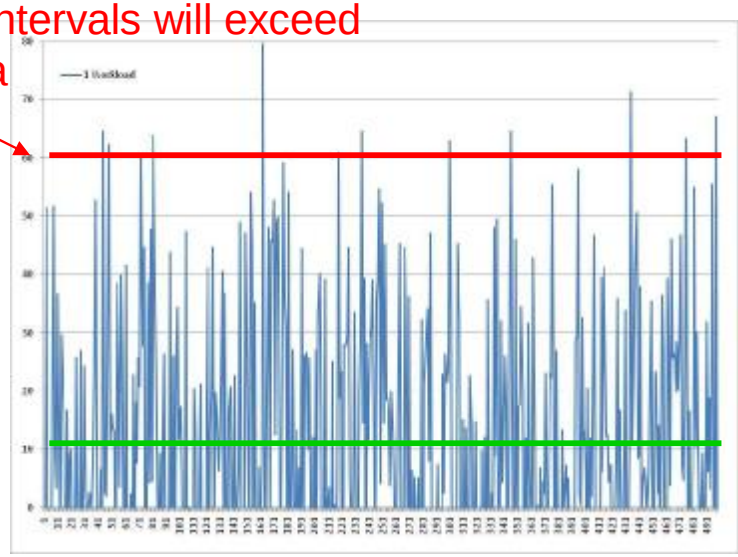

When Sigma = 2.5*Mean, a machine capacity of 6*M is required

The average utilization will be 17%!

* Normal probability distribution

When we consolidate…

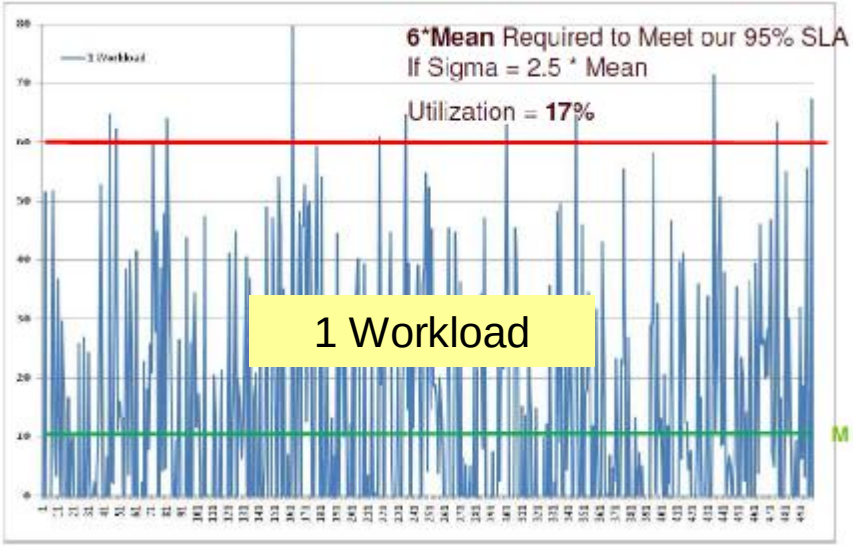

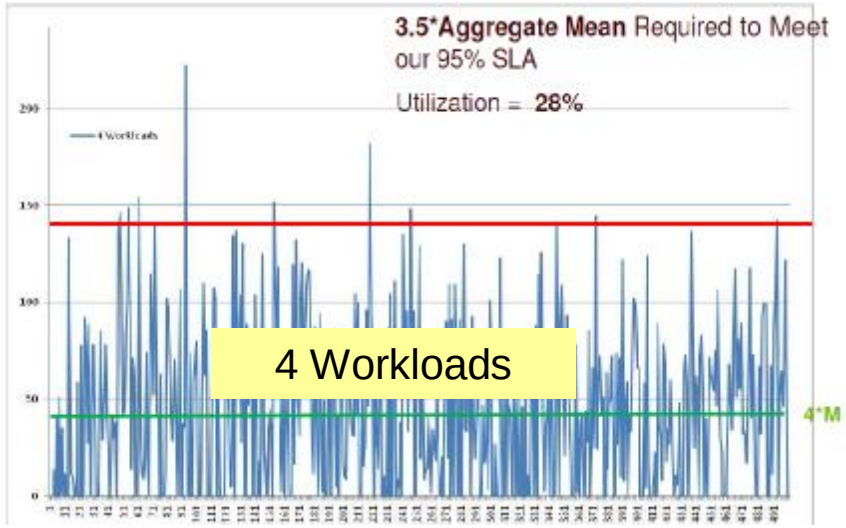

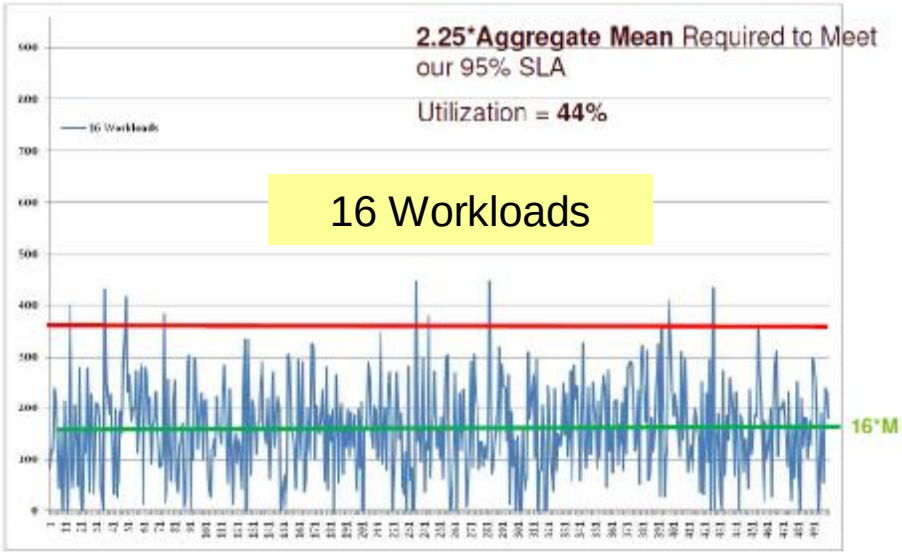

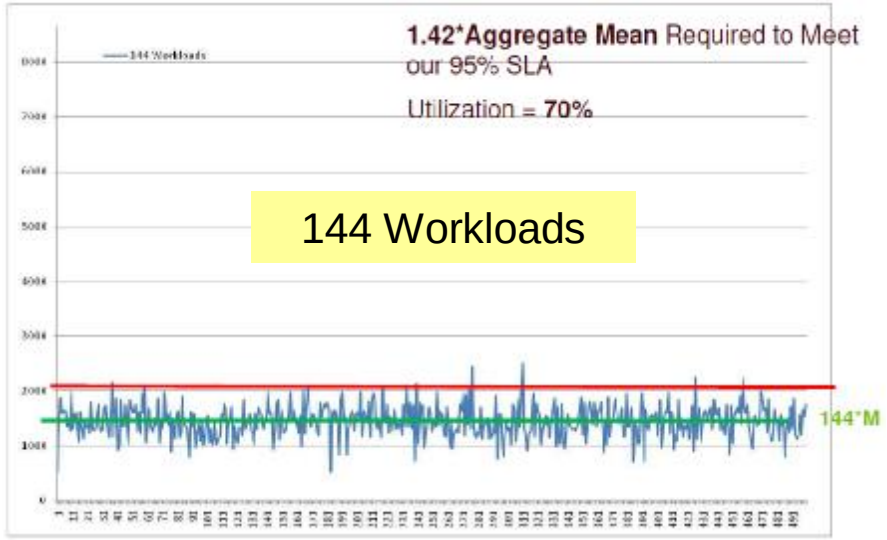

Apply utilization SLA to derive consolidation ratio for x86 Hypervisor

x86 Hypervisor results

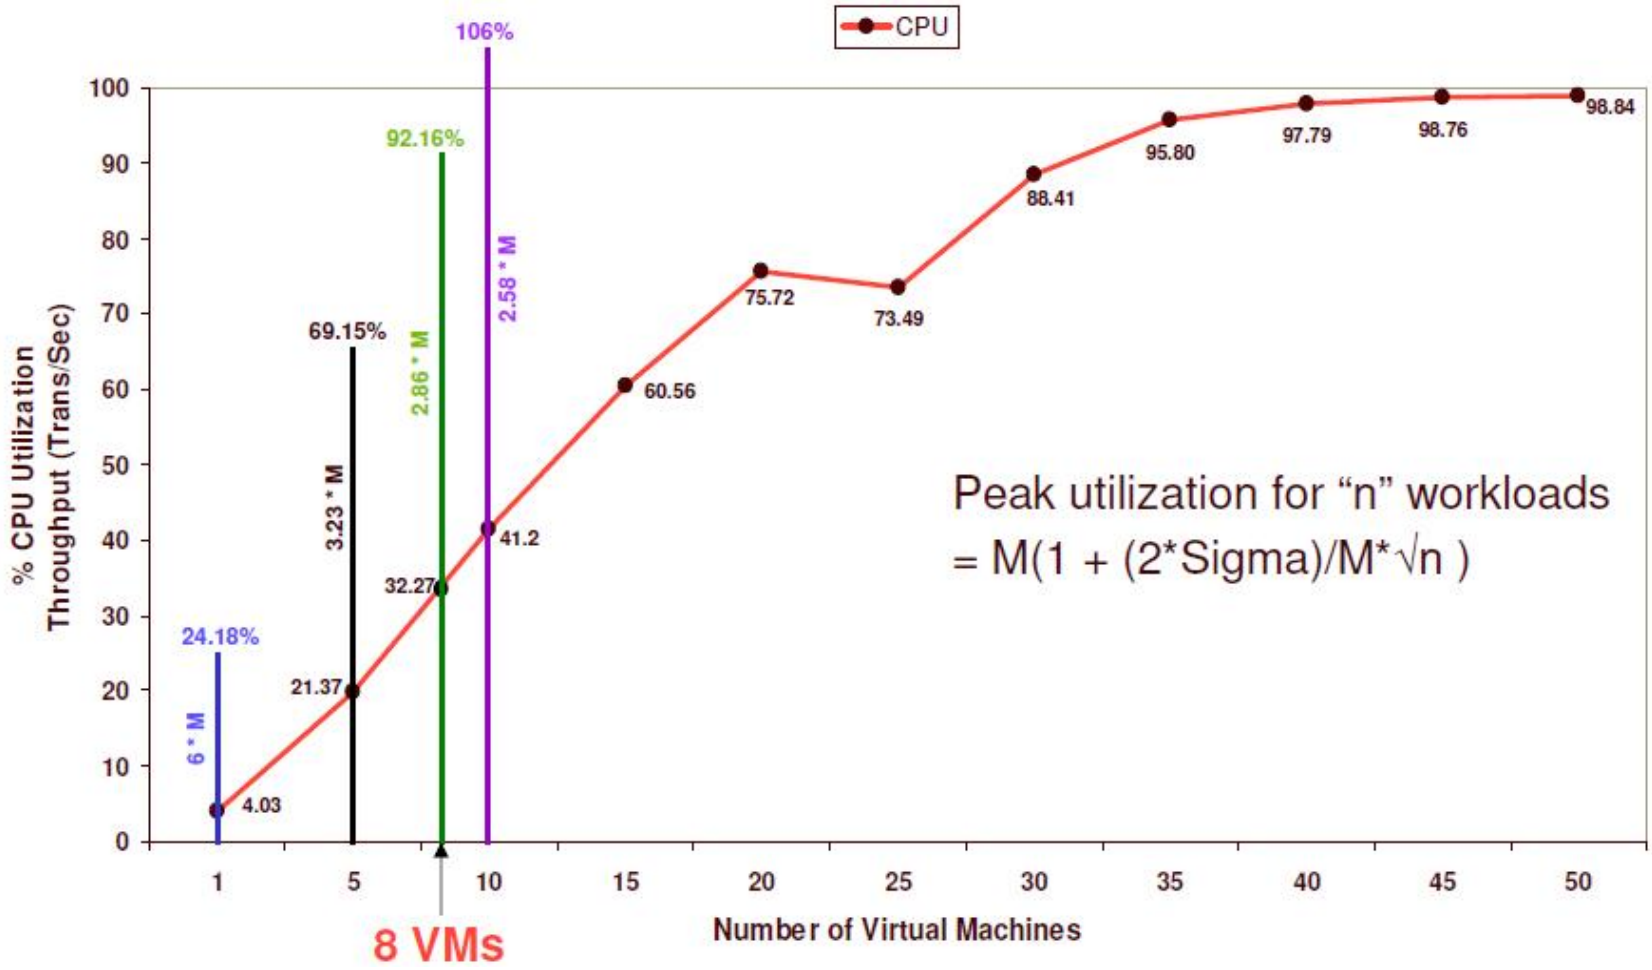

Apply utilization SLA to derive consolidation ratio for z/VM

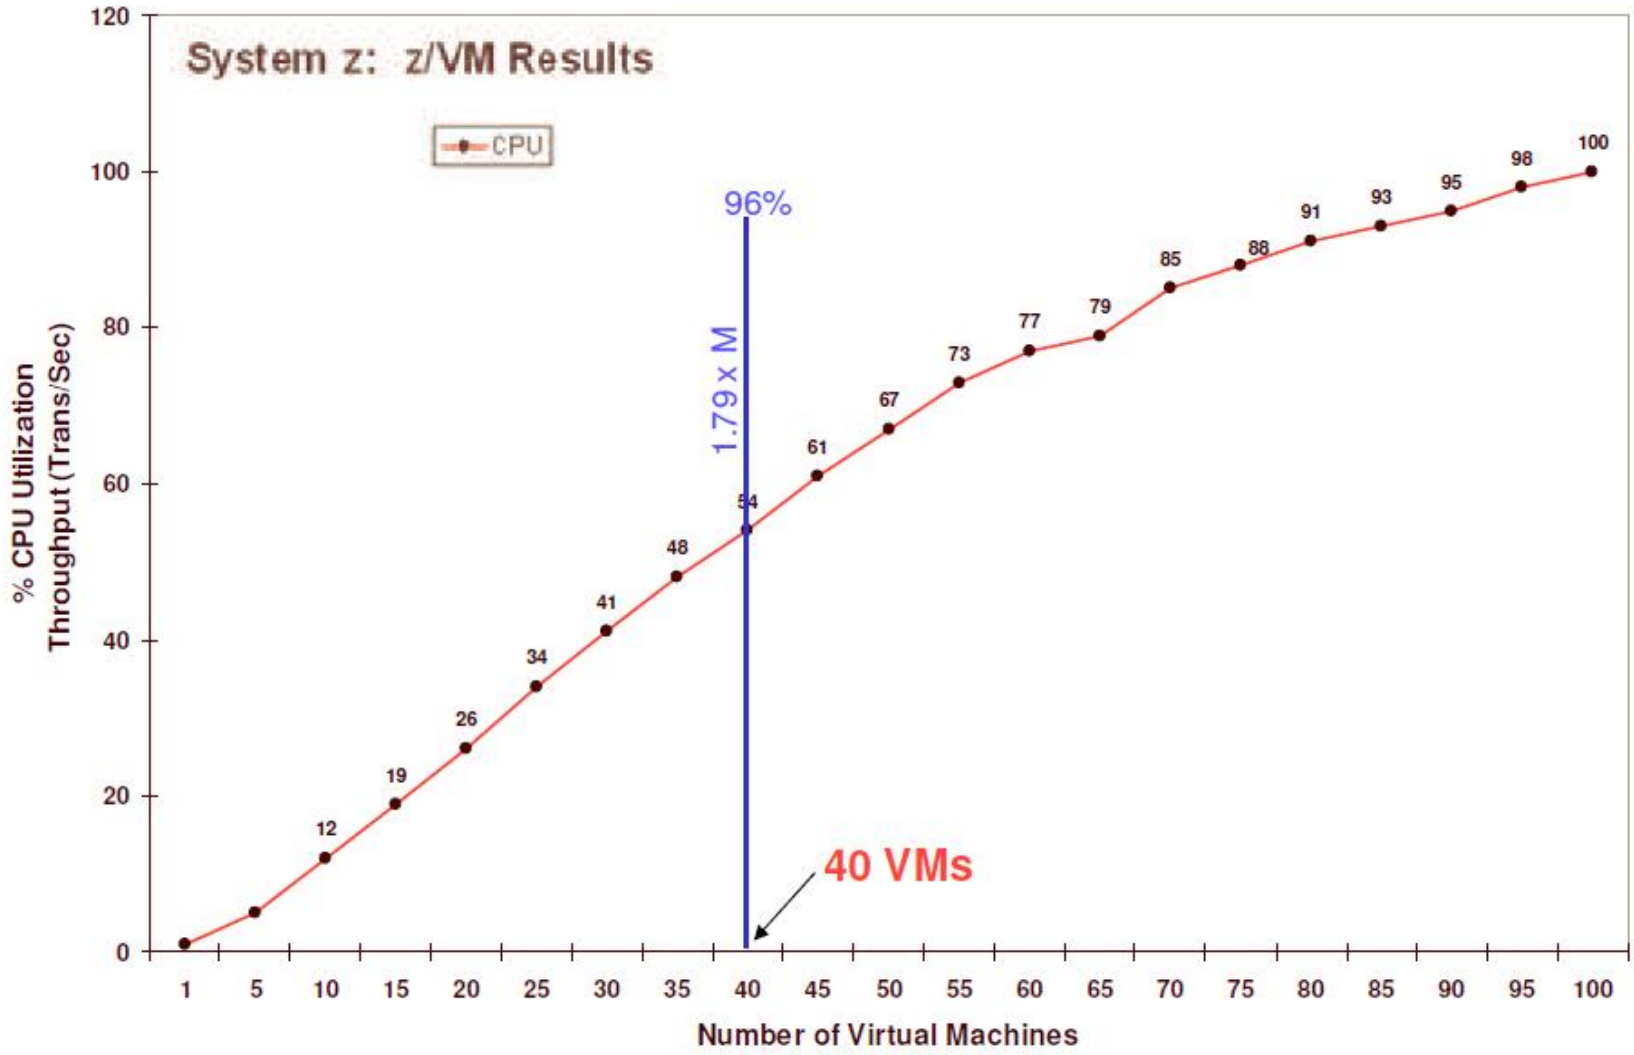

High Core-to-Core Ratios for Consolidations from Distributed IT-Environments to Linux on System z

Real customer examples with real workloads!

Industry	Distributed Cores	IBM System z10™ Cores	Core-to-Core Ratio*
Public	292	5	58 to 1
Banking	111	4	27 to 1
Finance	442	16	27 to 1
Banking	131	5	26 to 1
Insurance	350	15	23 to 1
Insurance	500+	22	22 to 1
Banking	63	3	21 to 1
Finance	854	53	16 to 1
Health care	144	14	10 to 1
Transportation	84	9	9 to 1
Insurance	7	1	7 to 1

* Client results will vary based on each specific customer environment including types of workloads, utilization levels, target consolidation hardware, and other implementation requirements.

IBM Consolidation: 3900 servers to Linux on System z

Similar distributed workload vs. System z Linux results in potential 60-75% Gross Costs Savings

Operating Cost: Distributed vs. Mainframe

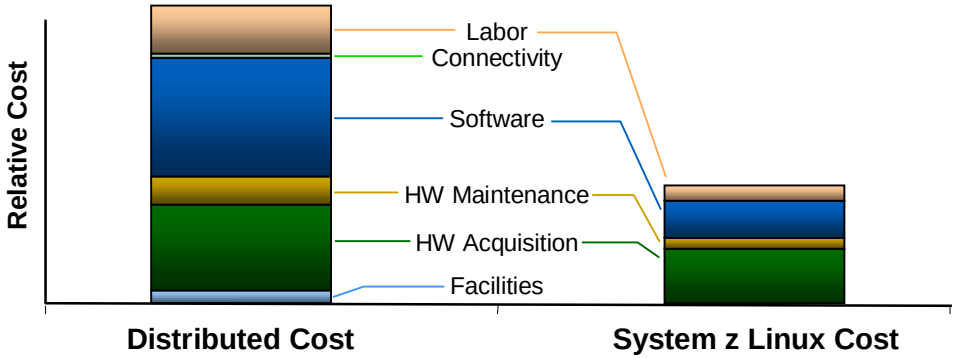

Potential Savings: Categories as a % of Gross Savings

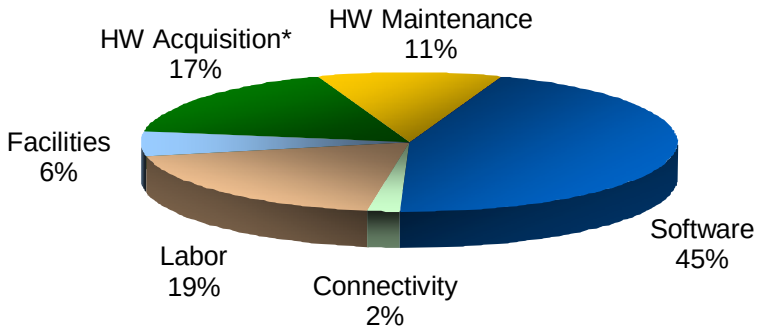

* HW Acquisition compares server/disk refresh of distributed environment to the cost of acquiring new mainframes/storage

Dramatic Simplification

Unit	Distributed	System z Linux	% Reduction
Software Licenses	**26,700**	**1,800**	**93%**
Ports	**31,300**	**960**	**97%**
Cables	**19,500**	**700**	**96%**
Physical Network Connections	**15,700**	**7,000**	**55%**

Results will vary based on several factors including # of servers and work load types

27

TCO Comparison of Real Customer Cases

Scenarios	Cost of Distributed	vs.	ELS	Distributed Cost Ratio	Migration Cost
Green Field Cases					
- Banking Benchmark	$43.3M	vs.	$18.2M	2.4x	No migration
Offloading Cases					
- Asian financial company	$119.0M	vs.	$53.0M	2.2x	6.0M
- Asian insurance company	$25.1M	vs.	$16.3M	1.5x	2.1M
- NA financial services	$58.9M	vs.	$34.0M	1.4x	5.0M
- European financial	$17.9M	vs.	$4.9M	3.7x	4.7M
- US County government	$8.1M	vs.	$4.7M	1.7x	2.9M
Offload Studies					
- European agency	€386.0M	vs.	€204.0M	1.9x	6.3M
- Restaurant chain	$56.3M	vs.	$23.3M	2.4x	10.0M
- Asian healthcare	$15.1M	vs.	$8.9M	1.7x	4.8M
- Asian bank	$31.6M	vs.	$23.5M	1.3x	6.0M
- US utility	$13.4M	vs.	$6.2M	2.2x	1.9M
- US manufacturer	$64.0M	vs.	$43.3M	1.5x	12.2M

IBM System z Virtualization Support

Saving Money and Reducing Complexity
Helping You "Do More with Less" with Linux

- Consolidate more x86 cores per CPU
- Spend less on software license fees
- Manage more virtual servers with fewer people
- Deploy new servers and applications faster
- Absorb workload spikes more easily
- Spend less on disaster recovery
- Occupy less floor space
- Save on energy

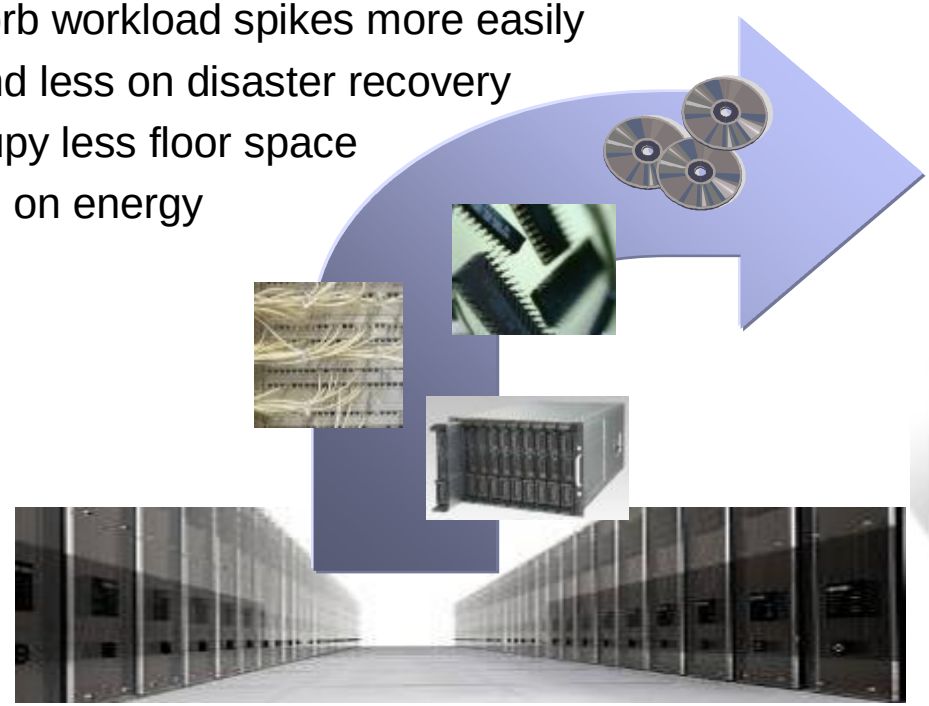

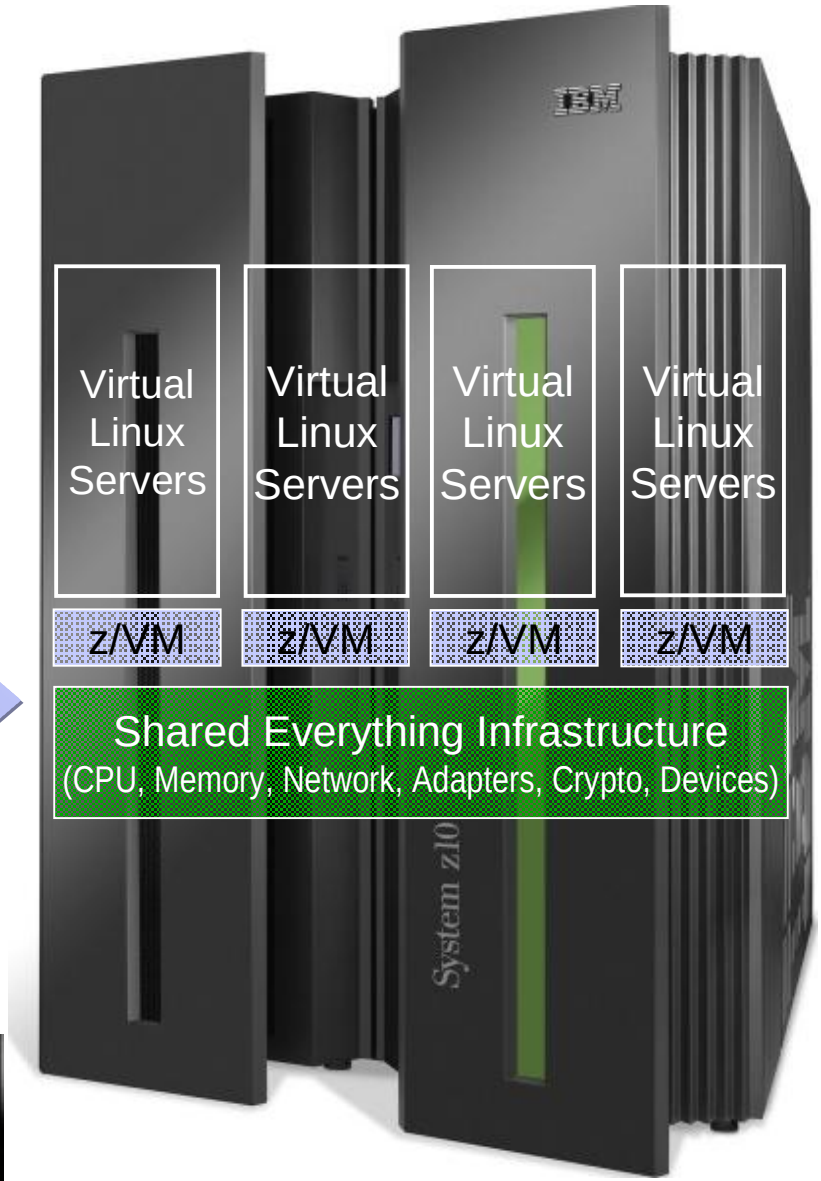

Questions?

Siegfried Langer
Business Development Manager
z/VSE & Linux on System z

IBM Deutschland Research & Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany

Phone: +49 7031 - 16 4228

Siegfried.Langer@de.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	FICON*	Parallel Sysplex*	System z10
BladeCenter*	GDPS*	POWER*	WebSphere*
CICS*	IMS	PR/SM	z/OS*
Cognos*	IBM*	System z*	z/VM*
DataPower*	IBM (logo)*	System z9*	z/VSE
DB2*			

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* 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.