IBM eServer Platform Selection

Rory Canellis
Mainframe Evangelist
IBM Software Group
rory@us.ibm.com

Agenda

System z and Best of Breed UNIX

- Differences and Similarities
  - History and Heritage and the differences they make
  - What is the Difference when the Terminology is the SAME?
    - Virtualization and Partitioning
    - Work Load Manager (WLM)
  - Other Significant System Attributes
    - Security
    - Availability
  - Total Cost of Ownership (TCO)
  - What Does this mean for Application Placement?

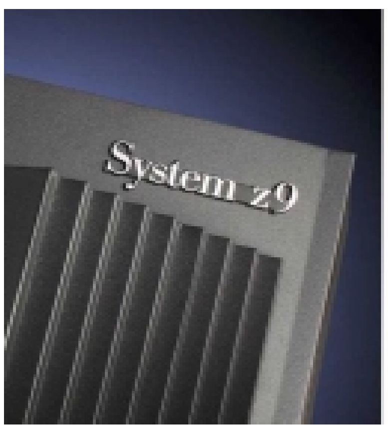


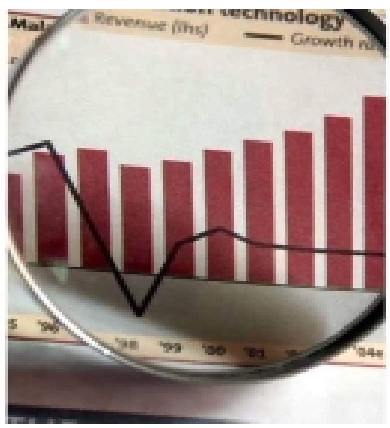
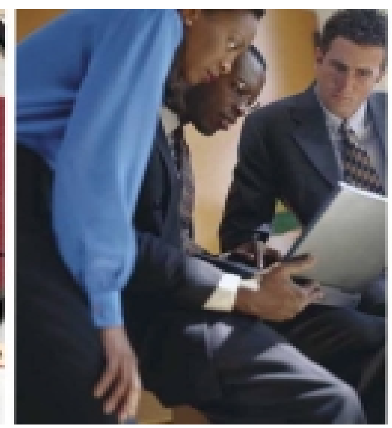
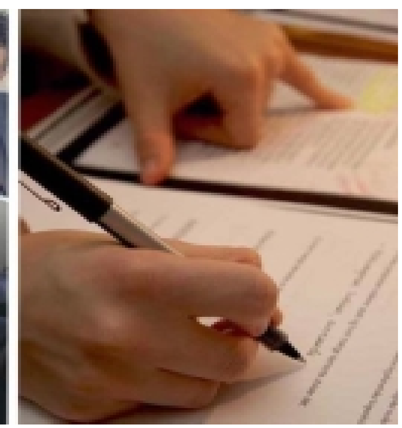

History and Heritage – The Formative Years

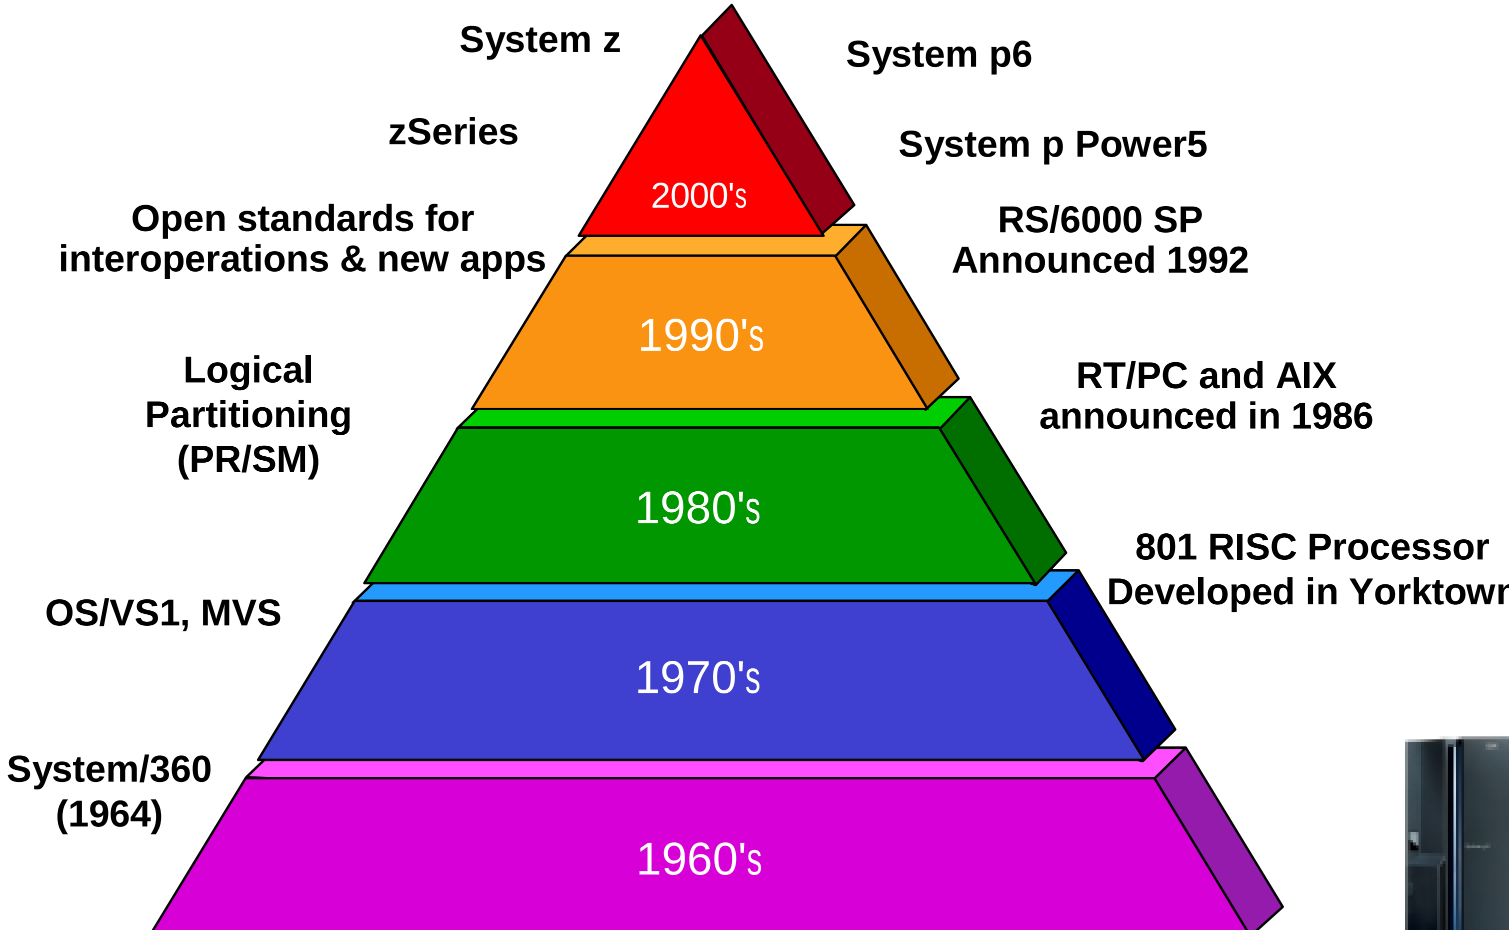

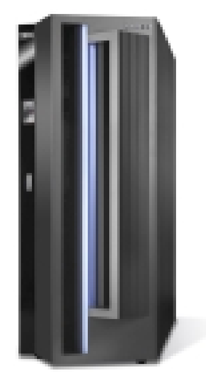

Architecture

- Basic Decisions on architecture
  - What are the decisions?
  - Why were they made?
- How do they effect system design?
- System z – a 54 way system?
  - It’s not just about numbers
  - It’s not just about chips
  - It’s about balanced design

Heritage Matters – OR – *Why Doesn't z run all those popular benchmarks?*

Faced with a blank sheet of silicon a designer must make fundamental choices

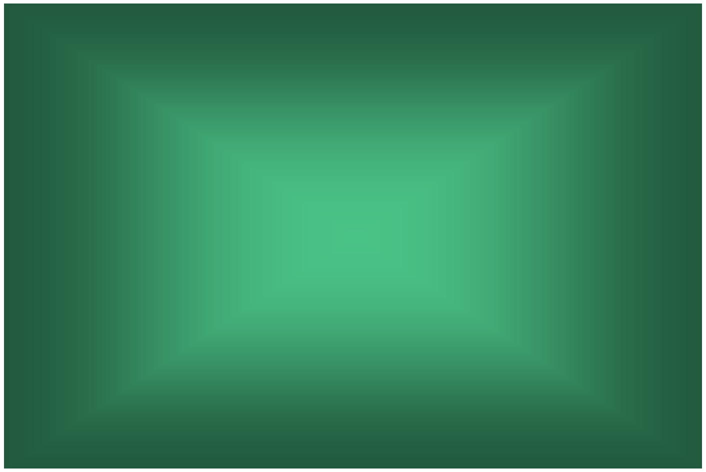

Because - Architecture Matters!

There are two fundamental strategies in the market today

RISC – Reduced Instruction Set Computing

Dedicated / Commodity

CISC – Complex Instruction Set Computing

Shared / Optimized

These strategies drive design tradeoffs…

The more silicon set aside for speed, the less silicon available for Cache and RAS (and vice versa)*

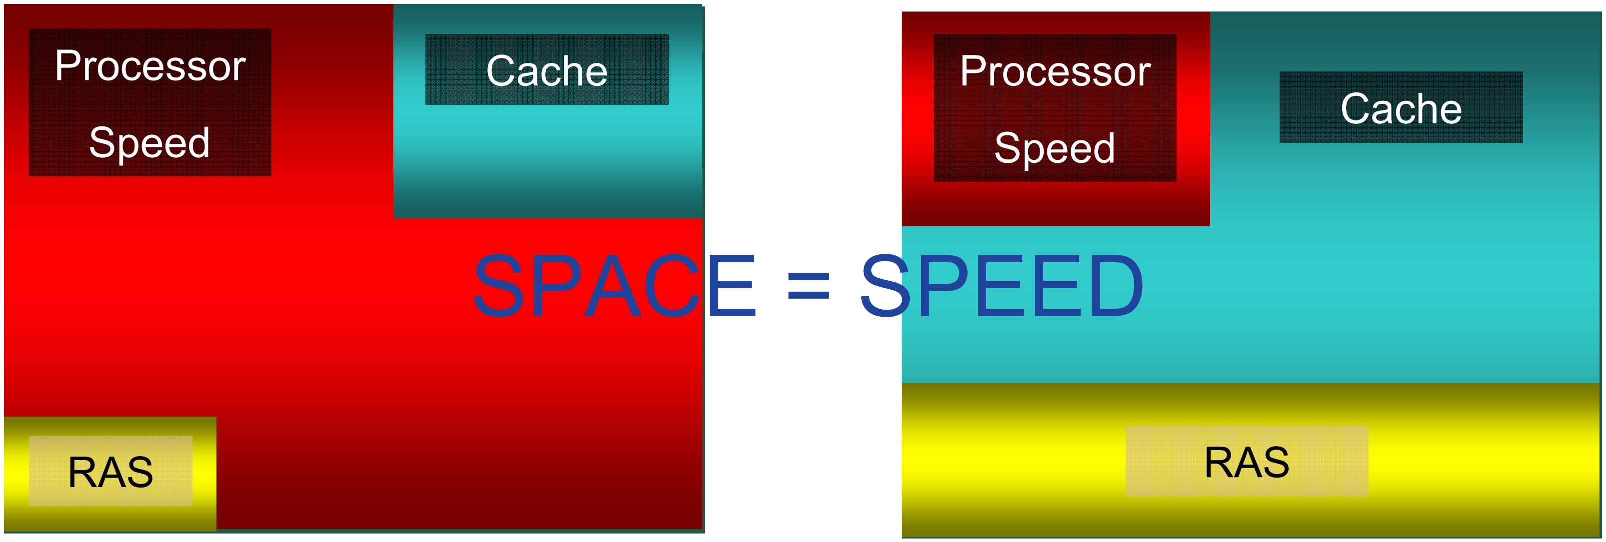

- Small Workload – Small Working Set
- RAS through replication
- Longer Execution Path Lengths
  - Cheap MIPS culture

Replicated

- Mixed Workload – Large Working Set
- Built-in RAS is required
- Shorter – “tuned” workloads
  - Valuable MIPS culture

Consolidated

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

Design choices affect "benchmark" performance SPECint example:*

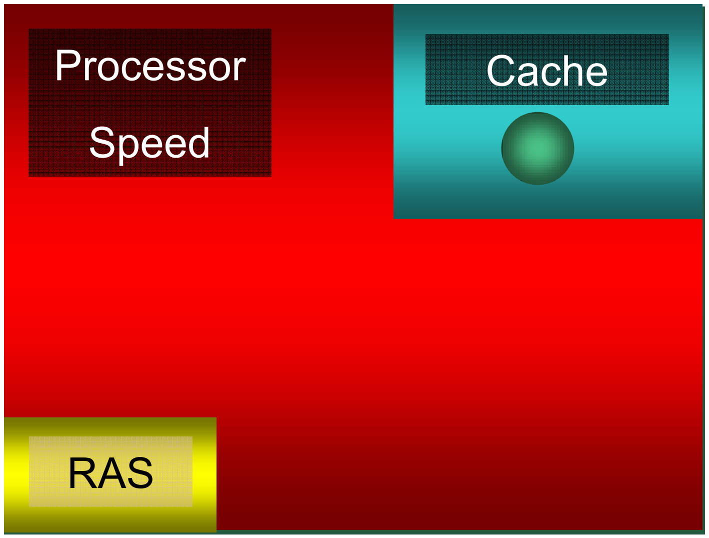

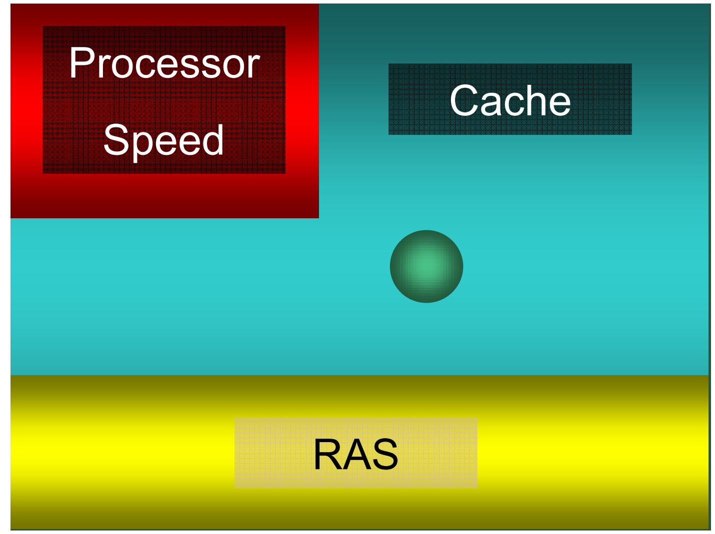

Small Working set fits easily in cache

Maximized Processor speed is leveraged

RAS not measured

Minimized RAS is not "penalized"

Replicated

Large Cache is "too big"

Processor speed / Cache tradeoff penalty

RAS Space is not "valued"

Maximized RAS and Cache are "penalized"

Consolidated

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

Design choices affect "benchmark" performance TPCc example:*

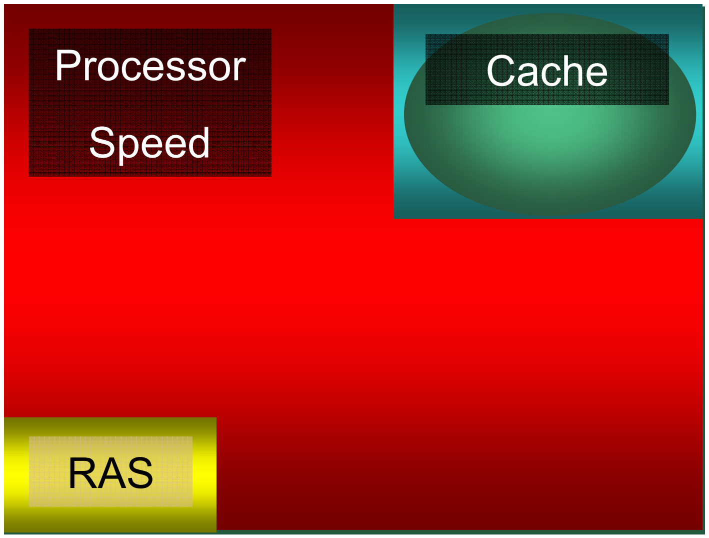

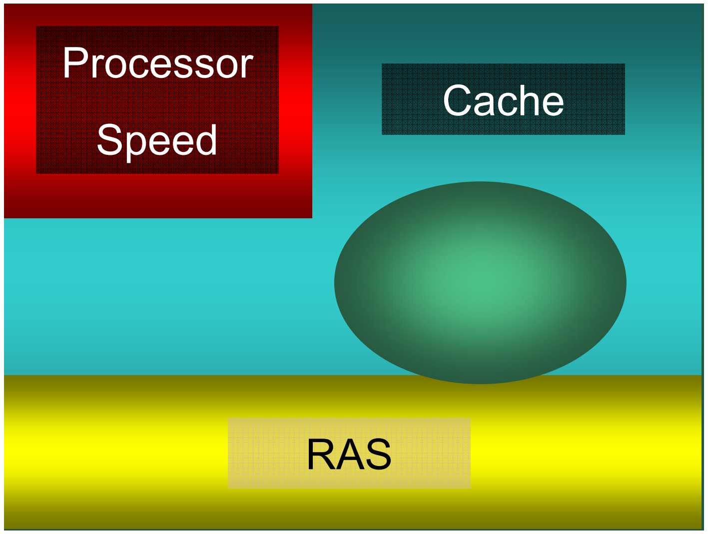

Cache is optimized for this working set

Maximized Processor speed is leveraged

RAS not measured

Minimized RAS is not "penalized"

Replicated

Large Cache is still "too big"

Processor speed / cache trade-off penalty

RAS Space is not "valued"

Maximized RAS and Cache are "penalized"

Consolidated

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

Design choices affect REAL WORLD performance too! - Mixed/WLM/Virtualization*

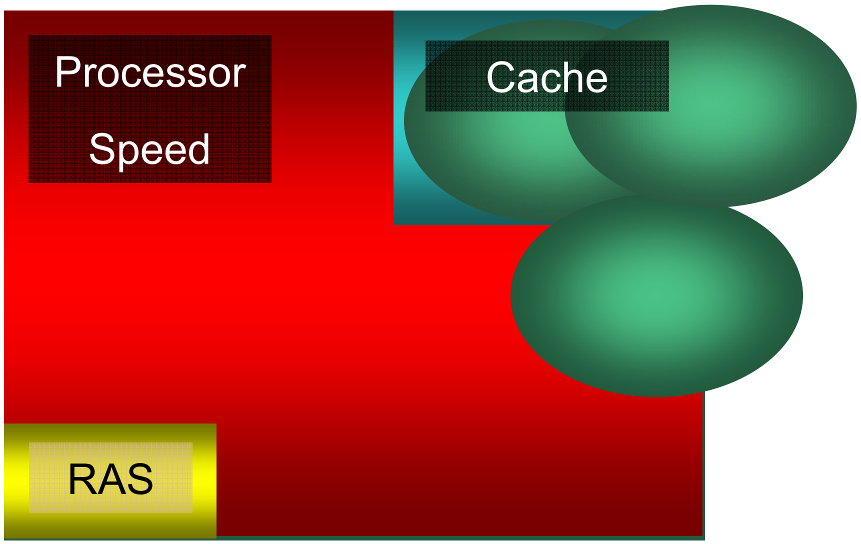

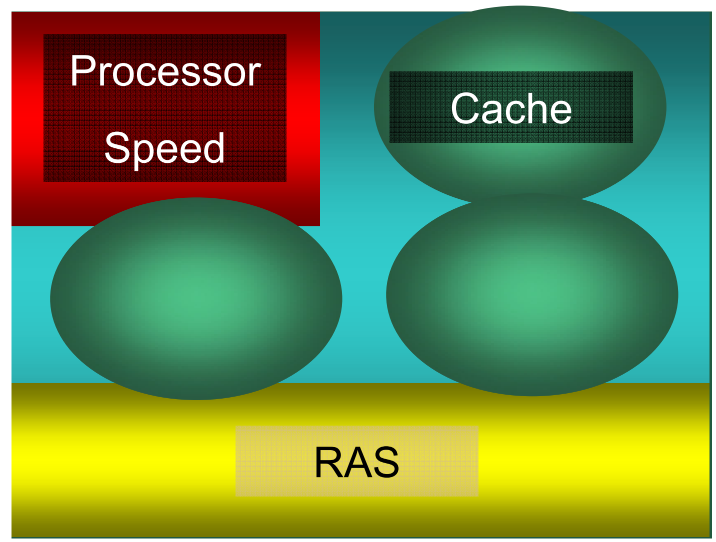

Replicated	Consolidated
Working set(s) too large for cache	Cache contains multiple working sets
Maximized Processor speed penalized	Processor speed optimized by cache
Minimized RAS is "penalized"	RAS Space is "valued"
"Fast" processor is under-utilized	All of "slow" processor is used

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

System z – just a 54 way box?

Unix Servers – 128 Way

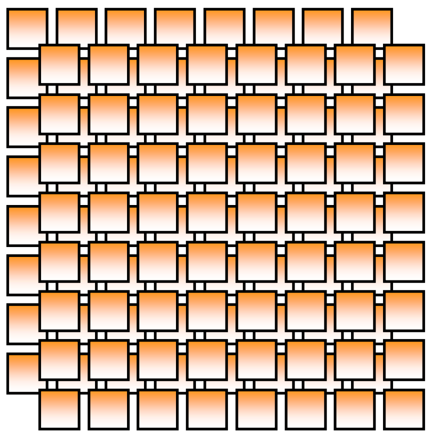

IBM eServer z9-109
CMOS 54-way SMP

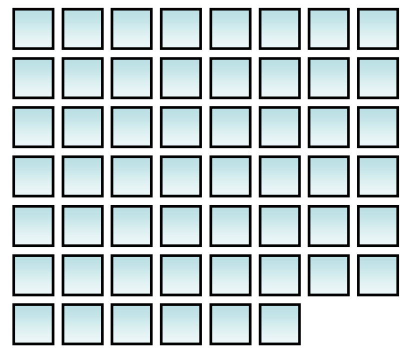

RISC/UNIX Implementation

Up to 128 processors

What do they do?

EVERYTHING!

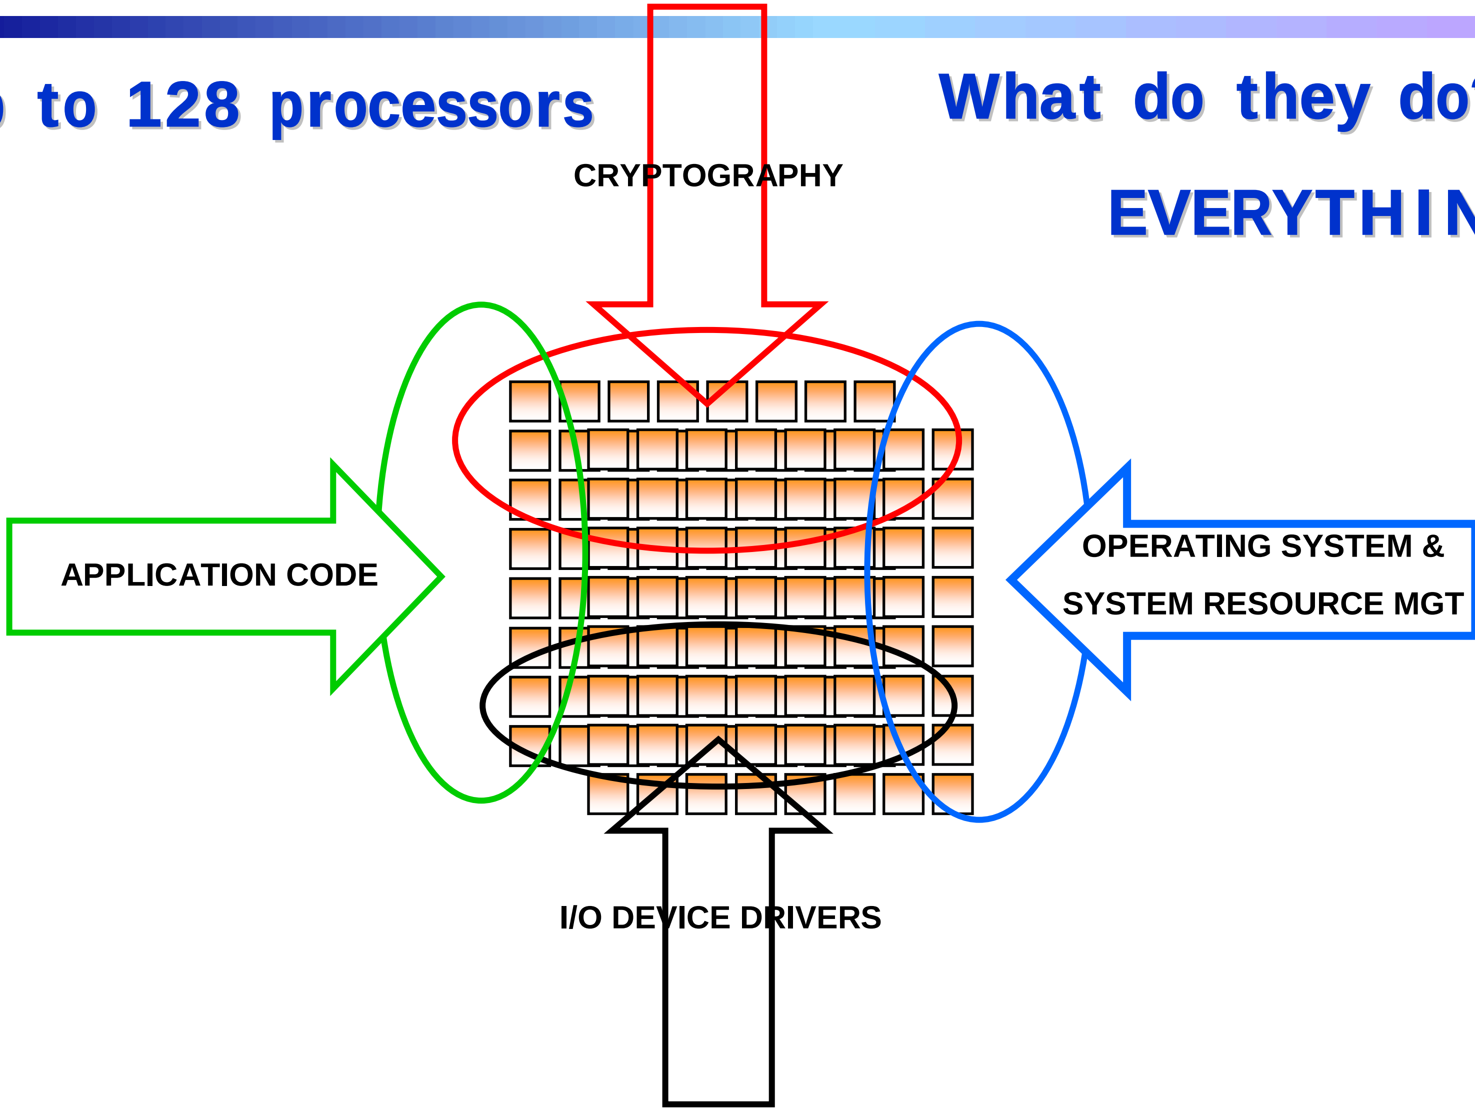

System z Architecture – the 16 processor book is the basic building block

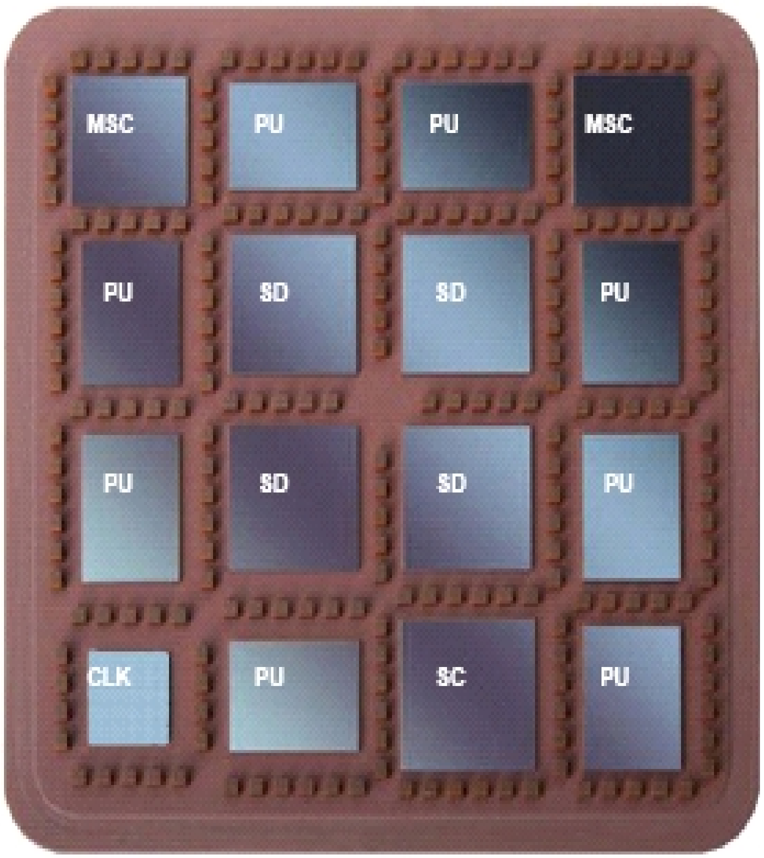

System z Implementation
"How to build a 54-WAY Mainframe"

4 Books, 16 processors per book

4x16 = 54???

Where are the missing 10?

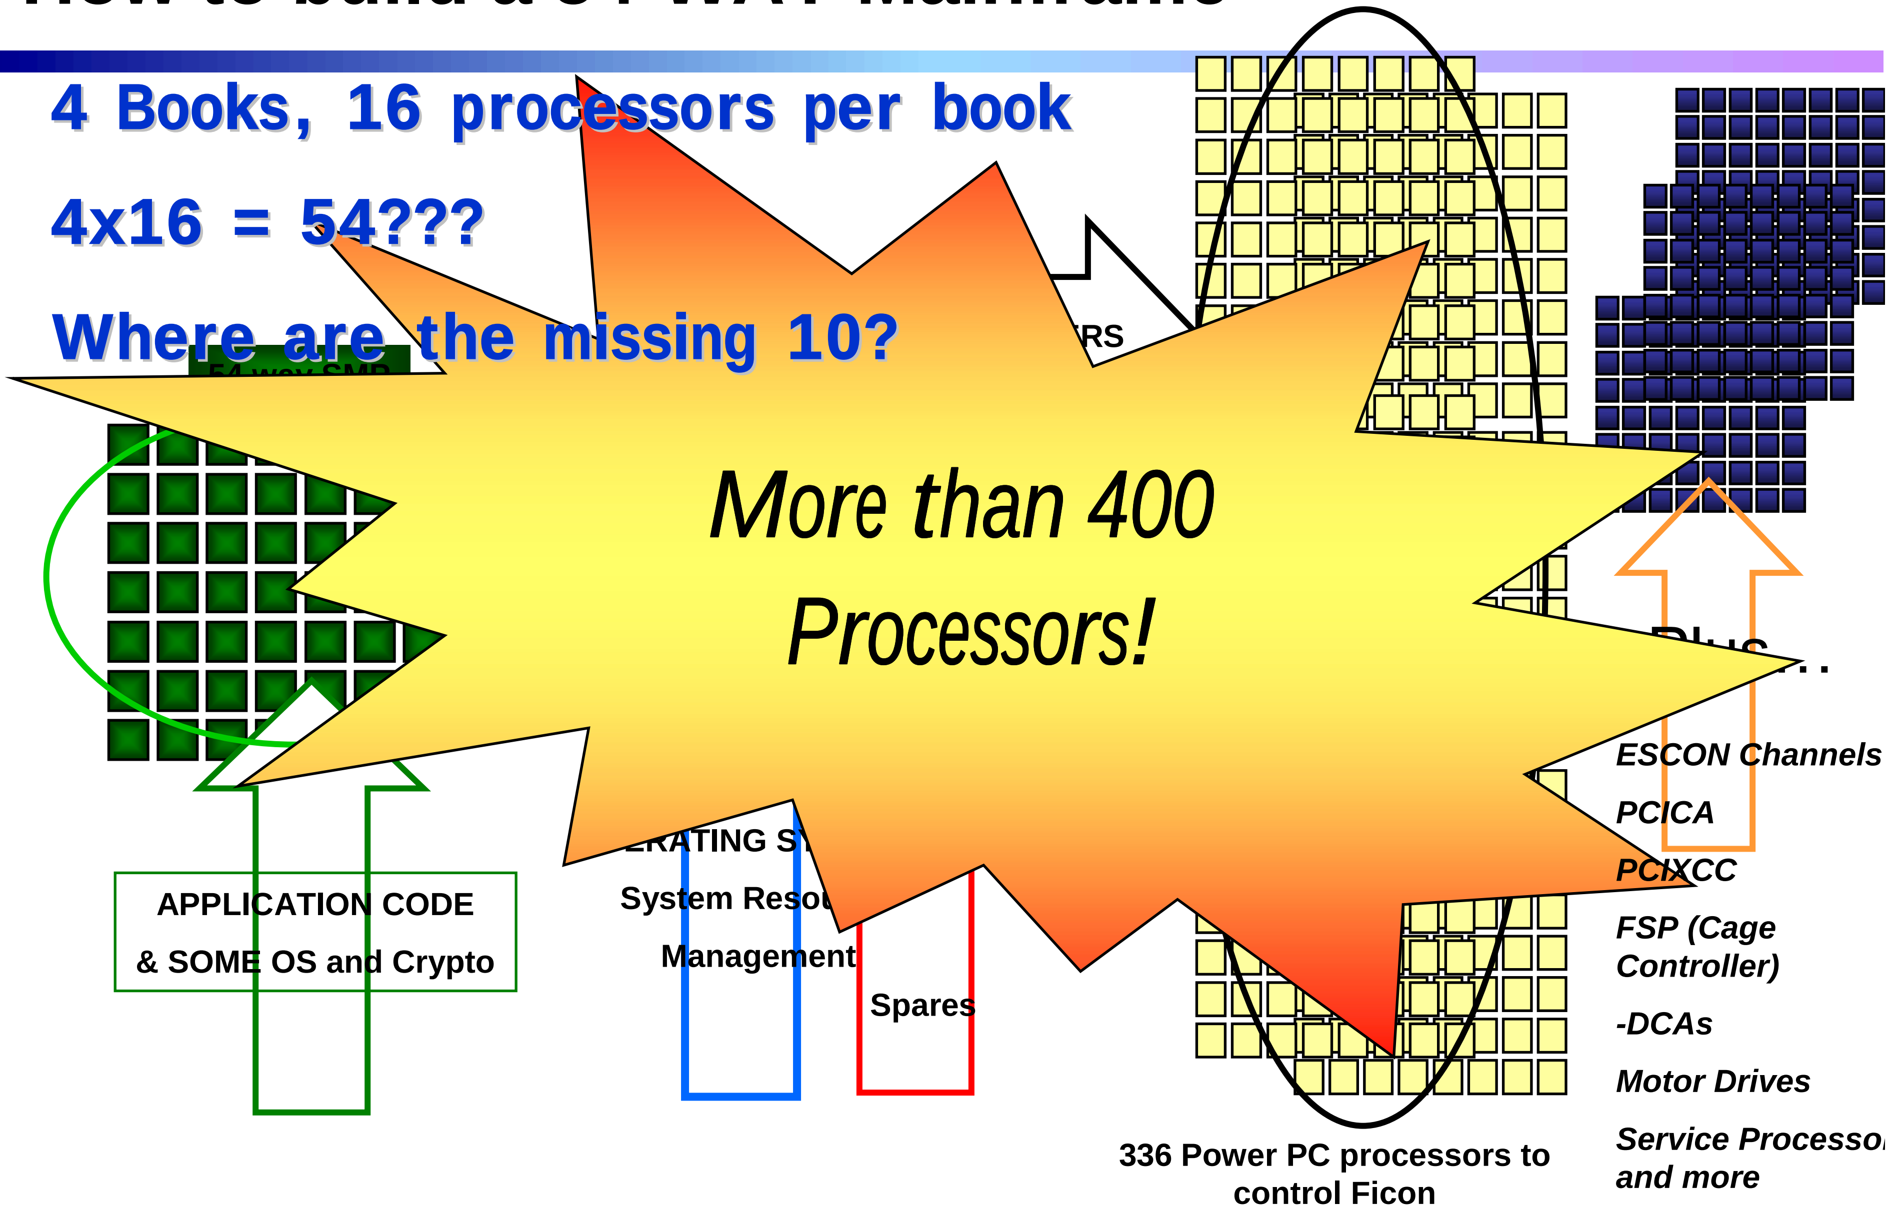

And It's just not about processors …
but balanced resources, well managed

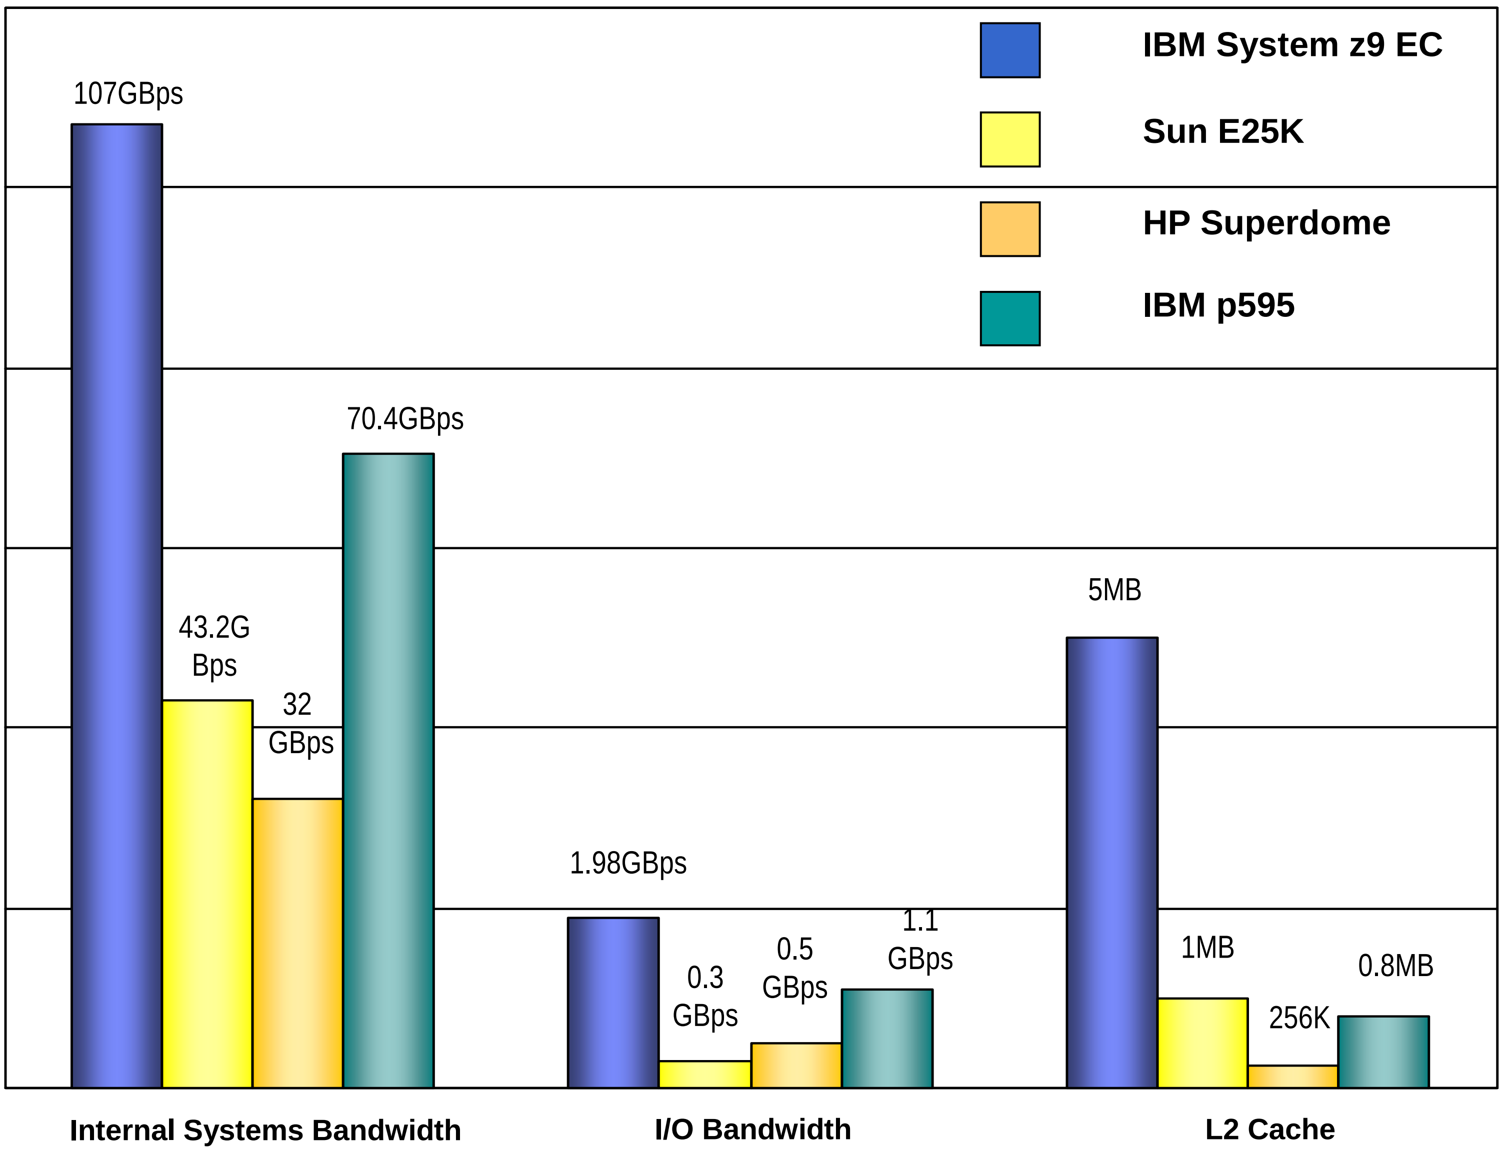

Per Processor Capacity

System p Architecture

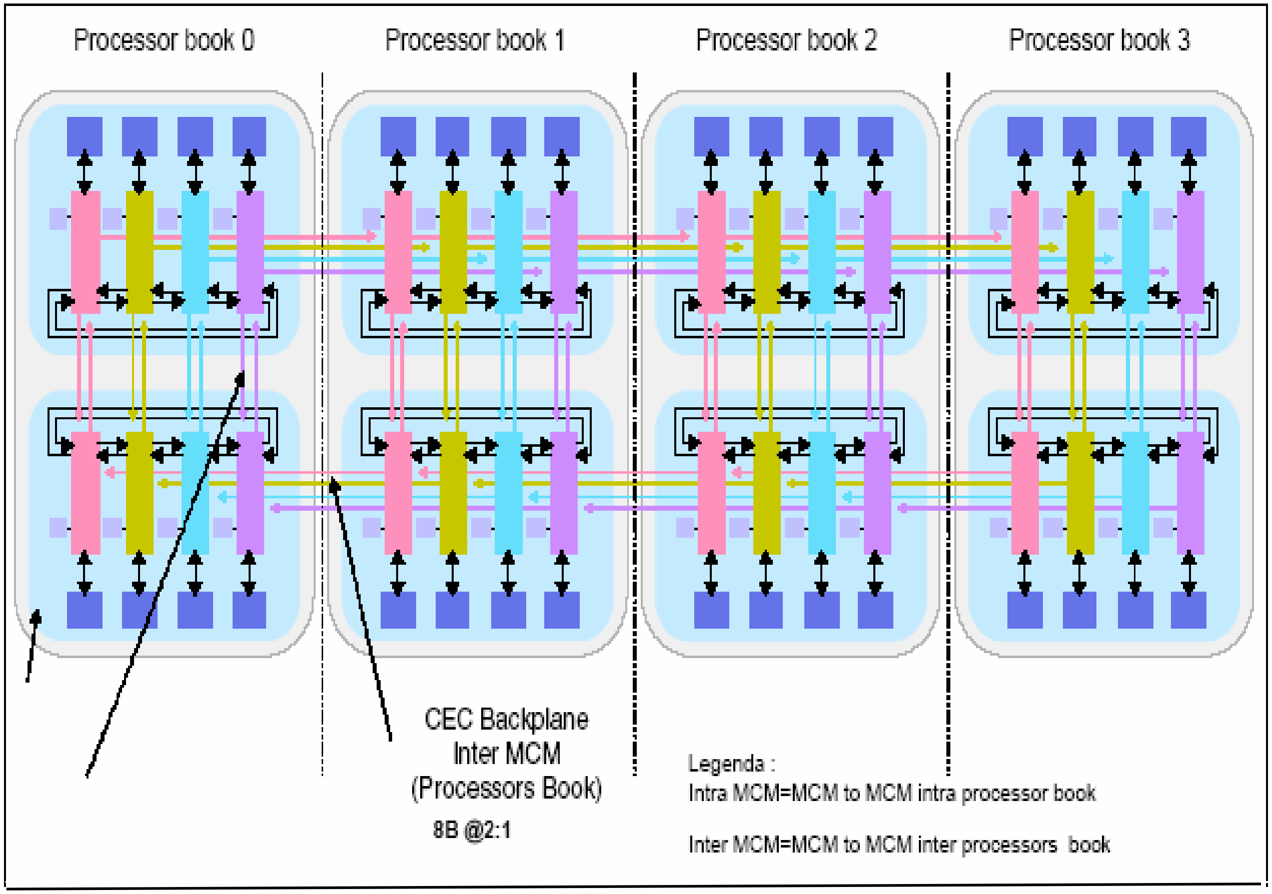

- **CPU, cache, memory controllers on MCM**
- **Memory inside CPU books for local access and control**
- **Interconnect fabric supports "many to many" connections**
- **Designed to minimize traffic, increase throughput and scalability**

System z Architecture

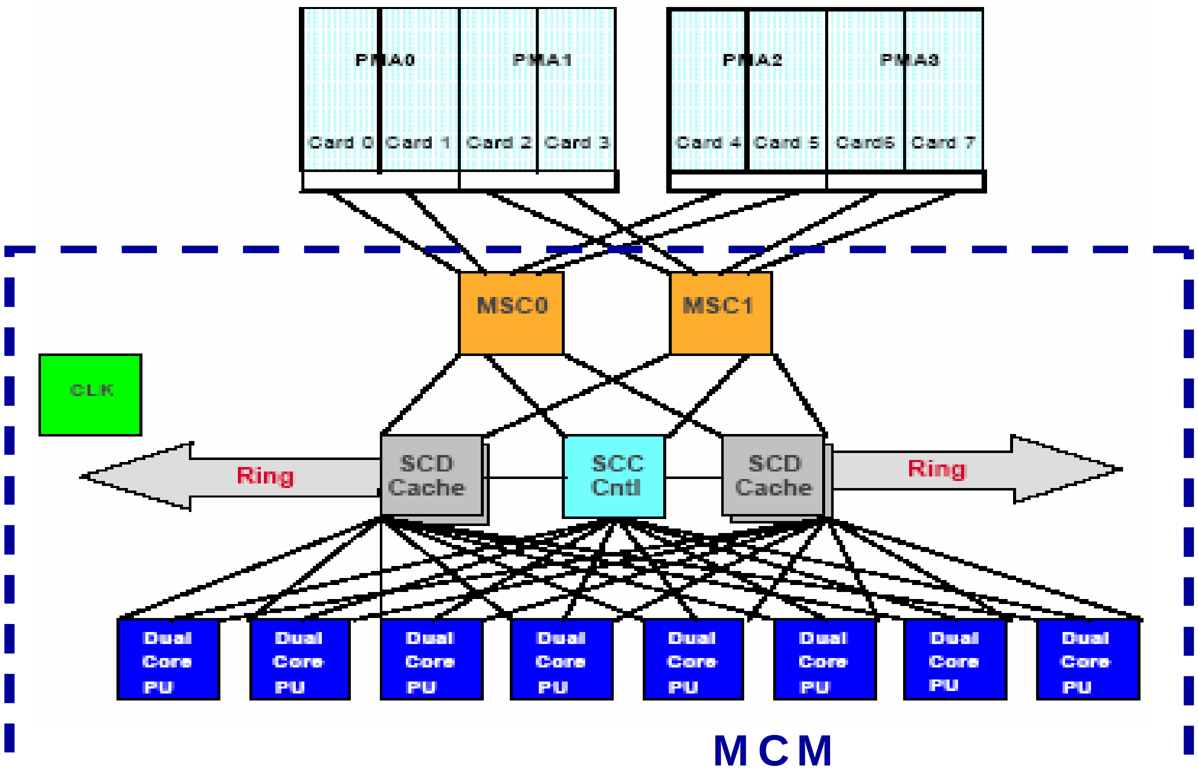

"Shared Everything approach"

- **Flat cache/memory ring crossbar interconnect creates "any to any" connections**
- **Large L2 cache on each MCM with redundant high speed interconnect – runs at processor speed**
- **Bidirectional ring cross point switch between all L2 cache and memory**
- **Near flat cache and memory access across system**

Design Causes Differences

- System Utilization
  - Virtualization
  - Systems Management
- System Availability
  - Planned vs Unplanned
  - System vs Application
- Security

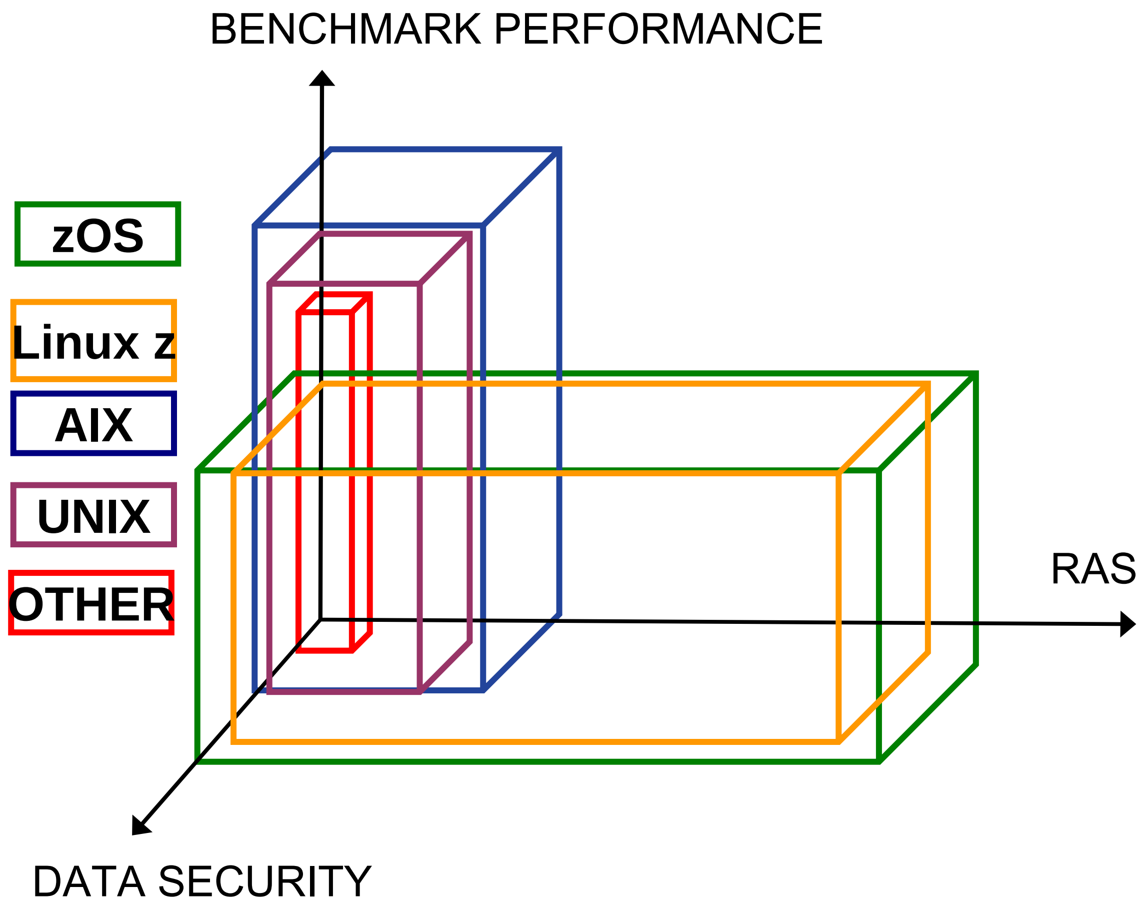

We still hear a lot of terminology

- Logical Partitioning (LPAR)
- Virtualization
- Systems and Workload Management
  - z Continues to Raise the Bar…
  - p Will Continue to adopt leading edge technology

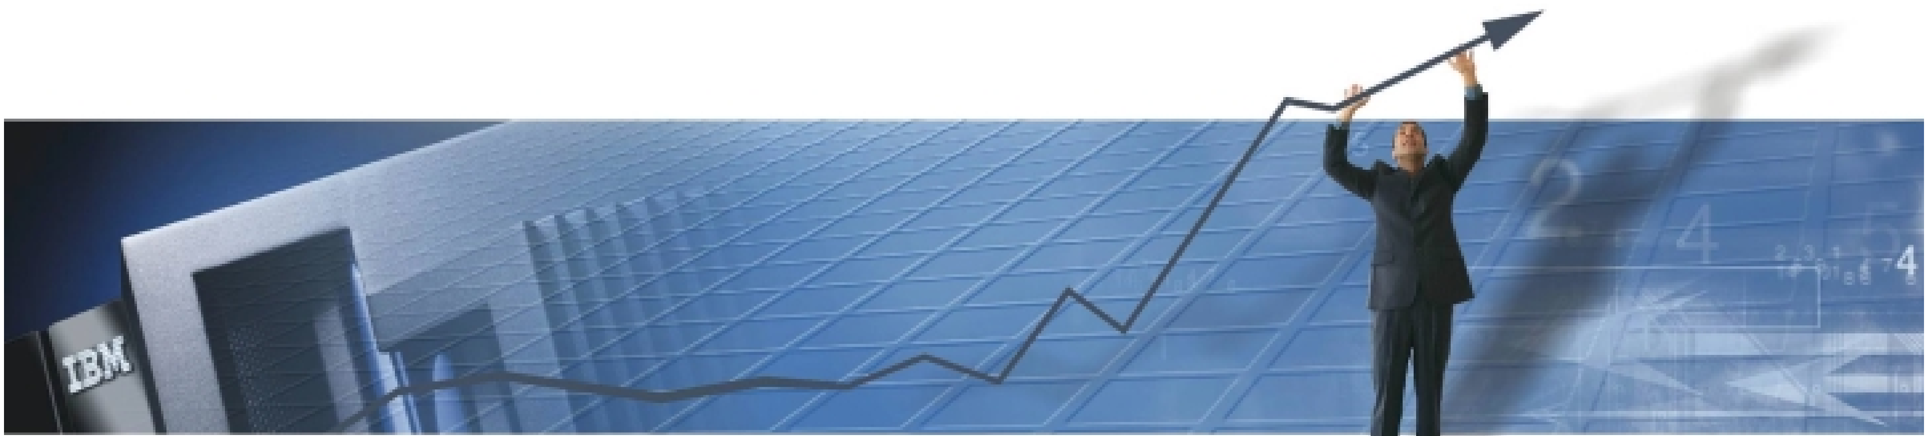

What's the difference when the words are the same?

Sophisticated System Management

Input/Output

- *Unified and enforced I/O model – with Unix I/O Handling is in the Kernel*
- *Cascading of processors, minimum CP involvement*
- Multipath
- Shared DASD standard – with UNIX it's up to you to install
- Channel arch. – CKD
- MIDAW Facility
  - Improves FICON performance for extended format datasets
  - Dramatically reduce channel utilization, fabric traffic and control unit adapter overheads
  - Significantly improve I/O response times

Workload Manager (WLM) – the key to maximizing use of resources

- *Service Level Agreement Driven*

- *Goal Mode based on Business Priorities*
  - *Dynamically assign processors, I/O, memory, TCP/IP connections*
  - *Within and across multiple z/OS images by transaction type, user, and time of day*
  - *End-to-end business management*
  - *Extended to z/OS Initiators*
- VM/RM
  - *Goal mode for multiple guests*
  - *highly granular VM I/O Throttle*

- Coordination of DirMaint and RACF changes for simplified user administration

Workload Management – Dispatching

- Preemptive dispatcher
- Numerous methods of establishing priority and long-term vs. short-term weighting
  - *3 levels of declining time slices*
- ***Intelligent Resource Director*** *extends management to storage, network*
- *Resources balanced and capable of being reassigned dynamically* ***and automatically***
- ***Business goal directives (WLM)***

Workload Manager (WLM)

- Mixed workloads handling is ***built into*** the hardware, hypervisors, operating systems, and clustering
  - Minimizes pathlengths
  - Responsive and stable
  - Efficient and effectively demonstrated by 25+ years of production
  - Availability and systems management of a dynamically changing workload included

System z Virtualization: a Multidimensional Solution

Virtualization is Built in, not Added On

On demand scale *up and out* solutions consist of multiple dimensions of function:

- **Application Support Dimension (open, stable)**
  - Open, stable operating system
  - Virtual server awareness infrastructure
  - Enterprise applications
- **Hypervisor Dimension Capabilities (powerful, flexible)**
  - Shared-memory based virtualization model
  - Granular resource sharing and simulation
  - Flexible virtual networking
  - Resource control and accounting
  - Server operation continuity (failover)
  - Server maintenance tools and utilities
- **Hardware Dimension Capabilities (robust, reliable)**
  - Legendary reliability, scalability, availability, security
  - Logical partitioning (LPAR)
  - Processor and peripheral sharing
  - Interpartition communication
  - Virtualization support at the hardware instruction level (e.g., SIE)

How does the Virtualization stack up?

	IBM		HP		SUN		VMWare
Family	**System z**	**System p,i**	**Superdome**	**Superdome**	**3900-6800 Sunfire**	**Solaris Sparc**	**Intel/AMD**
Solution	***PR/SM z/VM***	**Micropartioning, VIOS**	**HPnPar**	**HPvPar**	**Domains**	**Containers**	**VMWare/ESX**
OS Support	*z/OS, VSE, TPF, Linux, z/VM*	AIX 5.3, 5.3 iOS,Linux	HPUX 11	HPUX11	Solaris 8,9,10	Solaris 10	Windows, Linux, Netware
Granularity	*Small % cpu*	1/10 cpu	Min 4 cpu on board boundaries	1 cpu	Min 4 cpu on board boundaries	>8k partitions per Solaris image	1/8 cpu
Security/Fault Isolation	*HW/Firmware, Hipervisor, SW*	HW/Firmware/Hipervisor	Electrical isolation	Not HW enforced	Electricl isolation	Single os kernel exposed	Hipervisor
Scalability	*54*	64	64-128	64-128	72(144 cores)	72(144 cores)	4
Max partitions	*60/>10000*	256	16	128	18	18(per OS)	32

How does the Virtualization stack up?

	IBM		HP		SUN		VMWare
Family	**System z**	**System p,i**	**Superdome**	**Superdome**	**3900-6800 Sunfire**	**Solaris Sparc**	**Intel/AMD**
Solution	***PR/SM z/VM***	**Microparti oning, VIOS**	**HPnPar**	**HPvPar**	**Domains**	**Container s**	**VMWare/ESX**
CPU Capacity Sharing	*Yes*	Yes	No	No	No	Yes	Yes
Dynamic Reconfigurat ion	*Yes*	Yes	No	Yes	Yes	Yes	No
Support for Dedicated I/O	*Yes*	Yes	Yes	Yes	Yes	No	No – all I/O must be virtualized
Shared Resources	*CPU, Memory, I/O Channels*	CPU,I/O	No	No	No	CPU, Memory, I/O	CPU, Memory, I/O
Interpartition Communicat ions	*Yes*	Yes	No	No	No	No	Yes

System z Virtualization and Partitioning

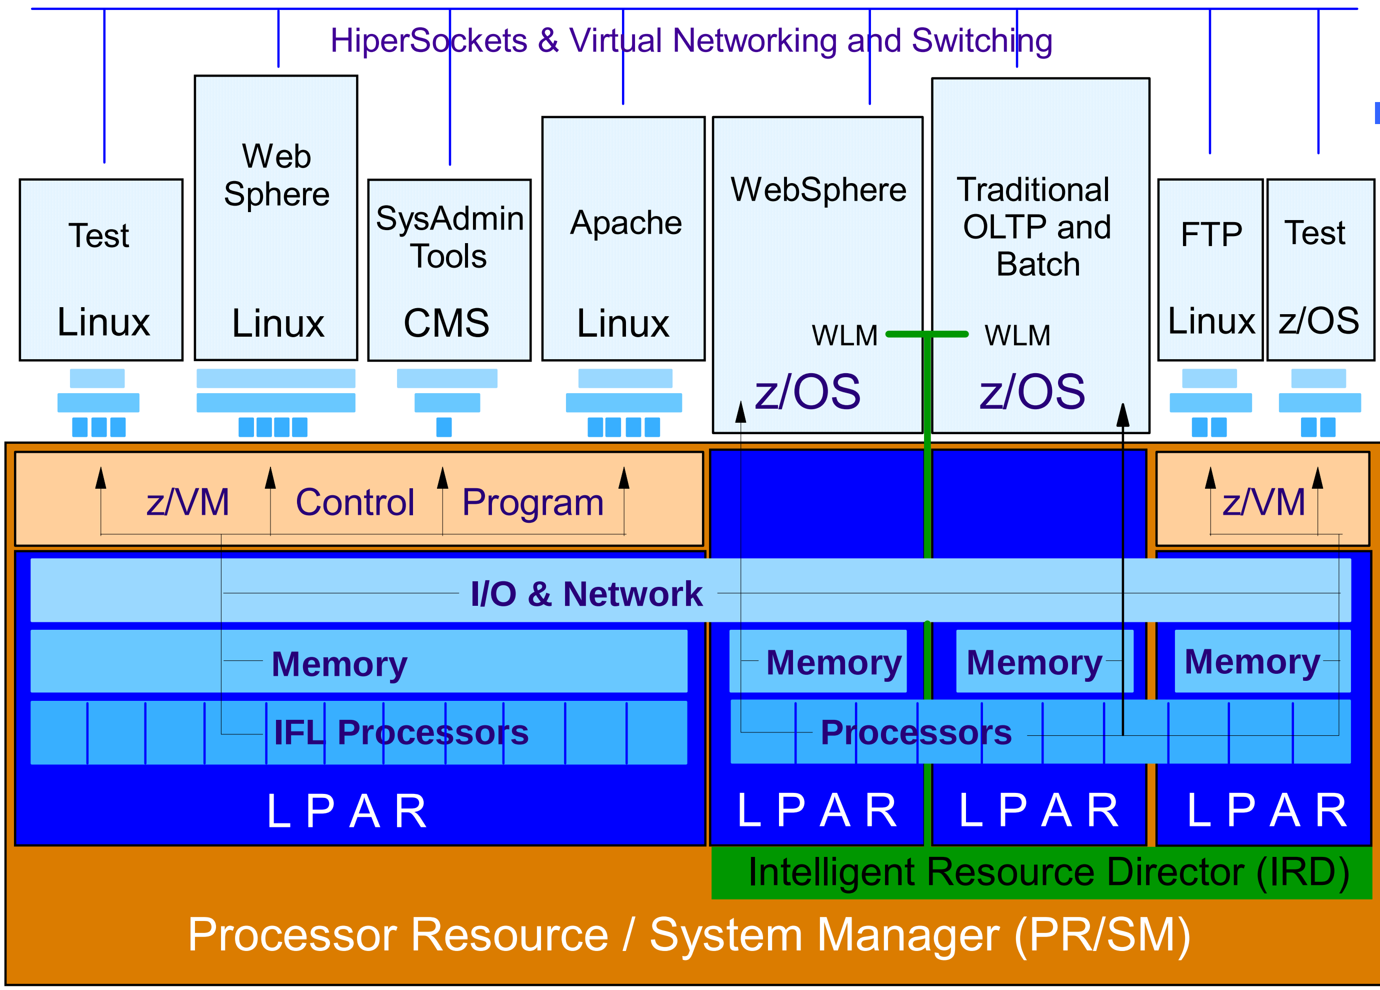

- Multi-dimensional virtualization technology
  - System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS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

* Excluding non-shared resources such as Integrated Facility for Linux

Differences in Resource Management and Virtualization

- Unix WLM allocates resources in response to system *internal* utilization rules established by the system administrator.
- System z can dynamically reallocate resources to maintain “end-user service levels” –prioritized by business goals– using *external* performance measurements.
- This results in higher utilization, since the system can schedule work on an ad hoc basis to sustain utilization levels approaching 100%.

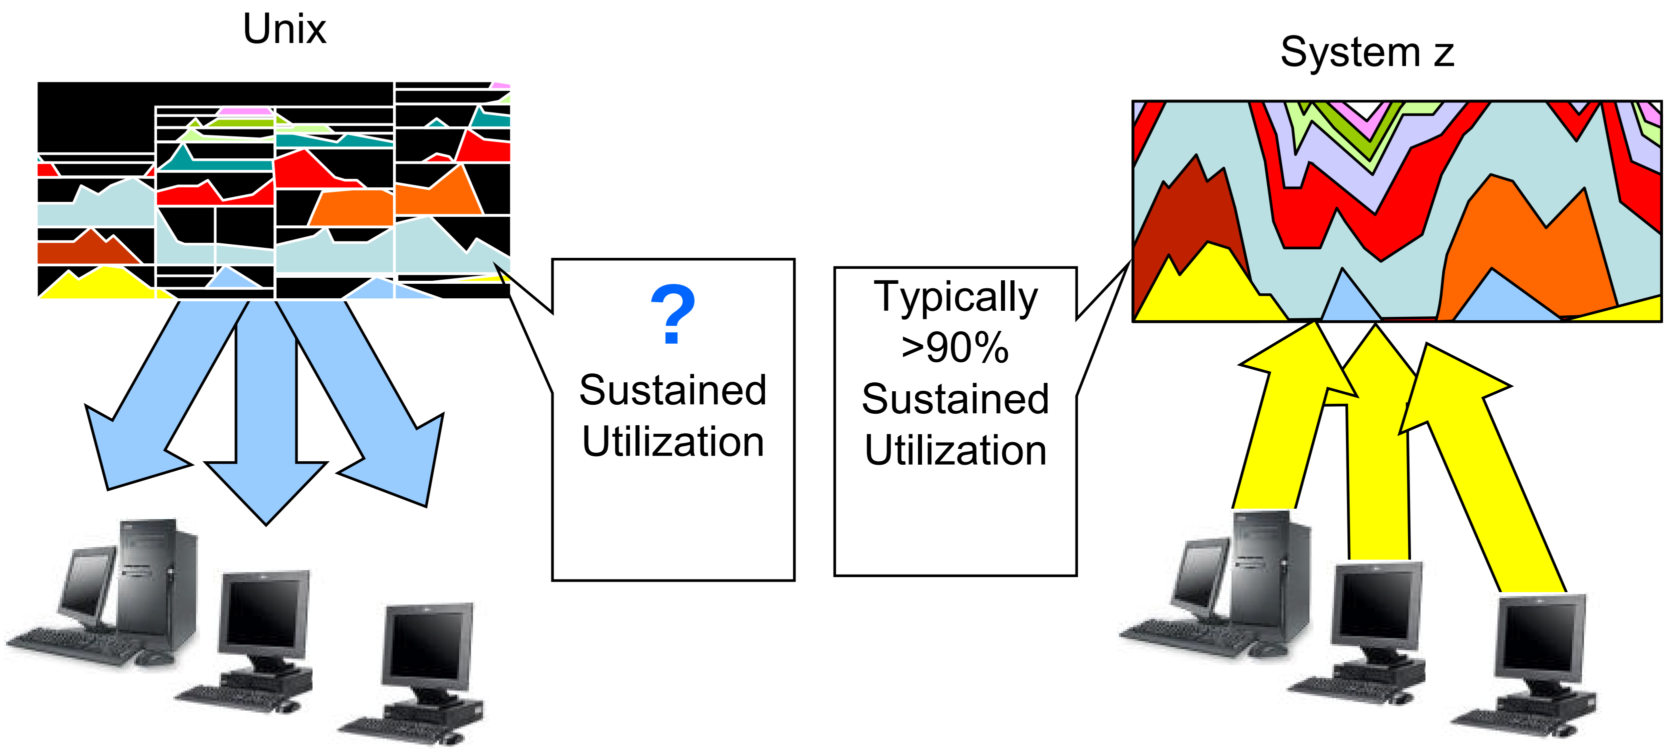

RESULT: Potentially Substantial differences in system utilization and end-user service level assurance/end-user productivity

Outage Avoidance and Security

What does the z have that can help avoid expensive downtime?

Multiple Features

- *FRRs, error percolation*
- *Pageable nucleus*
- *Application Recovery*
- *HyperSwap*
- *GDPS Open LUN Management*
- *Control Unit Initiated Reconfiguration*

- *Data Recovery*
  - *FlashCopy at dataset level*
  - *Concurrent Copy*
- Redundant I/O Interconnect
  - Potentially Reduce Unplanned Outages
- Concurrent Book and Memory Add/Replace/Repair
  - Potentially Reduce scheduled downtime
  - *Capacity Backup Upgrade (CBU) for all engines including* IFLs, zAAPs or ICFs

System z can help avoid downtime

System z

- Storage Protect Key

- Spare processor is ***always present***

- Sparing concurrent
  - If found before hard fail
  - ***If detected during execution because of Dual Instruction with correct and error detect and sharing of processor state***

IBM System z9 Concurrent Upgrade
Customer Controlled

- On/Off Capacity on Demand – Temporary upgrade
  - Non-disruptive temporary addition of CPs, IFLs, ICFs, zAAPs, and zIIPs
  - Customer orders and installs upgrade via Resource Link and IBM RSF
  - Non-disruptive removal when capacity is no longer wanted
- Customer Initiated Upgrade (CIU) – Express – Permanent upgrade
  - Customer capability to order and install permanent upgrade
  - Customer orders and installs upgrade via Resource Link and IBM RSF
- Capacity Backup (CBU) – Temporary emergency capacity upgrade
  - Non-disruptive temporary addition of CPs, IFLs, ICFs, zAAPs, and zIIPs in an emergency situation
  - CBU contract required to order CBU features and CBU LIC CC
  - Customer activates upgrade for test or temporary emergency
  - Concurrent downgrade after test or recovery completed

The System z Security

- RACF Security
- Specially designed microprocessors
- *CP Assist for Crypto Function (CPACF)*
- Integrated cryptography features offer more security options on System z
  - Integrating Advanced Encryption Standard (AES) support in System z hardware
  - Stronger hashing algorithm with SHA-256
  - Integrated Pseudo Random Number Generator
- z/OS feature complements TCP/IP based sensors
- Storage Protect Key
- System z has led in SSL benchmarks for the past 8 years
- Crypto Express2
  - configurability options on the System z for the PCI-X adapters - define either two coprocessors, two accelerators or one of each
  - ***Accelerator mode can improve SSL performance over being configured for coprocessor mode***
- *Only encrypt and decrypt your critical data with DB2 on System z*

Security Certification - The only platform with both *EAL5AI and FIPS 140 level 4 security classification from the US Federal Government*

So………………..
Why Do We Care????

System z can provide cost advantages

TCO -- Components

TCO is comprised of direct and indirect costs.

- Acquisition Cost (Direct)– Cost of initial acquisition and planned upgrades
- Maintenance Cost (Direct) -- Cost of required maintenance above warranty
- Operating System SW Cost (Direct) – Cost of operating system for servers
- Application and Database SW Cost (Direct) – Cost of application and database SW
- Special Management SW Cost (Direct) – Cost of any specific/unique SW
- Facilities – Power/Space/Cooling (Direct) – Cost of facilities
- Networking Cost – (Direct) – Cost of communications
- Application Admin Cost (Direct) – Cost of SysAdmin personnel
- Operations Cost (Direct) – Cost of standard Operations
- Training Cost (Indirect) – Cost of any specialized training
- Conversion Cost (Direct)– Cost of converting existing systems
- Implementation Cost (Direct) – Cost of any roll out plan

Many of the qualities of the System z relate to Indirect Cost Items and when they can be quantified to monetary terms are valuable to include – some examples:

- ***Availability***
- ***Reliability***
- ***Productivity Impact***

TCO is more than just the price of hardware

Main components of TCO

- **Hardware**
- **Software**
- **Other (environmentals, security, productivity-loss…)**
- **People**

People expense is now the dominant component!

… and the profile is changing

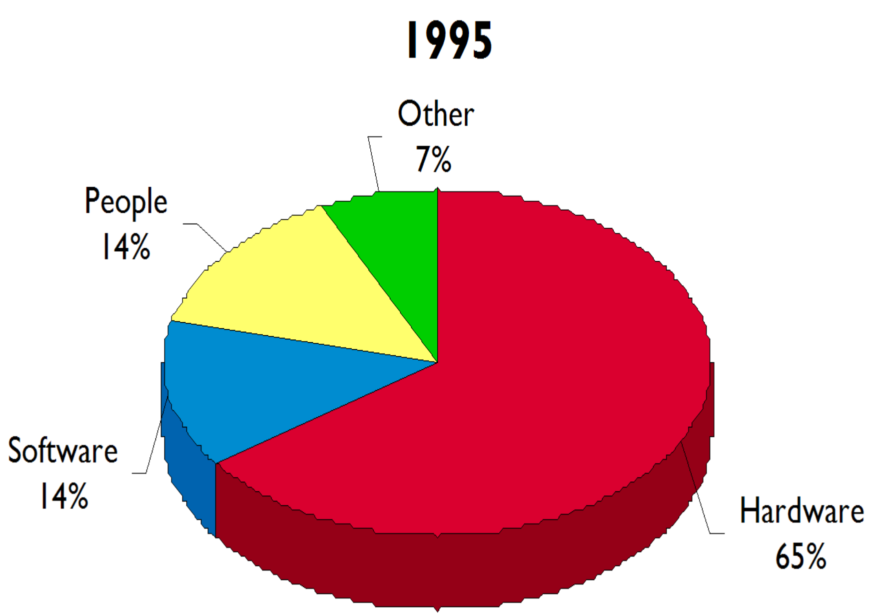

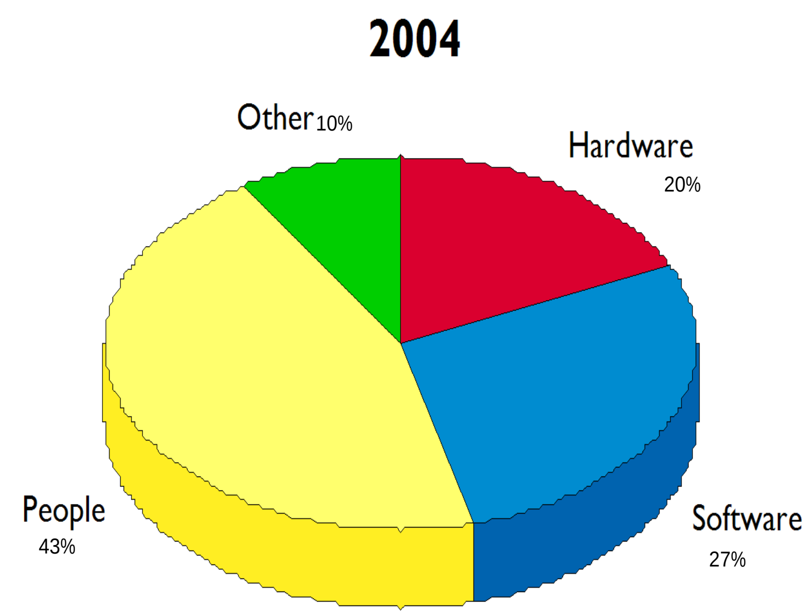

Based on IBM Scorpion customer analyses

System z managing growth and complexity

Volume of Workloads Processed Have Never Been Larger

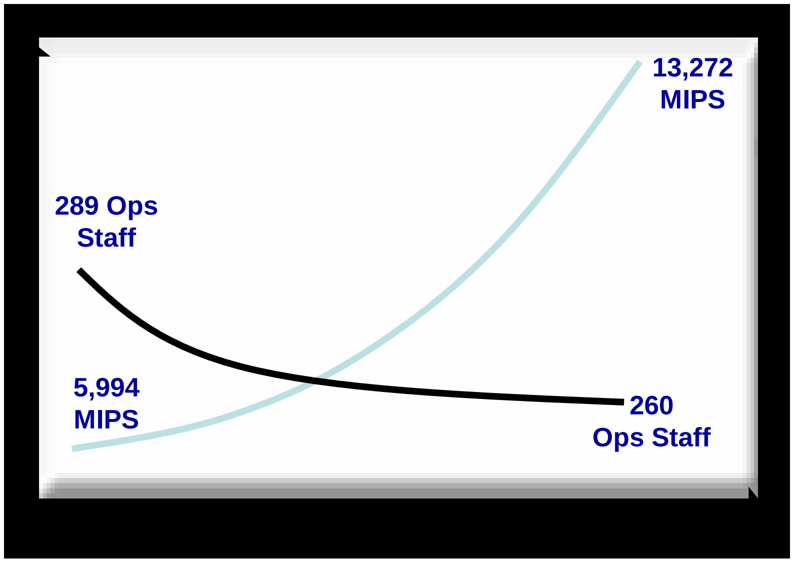

Data center staffing levels have not significantly changed despite large increase in volume.

"Since we published our last high-level perspective of the ratio between MIPS and head count in 2001, the largest ***z/OS installations have more than doubled their 'MIPS to head count' ratio****."*

L. Mieritz, M. Willis-Fleming – Gartner, 2004

" ***the number of operators and systems programmers required per mainframe MIPS has fallen ten-fold in the past seven years, and is expected to at least halve again in the next five years. "***

Arcati – The Dinosaur Myth, 2004 Update

Large Administrative Overhead has other consequences

The Causes of Downtime

Avoid Downtime – Minimize human interaction with the system!

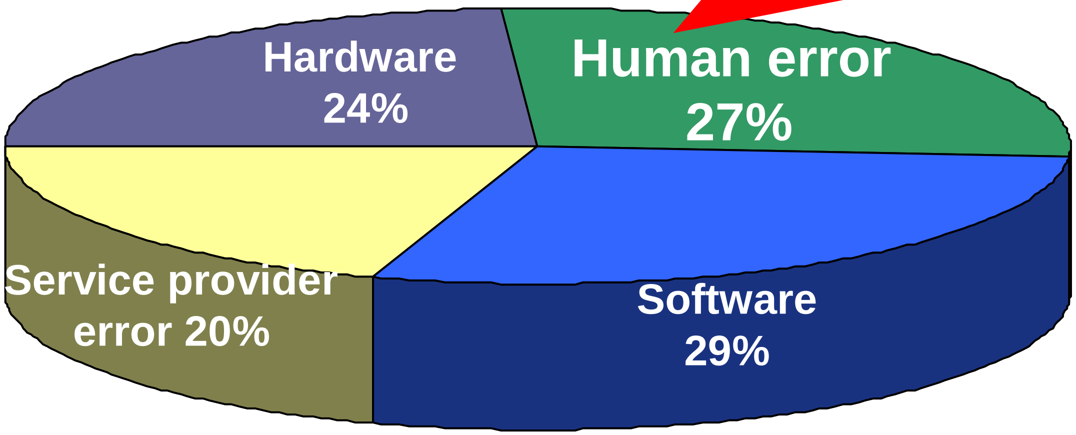

Annual hours 1,180

… there isn't a single problem area that organizations need to focus on, which would be a simpler fix. Every decision is critical, from hardware selection, to product setup and from employee training to SLAs with service providers. ***Human error is the most troubling, because fixes for human error are elusive and require process changes and retraining, which can take a long time and be very expensive***

The Costs of Enterprise Downtime:
North American Vertical Markets 2005
January 2005 – Infonetics Research

Hidden Costs – Downtime and Security breaches

- The effect on the business:
  - Escalating costs
  - Customer loyalty
  - Market competitiveness
  - Regulatory compliance
  - Business Reputation

- 90% of companies surveyed detected a security breech in the last twelve months.
- 80% admitted that the breech lead to new financial costs.
- 44% (223 companies) reported total losses of $455.848,000 as a result of cyber incidents.
- 74% reported that their Internet connection was frequently attacked.
- 34% reported cyber attacks to law enforcement.

Source: The 2002 FBI/CSI *Computer Crime and Security Survey*

Financial Impact of Downtime By Industry

Energy	**$2,817,846**
Telecommunications	**$2,066,245**
Manufacturing	**$1,610,654**
Financial	**$1,495,134**
Information Technology	**$1,344,461**
Insurance	**$1,202,444**
Retail	$1,107,274
Pharmaceuticals	$1,082,252
Banking	$996,802
Consumer Products	$785,719
Chemicals	$704,101
Transportation	$668,586

Source: Robert Frances Group, 2005

On Security…

- **"By one estimate,** *53 million people -* **including consumers, employees, students, and patients –have had data about themselves exposed over the past 13 months."**
- **"In the past few weeks, some of the largest U.S. banks have had to reissue debit cards, all because of data theft."**
- **"Last month, one IT worker at a major health system was fired and three others quit following the theft of computer tapes that held sensitive patient information, including addresses, phone numbers, and Social Security numbers. Staffers at the health care concern routinely brought the tapes home to back up the data on systems there.** *Predictably, the tapes were stolen from a van parked outside one of their residences."*
- **"This gaffe shows how the financial and legal aftereffects of a breach spread like a bad rash. A company** *official estimated that he expects the case, not including litigation costs, to cost from $7 million to $9 million."*

Source: Information Week, 3/20/06, "The Losers"

TCO trends driving Virtualization – and z knows Virtualization!

- Simplify Systems Management radically
- Basis for new IT delivery and provisioning paradigm
- Power is the fastest growing component of cost

Source: IDC, Virtualization 2.0: The Next Phase in Customer Adoption, Doc #204904, Dec 2006

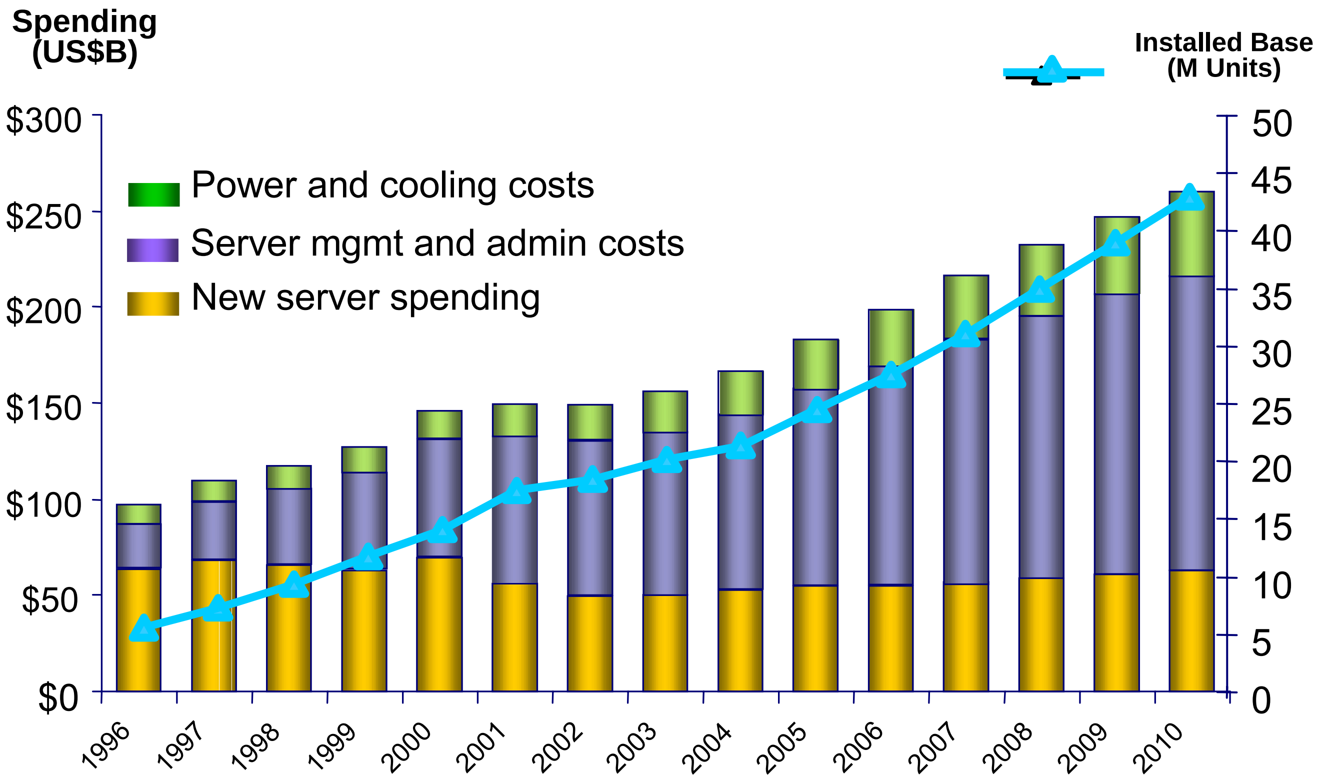

1995
HW = 65%
SW = 14%
People = 14%
Other = 7%

IT budget trends…

2004
HW = 18%
SW = 28%
People = 45%
Other = 9%

Data centers are at a tipping point

- Left unchecked, the cost to power and cool servers in the future may well equal the cost of acquisition.
- If IDC 2010 forecast holds, the cost to power and cool servers in the data center will increase 54%.
- IT executives now rank power and cooling in the top 5 among current concerns.

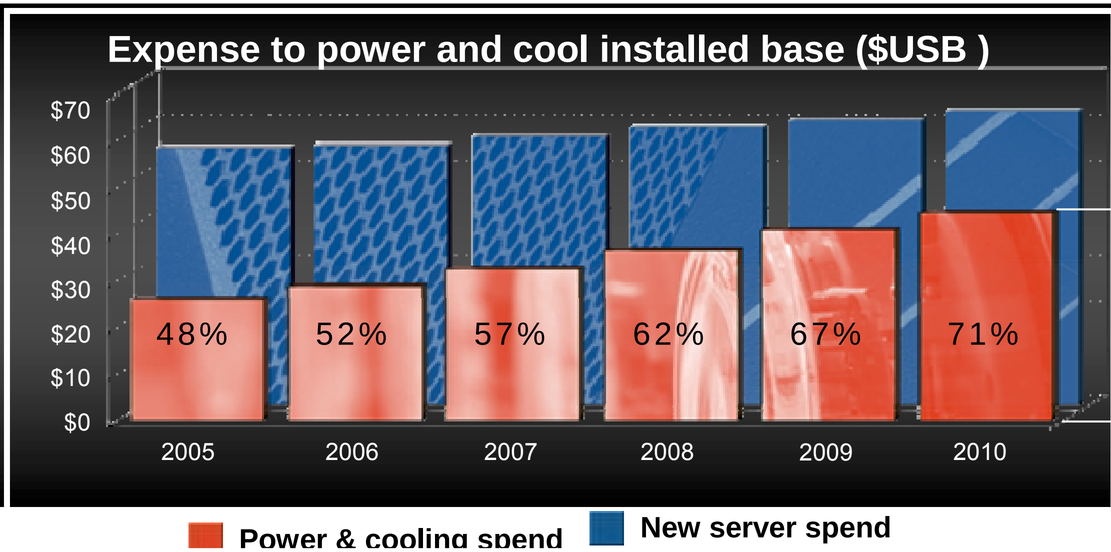

The expense to power and cool the installed base of servers is projected to grow four times compared with the growth rate for new server spending

Source: IDC, 'Worldwide Server Power and Cooling Expense 2006-2010,' Document #203598, Sept. 2006

The IT economics of becoming Lean and Green

- Distributed servers typically run at utilization levels in the range of 5% to 20%
  - Production servers, development servers, test servers
- Virtualization and workload management enables efficient consolidation on the mainframe
  - Run multiple images on fewer processors
  - Achieve utilization levels of 85% or more
- Become Lean and Green through IT consolidation and simplification

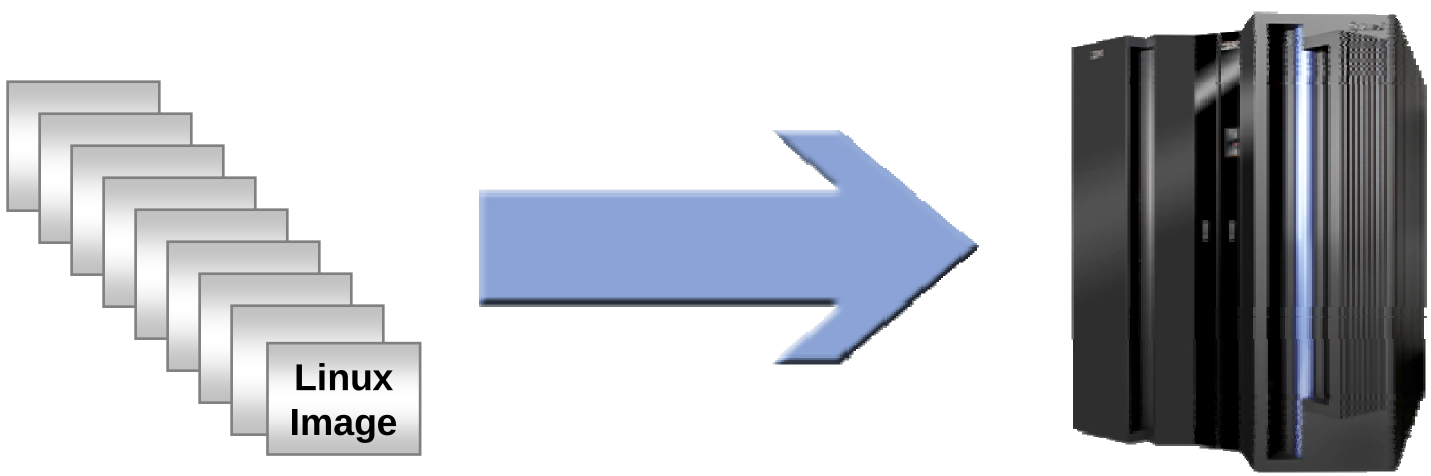

Be Lean and Green through System z technology

System z energy efficiency gains

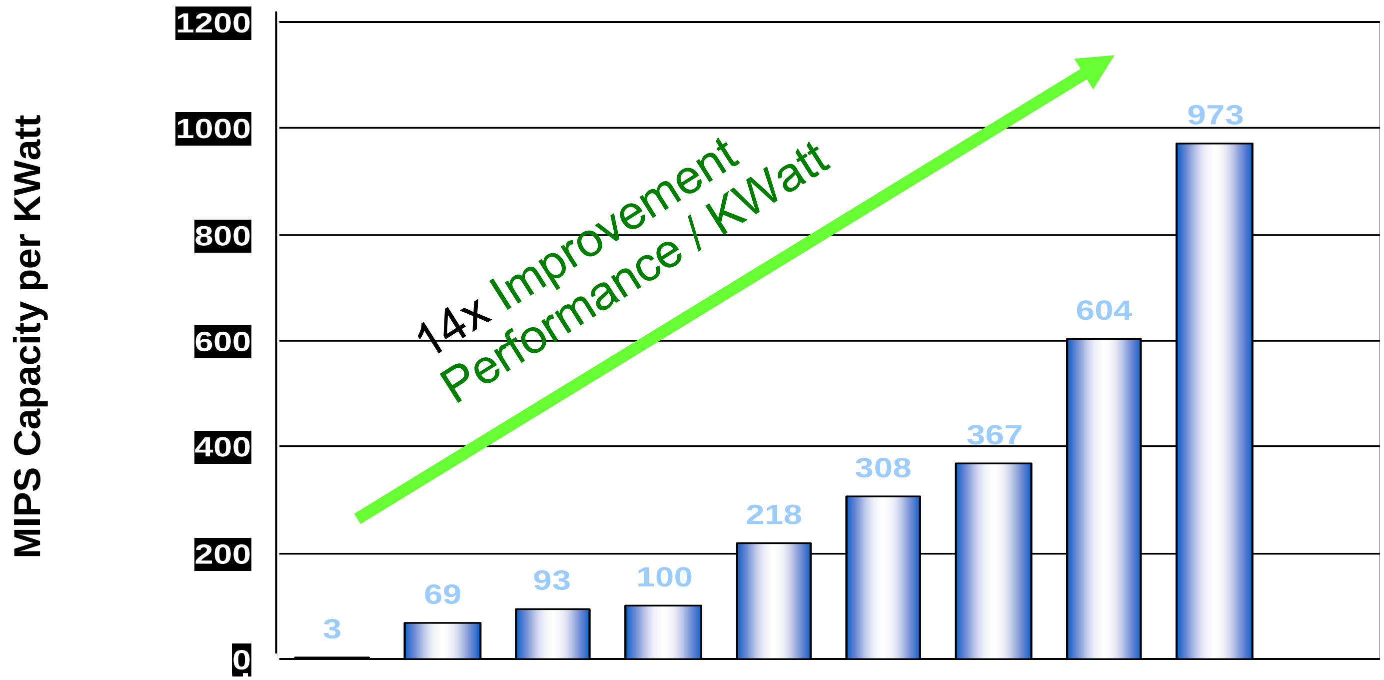

13 years (CMOS Evolution) G2 – z9 EC	Net Effect: CMOS G2 to z9 EC
Power increase 20%/year	**Performance increased by 192x**
Performance increase 50%/year	Performance/KWatt increased by 14x
Power density increase 17%/year	**Performance/sq ft increased by 4,000x**

When consolidating low utilization Linux on Intel servers, the System z9 Mainframe's ability to provide high utilization may help to reduce both power and facility costs

Power and Space Consumption

System z9 EC

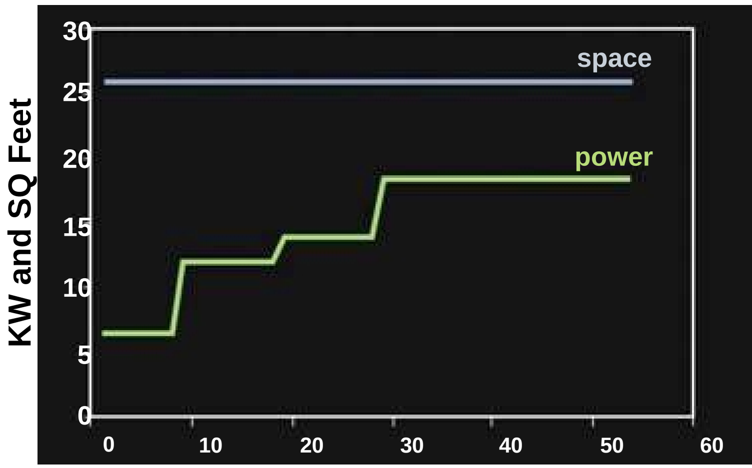

Processors

In a consolidation, the System z9 EC may provide up to 4 times the same work in the same space and may provide up to *12 times the work for the same power consumption*

Linux on Intel

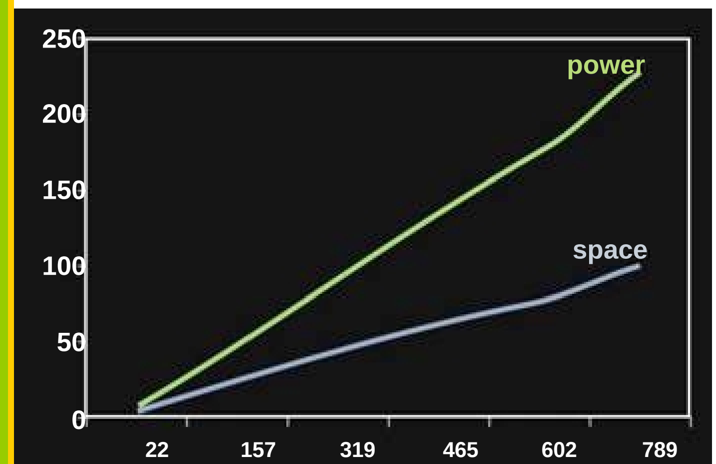

Processors

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 z9 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 * Comparison made versus SUN X2100 Opteron servers.

System Attributes to Consider

- System Design Characteristics
  - Does the CEC do all the work?
    - The roll of system assist processors
  - How is data moved around the system?
    - How efficiently, and at what speeds?
  - How automated is the systems management?
    - Minimize operator intervention
  - How automated is the workload management?
    - Based on internal standards, or business goals?
  - How is RAS implemented?
    - Fail over or DON'T FAIL
  - What is the security level supported?
    - Industry Standard (ex. Kerberos)
    - Certification Levels
    - Partition isolation
  - What is the "total system implementation'?
    - How integrated is the HW/SW stack

An approach to categorizing workloads

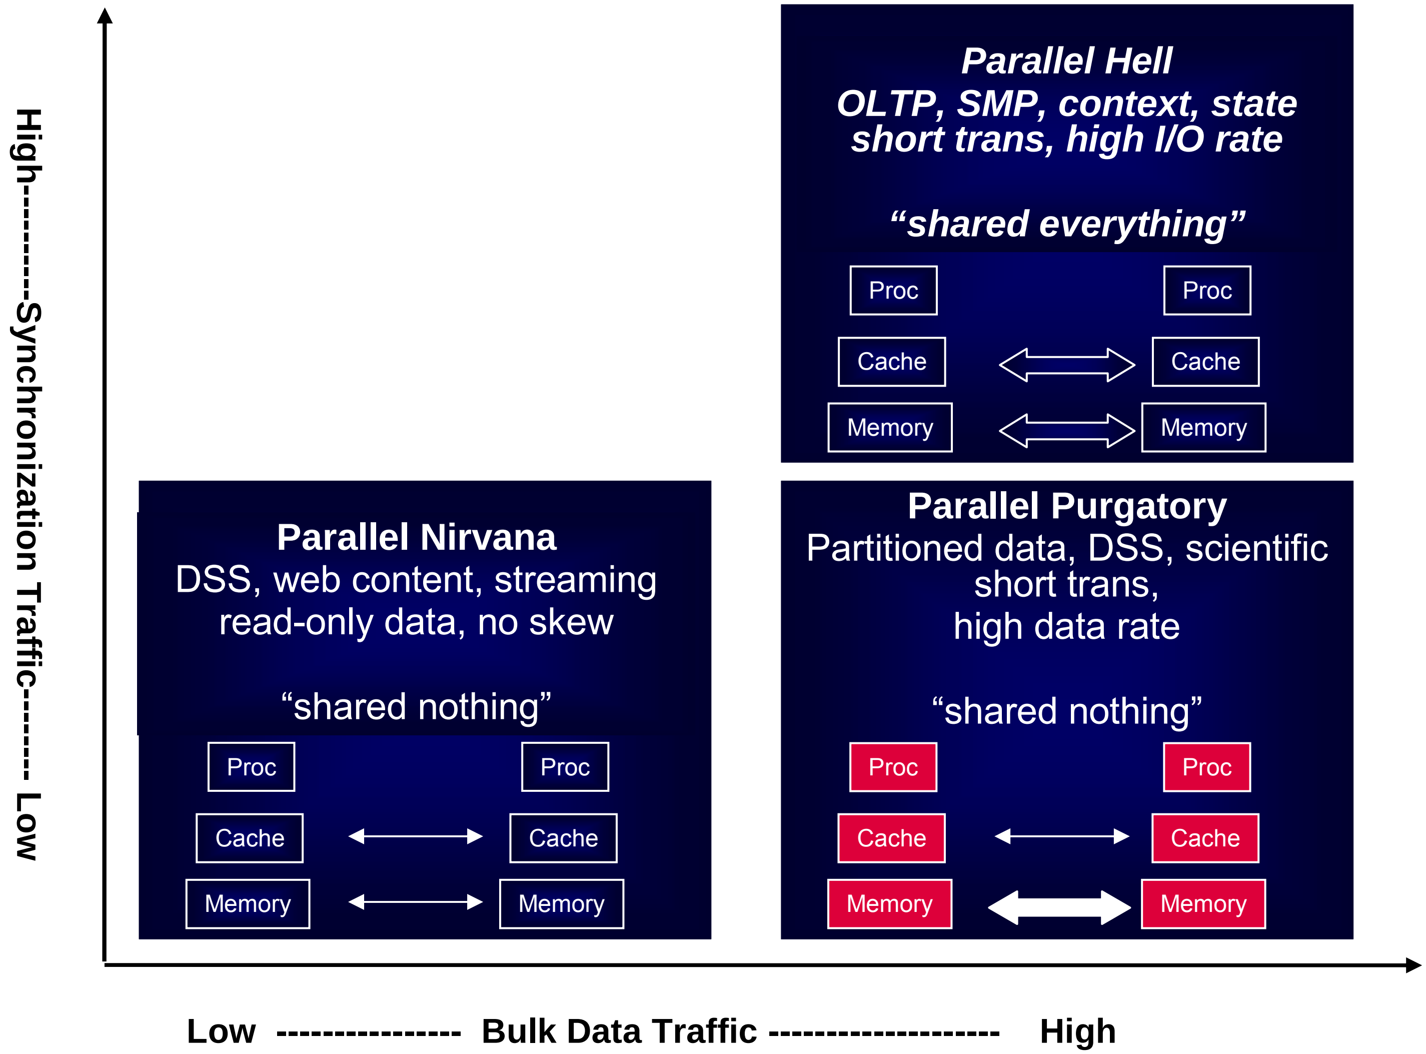

Example of mapping Pfister's quadrants based on how the workload uses the underlying server architecture

From: In Search of Clusters, The ongoing battle in lowly parallel computing by Greg Pfister, p461

Characterize and Profile the Workloads

These categories give a starting point to how the workload uses the server resources

I/O Intensive

Transactional - Response Time Critical

Data Intensive - I/O Critical

CPU Intensive

1. **I/O intensive workload -**
   - low to high elapsed time
   - low to moderate CPU utilization
   - high I/O count or rate
2. **Transactional –**
   - response time oriented
   - small “unit of work”, text based application
   - low to high elapsed time - transaction monitor?
   - moderate CPU time
   - moderate to high criticality or short RTO and RPO
3. **Data Intense -**
   - I/O critical
   - same criteria as above plus
   - data created by <u>other</u> applications accessed locally or through middleware (ODBC,JDBC, RMI, etc)
4. **CPU Intensive -**
   - low to moderate elapsed time
   - high CPU utilization
   - low I/O count or rate

Mapping of Platforms to Pfister's quadrants

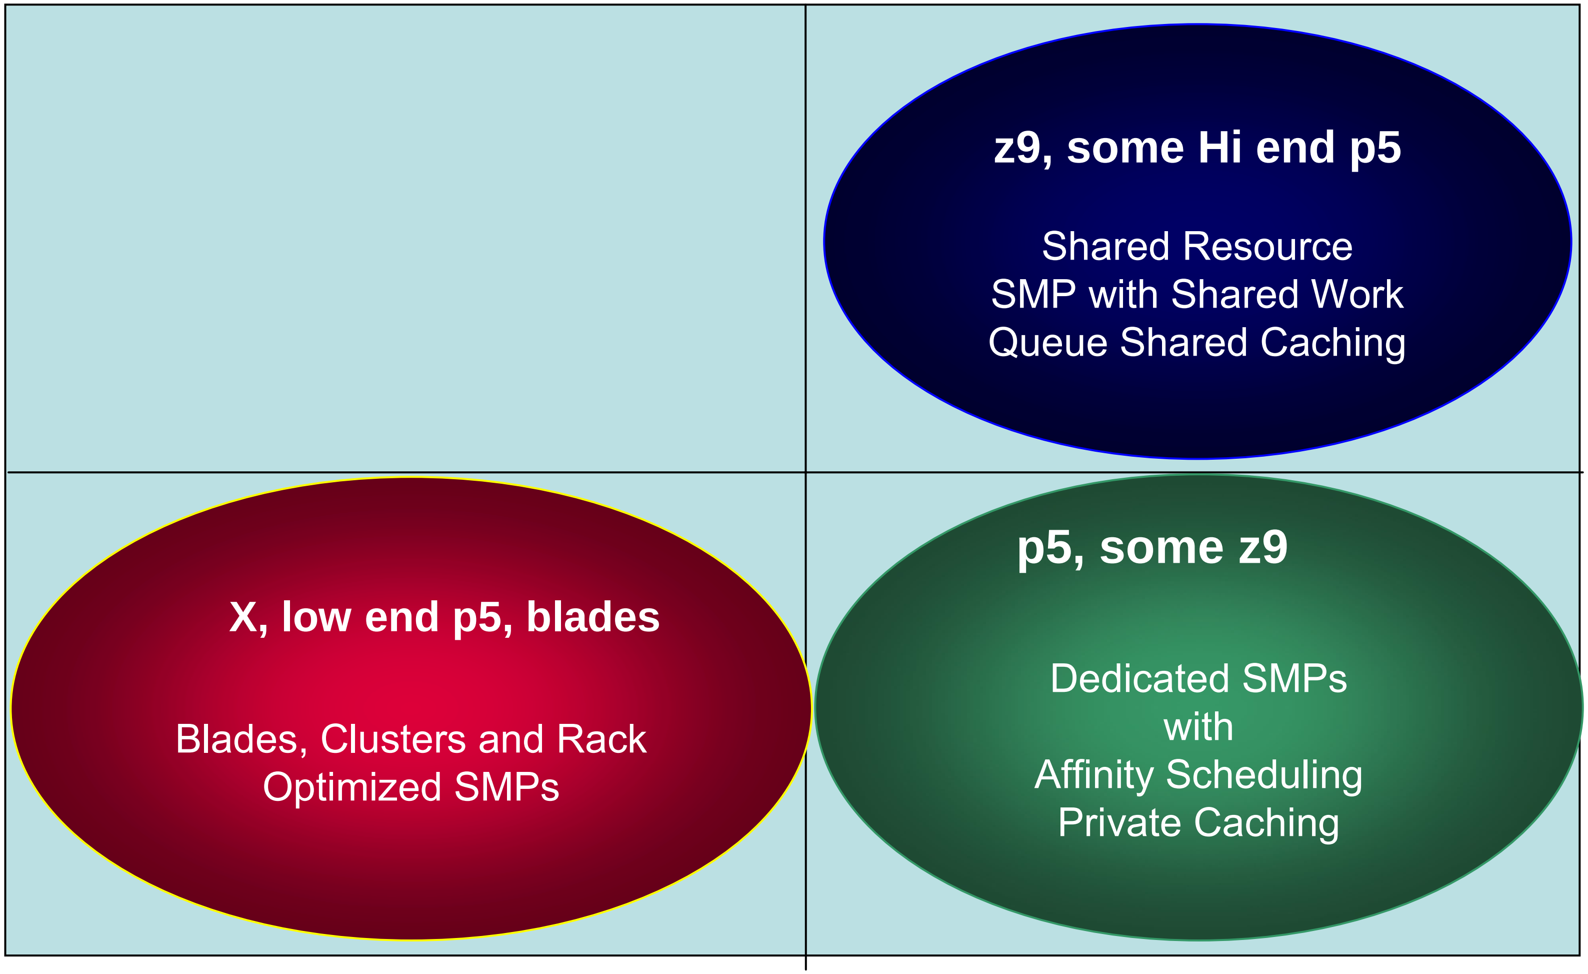

Example of mapping Pfister's quadrants to category of server

The System z Design

- System Utilization maximized through sophisticated “shared everything” architecture
- System Efficiency optimizes resource consumption such as power, space and people by sophisticated automation
  - Virtualization
  - Systems Management
- Cost of outages avoided through high System Availability
- Cost of Security breaches minimized through sophisticated encryption, RACF, TCPIP sensors

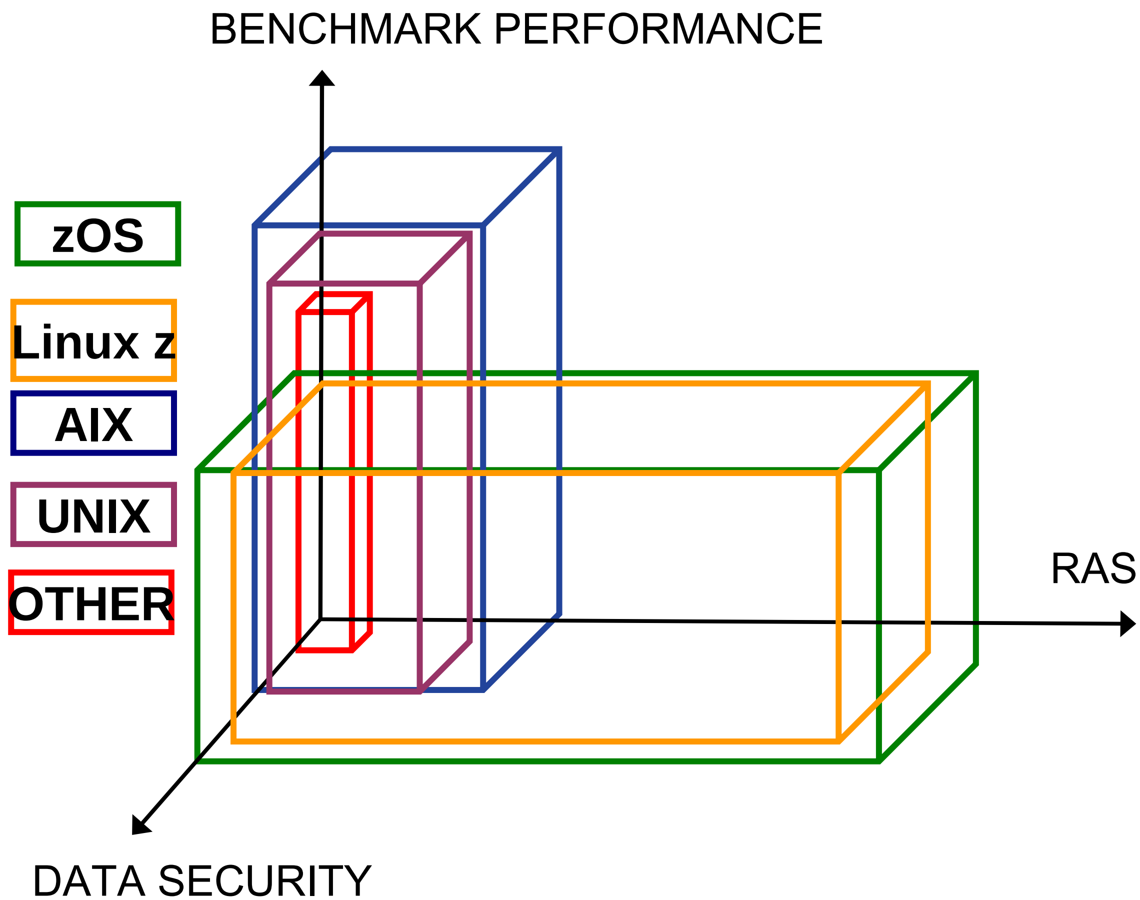

Comparative Implementations

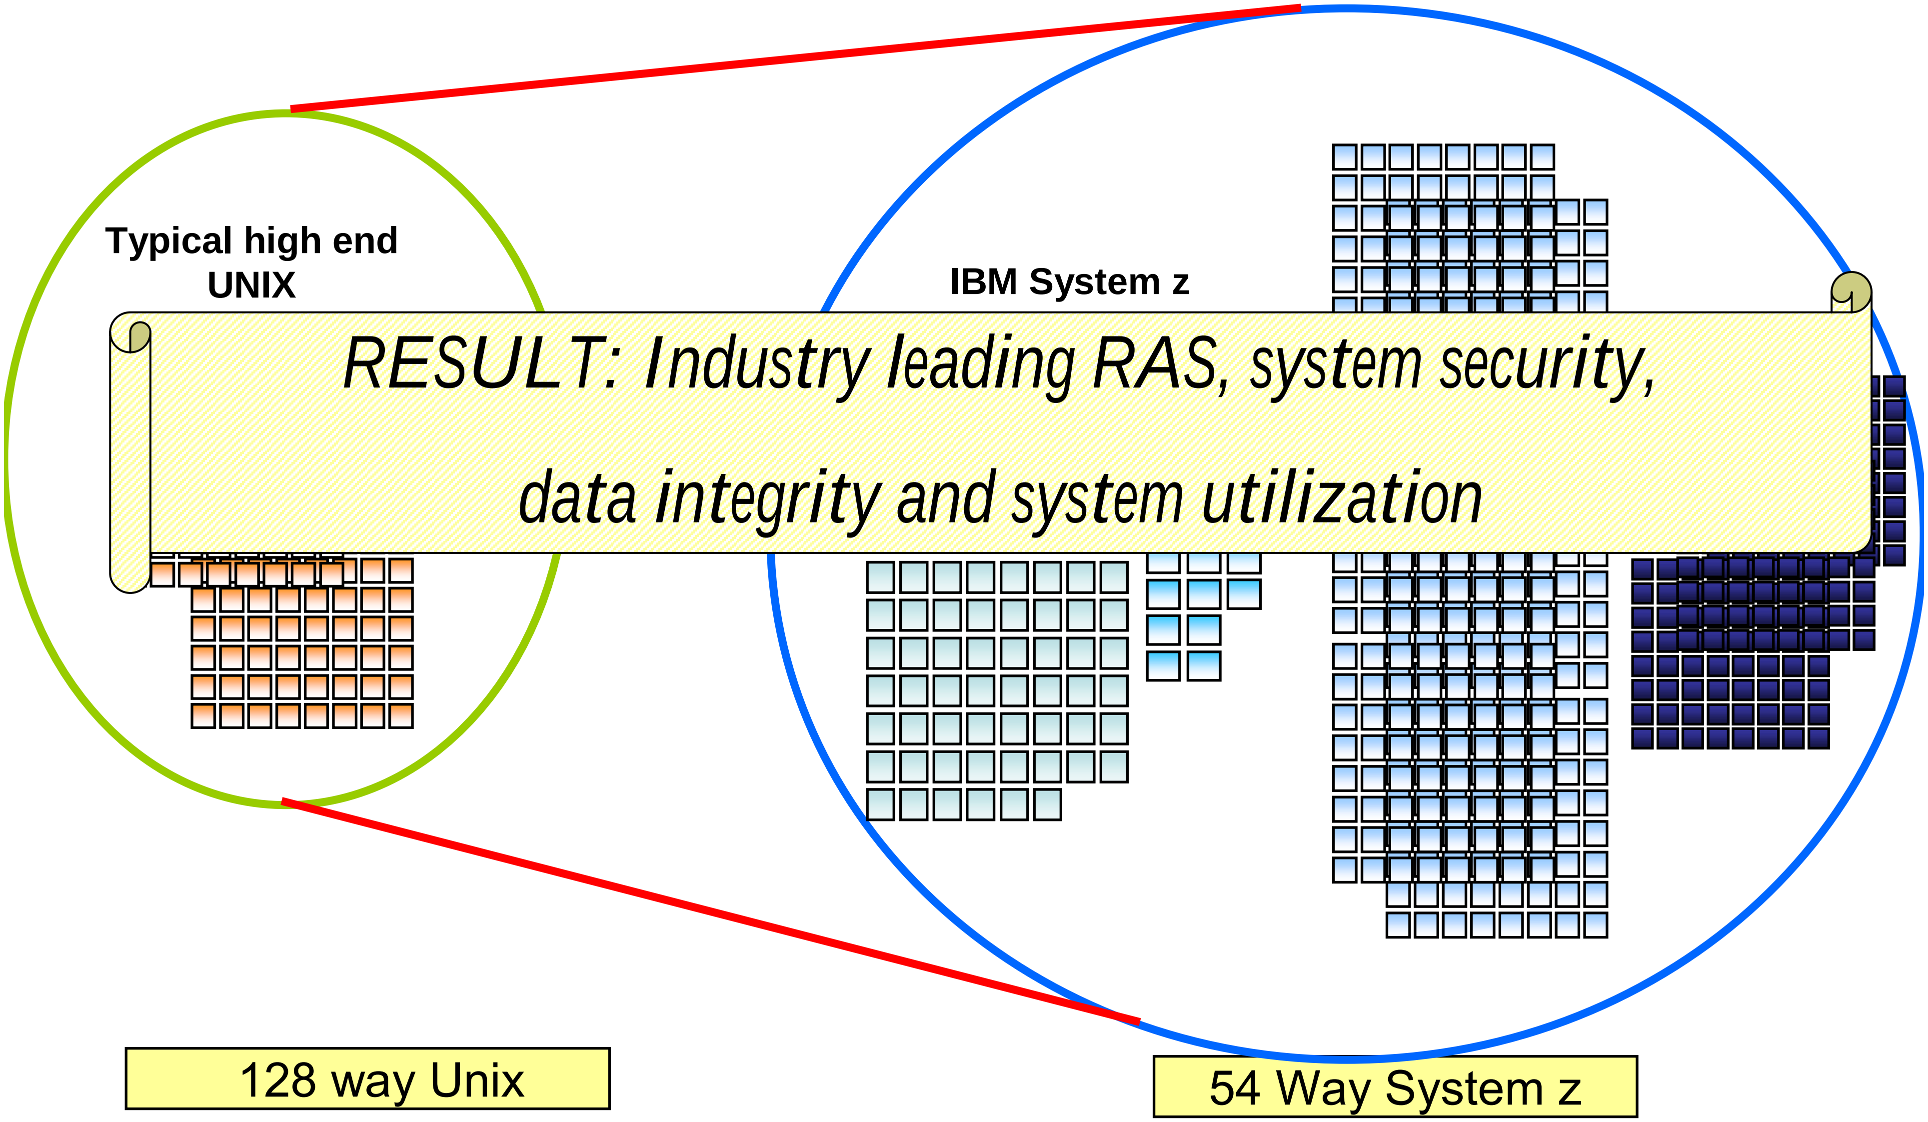

The IBM System z – the natural hub

…..the mainframe's capabilities, its time proven qualities of service and the tremendous collection of key business assets deployed there make System z the natural hub for a SOA enterprise. As such the WebSphere brand will continue to invest in System z as a premier platform for the WebSphere portfolio. Additionally, we will ensure that through our enablement of CICS TS and associated tooling, customers can continue to leverage their assets as part of a SOA implementation. ***System z's unique technical and business capabilities allow System z to stake a claim as the platform to serve as the hub of an enterprise for collaborative computing and SOA.***

Tom Rosamilia - General Manager of Application & Integration Middleware in IBM SWG

Trademarks

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

CICS*
DB2*
DB2 Universal Database
DirMaint
ESCON*
FICON*
GDPS*
HiperSockets
HyperSwap

IBM*
IBM eServer
IBM logo*
IBM System z9-109*
IMS
NetView*
OMEGAMON*
On Demand Business logo
Parallel Sysplex*

RACF*
S/390*
System z9-109
Tivoli*
TotalStorage*
VSE/ESA
VTAM*
WebSphere*
z/Architecture

z/OS*
z/VM*
z/VSE
System z*
System z Entry License Charge

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
Linux is a trademark of Linus Torvalds in the United States and other countries..
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft is a registered trademark of Microsoft Corporation in the United States and other countries.

* 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.