



# Advanced Technology Drives Differentiation

---

## *Power Architecture*

**Rick A. Kearns**  
**Advanced Solutions Consultant**  
**Oracle - IBM Enterprise Systems**  
**[rkearns@us.ibm.com](mailto:rkearns@us.ibm.com)**



## Customer Requirements Drive Technology Choice

**Number of users/applications accessing database**  
**Size and activity characteristics of database**  
**Potential for growth**

**Consolidation and Virtualization**  
**Reliability (SLA)**  
**Complexity**  
**Green Initiatives**

**Solution cost vs. component cost**

**Preference and Familiarity**

**There is no "One Size Fits All" Solution**

# Architectures Provide Different Capabilities



**System z™**

Throughput  
Quality of Service  
Resource Utilization  
System Virtualization  
SW License Reduction



**POWER™**

Work / Resource  
Scalability  
Throughput  
Resource Virtualization  
SW License Reduction



**System x™**

Standardization  
X86 Performance  
X86 Scalability  
VMWare  
Lowest HW Cost



**System Storage™**



# Primary Server Workloads 2009

## 18 Workloads in 7 Workload Categories


- Business Processing
  - ERP
  - CRM
  - OLTP
  - Batch
- Decision Support
  - Data Warehousing/Mart
  - Data Analysis/Mining
- Collaborative
  - E-Mail
  - Workgroup
- IT Infrastructure
  - File & Print
  - Networking
  - Proxy/Caching
  - Security
  - Systems Management
  - Web Infrastructure
    - Streaming Media
    - Web Serving
- Application Development
- Industrial R&D
- Other

YTY Revenue Growth	HP Fiscal 2Q09	IBM 1Q09
Total Company	-19% (HP+EDS)	-11%
Services	-15% (HP+EDS)	-10% (GTS)
Software	-15%	-1% (Tivoli), +9% (Rational)
Total Systems	-28%	-22%
Unix Servers	-29%	-2%
X86 Servers	-29%	-27%
Storage	-22%	-20%

# Power Values....Related To Customer Requirements

## Availability

- Low Resource Count
  - Self Monitoring
  - Self Correcting
- O/S & HW Integration
  - Hot Add
  - Hot Replace
- Workload Mobility
  - Planned Outage
- New Workloads



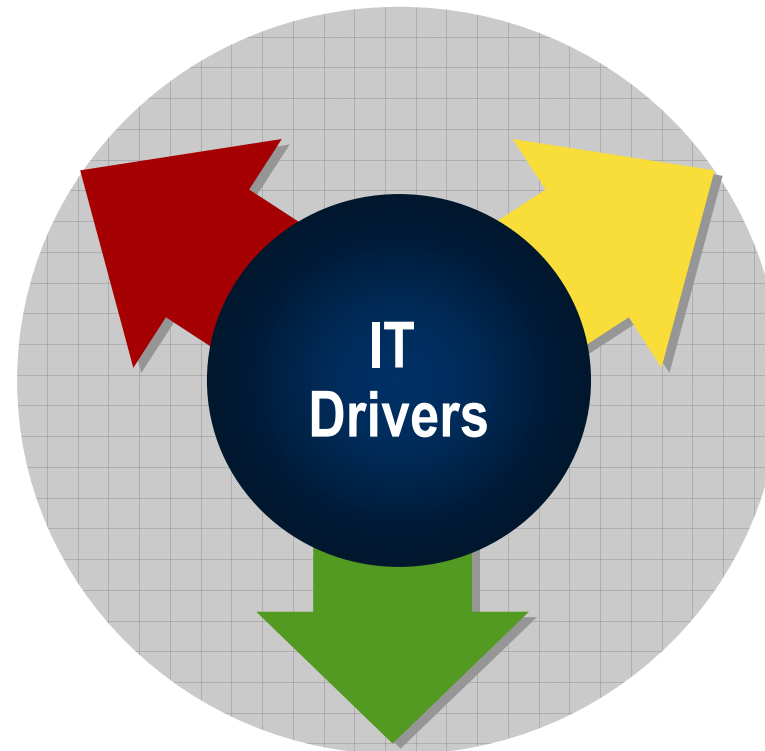
## Cost Reduction

- Performance / Resource
  - Utilization
  - Throughput
  - Efficiency
  - Near Linear Scalability
- Environmental
- Human Capital
  - Management
- Effect on HW & SW \$\$
  - TCO vs. TCA

## Business Flexibility/Responsiveness

- Dynamic Resource Allocation
- Virtual vs. Physical

# Power Values....Related To Customer Solutions



## Cost Reduction

- **Performance / Resource**
  - Utilization
  - Throughput
  - Efficiency
  - Near Linear Scalability
- **Environmental**
- **Human Capital**
  - Management
- **Effect on HW & SW \$\$**
  - TCO vs. TCA

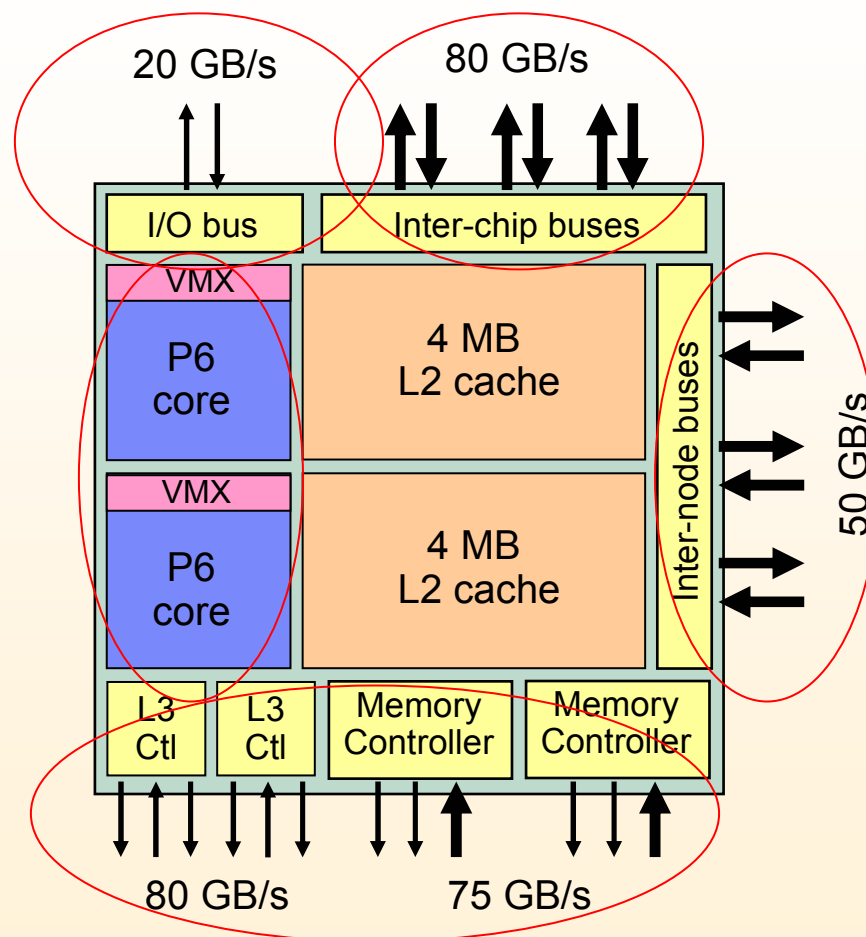
# The IBM POWER6 Chip

## ■ Innovative design

- New processor core technology
- In-order execution for high performance
- Tailored transistors for low power consumption
- Delivers more throughput and better response times
- Hardware decimal floating point
- Integrated VMX (AltiVec) unit
- Customizable configuration

## ■ Mainframe-inspired RAS



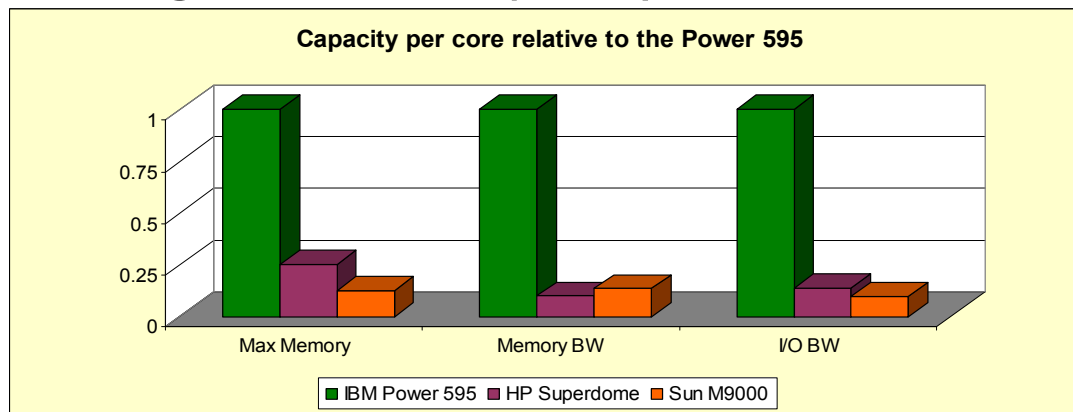
- Processor recovery unit
- Memory protect keys





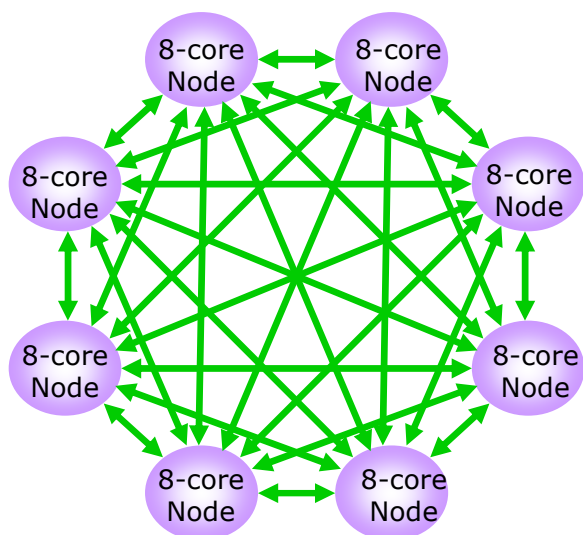
# IBM Power Servers

- *designed with capacity for consolidation*



- Shared Everything Design
- Flexible System Bandwidth
- Bandwidth Scales with Resources
- Enables High Utilization
- Supports Sustained Throughput

## IBM Distributed Switch vs. Cell Board Architecture



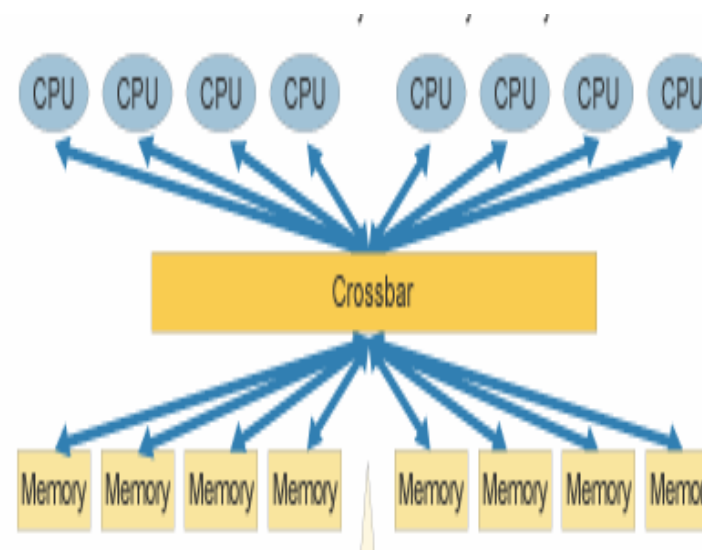
**Up To 32 socket**

**2.5GHz interconnect**

**Uniform memory Access**

**Point to Point connections**

### Sun/Fujitsu SPARC64 M-Class



**530MHz interconnect**  
**Long Memory Latency through Crossbar**

**(Sun E25K is 150MHz)**

# POWER6 EnergyScale™ Functions

## Data collection and reporting

- Power usage data
- System temperature

## Standby state for inactive resources

- Processor nap mode
- Memory chip power-down

## Throttling for rapid power savings

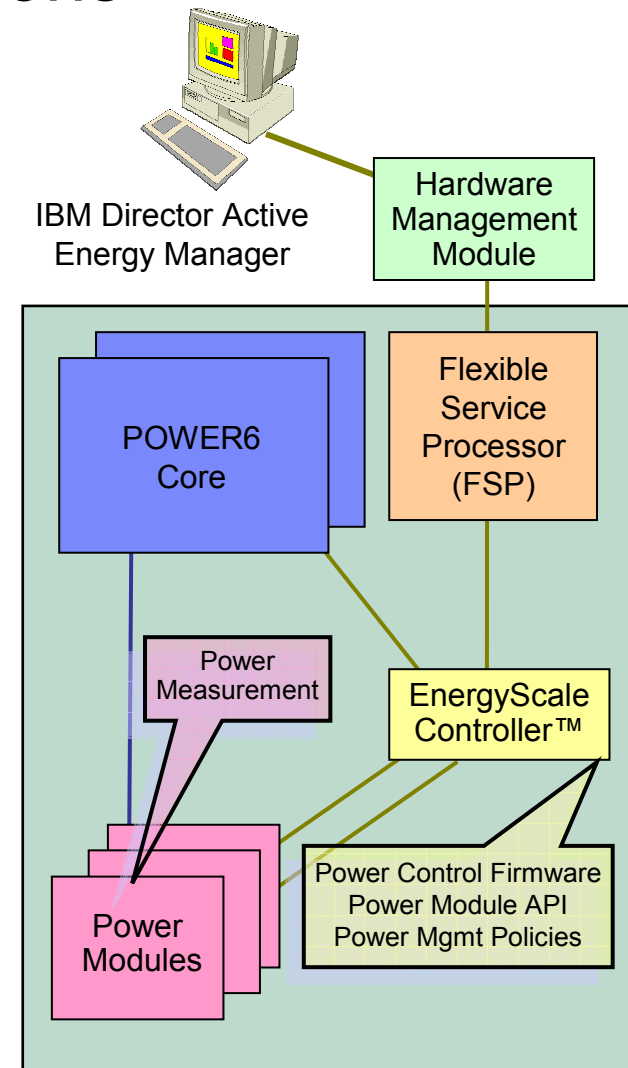
- Inject dead cycles into processor cores
- Reduce rate of memory accesses

## Static Power Save

- Drop core frequency to pre-defined level

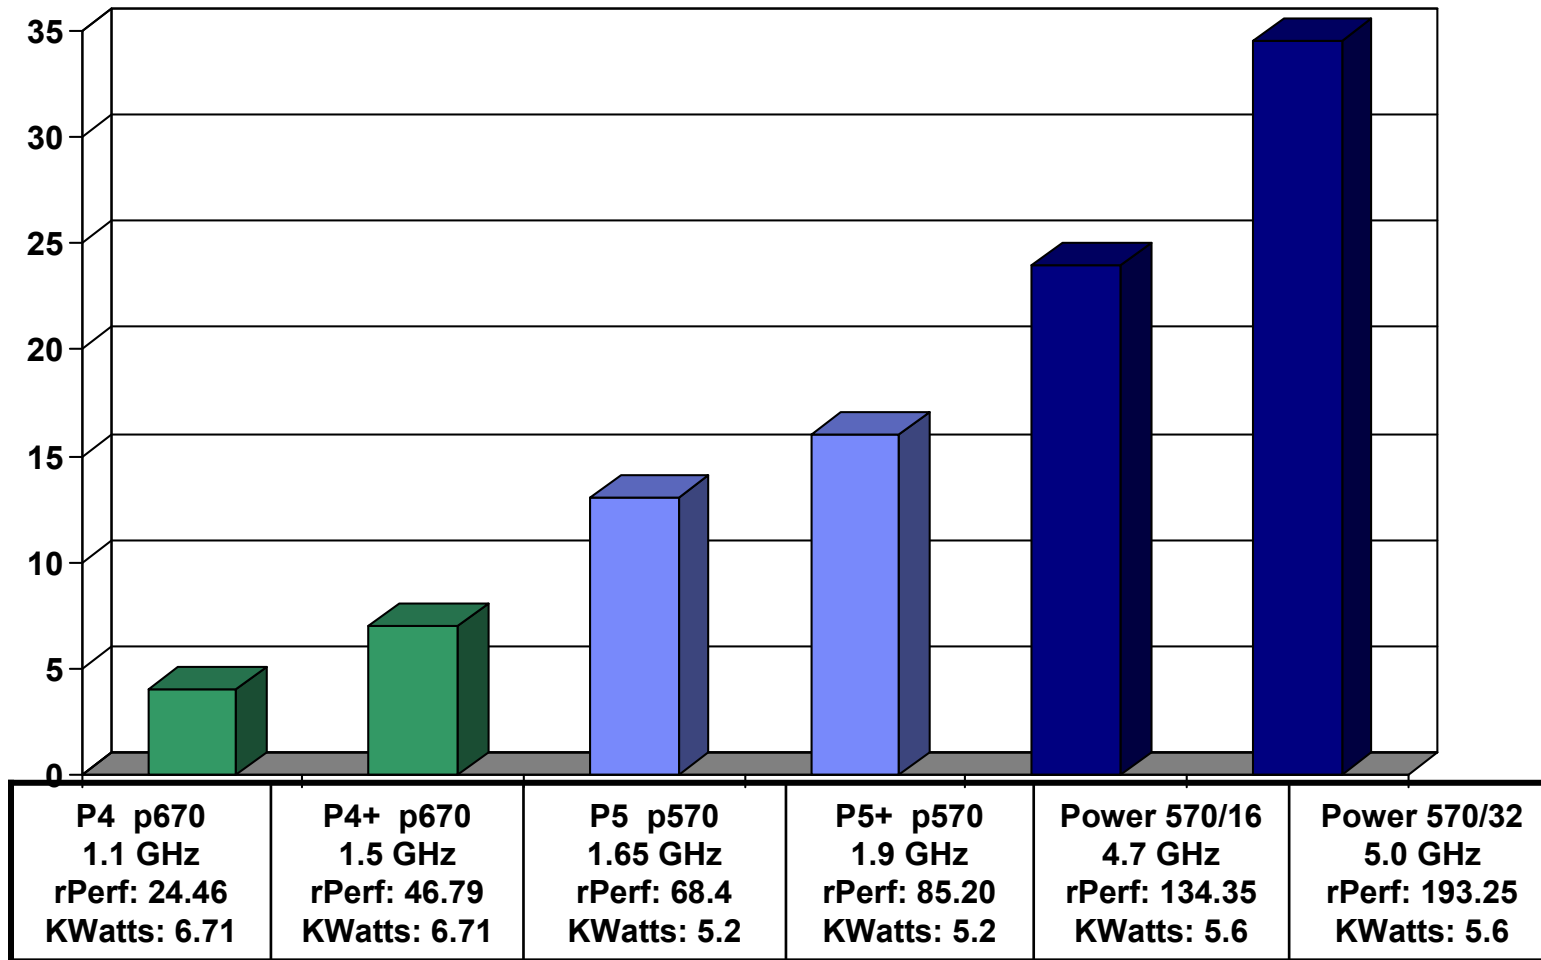
## Management policies

- Implemented by firmware, activated by HMC or Systems Director

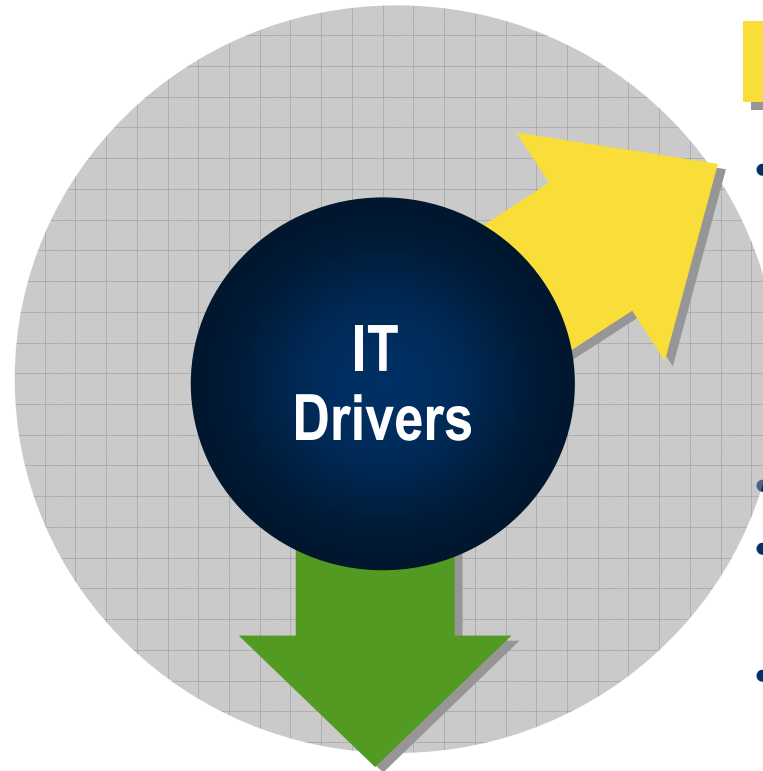


# Performance per KWatt

rPerf per KWatt



# Power Values....Related To Customer Solutions



## Cost Reduction

- Performance / Resource
  - Utilization
  - Throughput
  - Efficiency
  - Near Linear Scalability
- Environmental
- Human Capital
  - Management
- Effect on HW & SW \$\$
  - TCO vs. TCA

## Business Flexibility/Responsiveness

- Dynamic Resource Allocation
- Virtual vs. Physical



# ■ *Optimizing IT with industrial strength virtualization*

**64% of Power Systems shipped with PowerVM in 2Q08**



# Industry Leading Functionality



**Virtualization**  
 Workload Partitions  
 Workload Partitions Manager™  
 Multiple Shared Pools  
 Shared Dedicated Processors

**Security**  
 Role Based Access Control  
 Trusted AIX  
 Encrypting Filesystem  
 AIX Security Expert

**Continuous Availability**  
 Concurrent AIX Updates  
 Kernel Memory Protection keys  
 Hot Add - Cold Replace  
 Live Application Mobility  
 Live Partition Mobility

**Manageability**  
 System Director for AIX  
 WPAR manageability  
 Integrated Filesystem Snapshot  
 Graphical Install

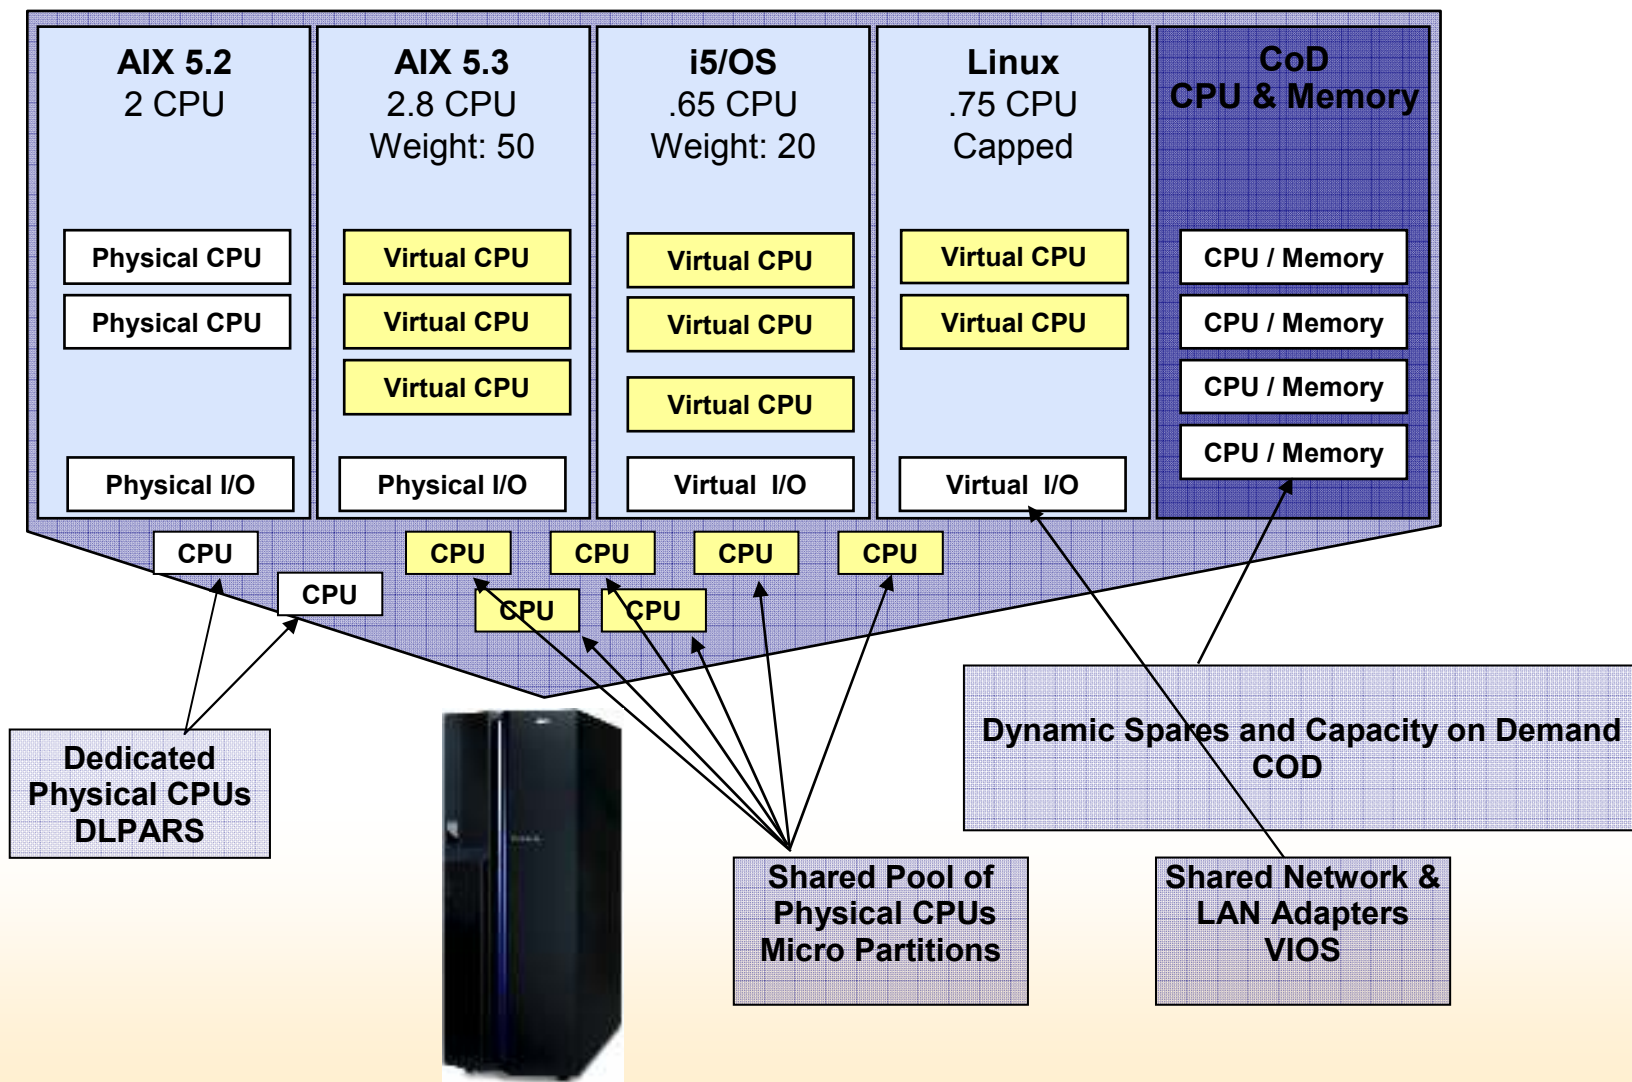


## Key Points to Remember about Virtualization

- **Virtualization is not about higher performance**
- **Higher utilization (virtualization) is directly proportional to available server throughput**
- **Virtualization will not overcome a weak architecture**



# POWER: CPU - Memory - I/O Virtualization





## Dynamically optimize memory utilization across virtual images

### PowerVM Active Memory Sharing



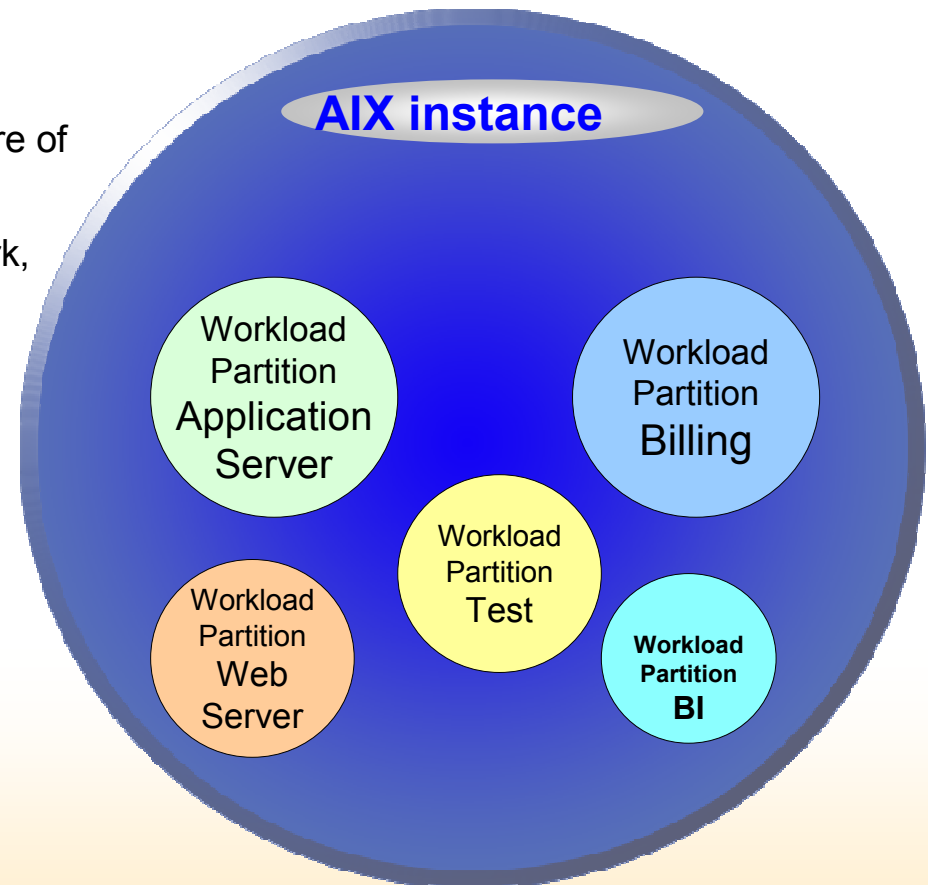
- Dynamically adjusts memory available on a physical system for multiple virtual images based on their workload activity levels:
  - Supports over-commitment of logical memory
  - Different workload peaks due to time zones
  - Mixed workloads with different time of day peaks (e.g. CRM by day, batch at night)
  - Ideal for highly-consolidated workloads with low or sporadic memory requirements
  
- Available with PowerVM Enterprise Edition
  - Supports AIX, IBM i and Linux workloads



# AIX V6.1 Workload Partitions

Improved administrative efficiency by reducing the number of AIX images to maintain


- **Software partitioned system capacity**
  - Each Workload Partition obtains a regulated share of system resources
  - Each Workload Partition can have unique network, filesystems and security
- **Two types of Workload Partitions**
  - System Partitions
  - Application Partitions
- **Separate administrative control**
  - Each System Workload partition is a separate administrative and security domain
- **Shared system resources**
  - Operating System, I/O, Processor, Memory




# Typical Scale-out Approach

- **Single workload on a single system**
  - Average Utilization: 20.7%
  - Peak: 79%
- **Eight separate workloads on eight identical systems**
  - Average utilization is still 20.7%
  - Peak is still 79%


Single Application Server




Single Application Server




Single Application Server




Single Application Server




Single Application Server




Single Application Server




Single Application Server



Single Application Server



Single Application Server



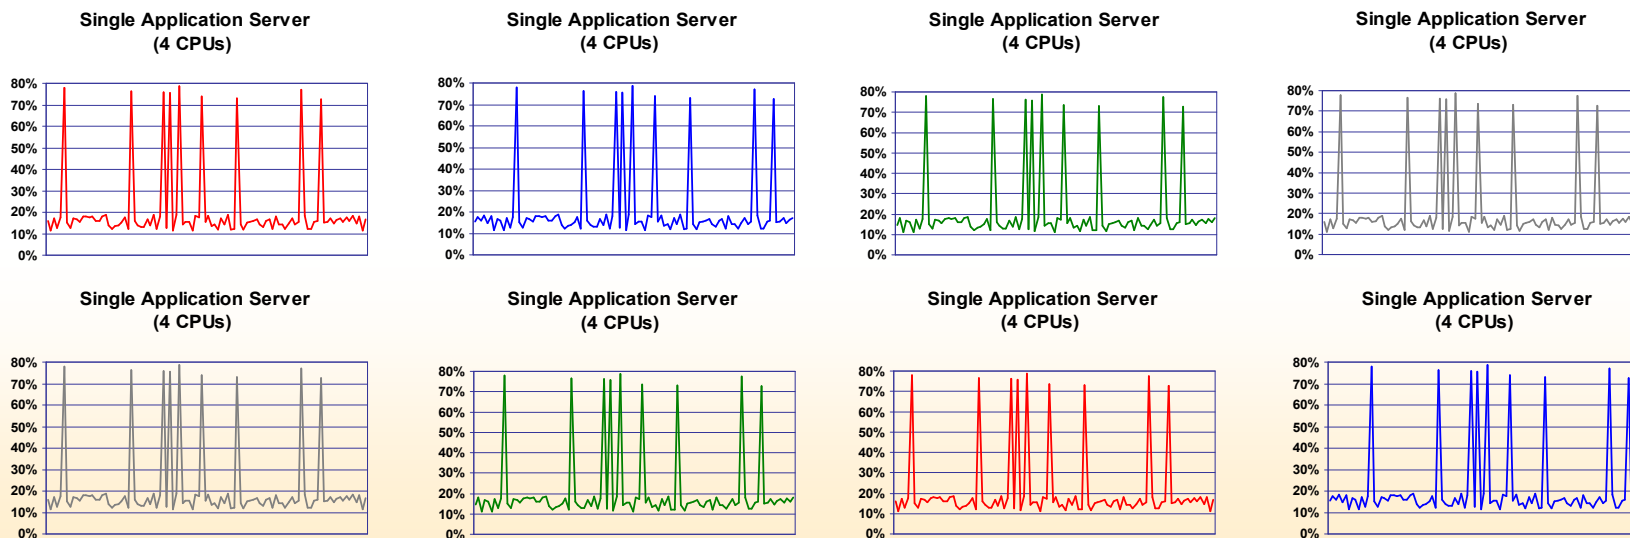

**Result is 8X the hardware, software, maintenance, floor space, and energy that you pay for, is wasted**

# Same Scenario with Virtualized Resource Consolidation

- **Eight separate workloads on eight identical systems**
  - Average utilization is 20.7%
  - Peak is 79%
- **Eight separate workloads on one system\***
  - Average utilization is 39%
  - Peak is 76%

**\* 32 CPUs reduced to 16 CPUs (2 to 1)**

8 to 1 Systems Consolidation  
(16 CPUs)



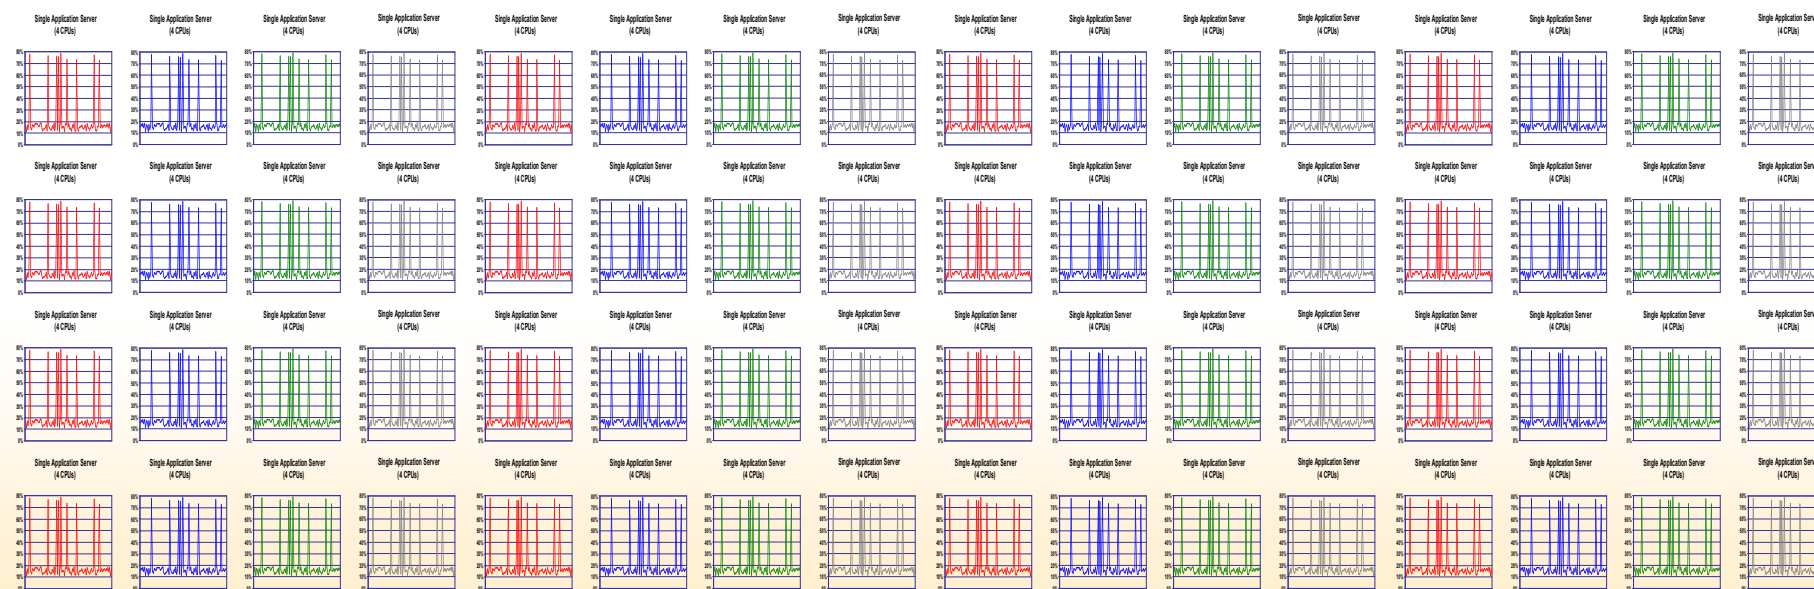

# Very Large Scenario with Virtualized Server Consolidation

- **64 separate workloads on 64 identical systems**
  - Average utilization is 20.7%
  - Peak is 79%
- **64 separate workloads on one system\***
  - Average utilization is 61%
  - Peak is 78%



**256 CPUs reduced to 64 CPUs (3.5 to 1)**

**64 to 1 Systems Consolidation (64CPUs)**





# Reduce Server Acquisition over \$400K & Save up to 91% of the energy use!


*By consolidating on Power 550 Express instead of Sun SPARC Enterprise M4000*

*-- Use up to 1 / 7th the rack space*

*-- Up to 90% fewer cores*

From:

- 10 V490 servers in two racks (20% utilized)
- Maximum energy of 17,500 WATTS




To:



- **One Power 550 Express server**
- **List Price of \$161,596**
- **Maximum energy requirement of 1,500 Watts**

or




Sun M4000 supports no more than 2 dynamic domains per system and would require Five M4000 servers to consolidate 10 V490s

- **Five M4000 servers**
- **List price of \$571,900**
- **Maximum energy requirement of 10,080 Watts**

# Power Values....Related To Customer Solutions

## Availability

- **Low Resource Count**
  - Self Monitoring
  - Self Correcting
- **O/S & HW Integration**
  - Hot Add
  - Hot Replace
- **Workload Mobility**
  - Planned Outage
- **New Workloads**



## Cost Reduction

- **Performance / Resource**
  - Utilization
  - Throughput
  - Efficiency
  - Near Linear Scalability
- **Environmental**
- **Human Capital**
  - Management
- **Effect on HW & SW \$\$**
  - TCO vs. TCA

## Business Flexibility/Responsiveness

- **Dynamic Resource Allocation**
- **Virtual vs. Physical**



# Industry Leading Functionality



## Virtualization

- Workload Partitions
- Workload Partitions Manager™
- Multiple Shared Pools
- Shared Dedicated Processors

## Security

- Role Based Access Control
- Trusted AIX
- Encrypting Filesystem
- AIX Security Expert

## Continuous Availability

- Concurrent AIX Updates
- Kernel Memory Protection keys
- Hot Add - Cold Replace
- Live Application Mobility
- Live Partition Mobility

## Manageability


- System Director for AIX
- WPAR manageability
- Integrated Filesystem Snapshot
- Graphical Install



# OS Downtime Comparison Survey

400 participants in 27 countries

Hours



The Yankee Group "2007-2008 Global Server Operating Systems Reliability Survey" as quoted in "Windows Server: The New King of Downtime" by Mark Joseph Edwards at [www.windowsitpro.com/article/articleid/98475/windows-server-the-new-king-of-downtime.html](http://www.windowsitpro.com/article/articleid/98475/windows-server-the-new-king-of-downtime.html), March 5, 2008 and in <http://www.sunbeltsoftware.com/stu/Yankee-Group-2007-2008-Server-Reliability.pdf>

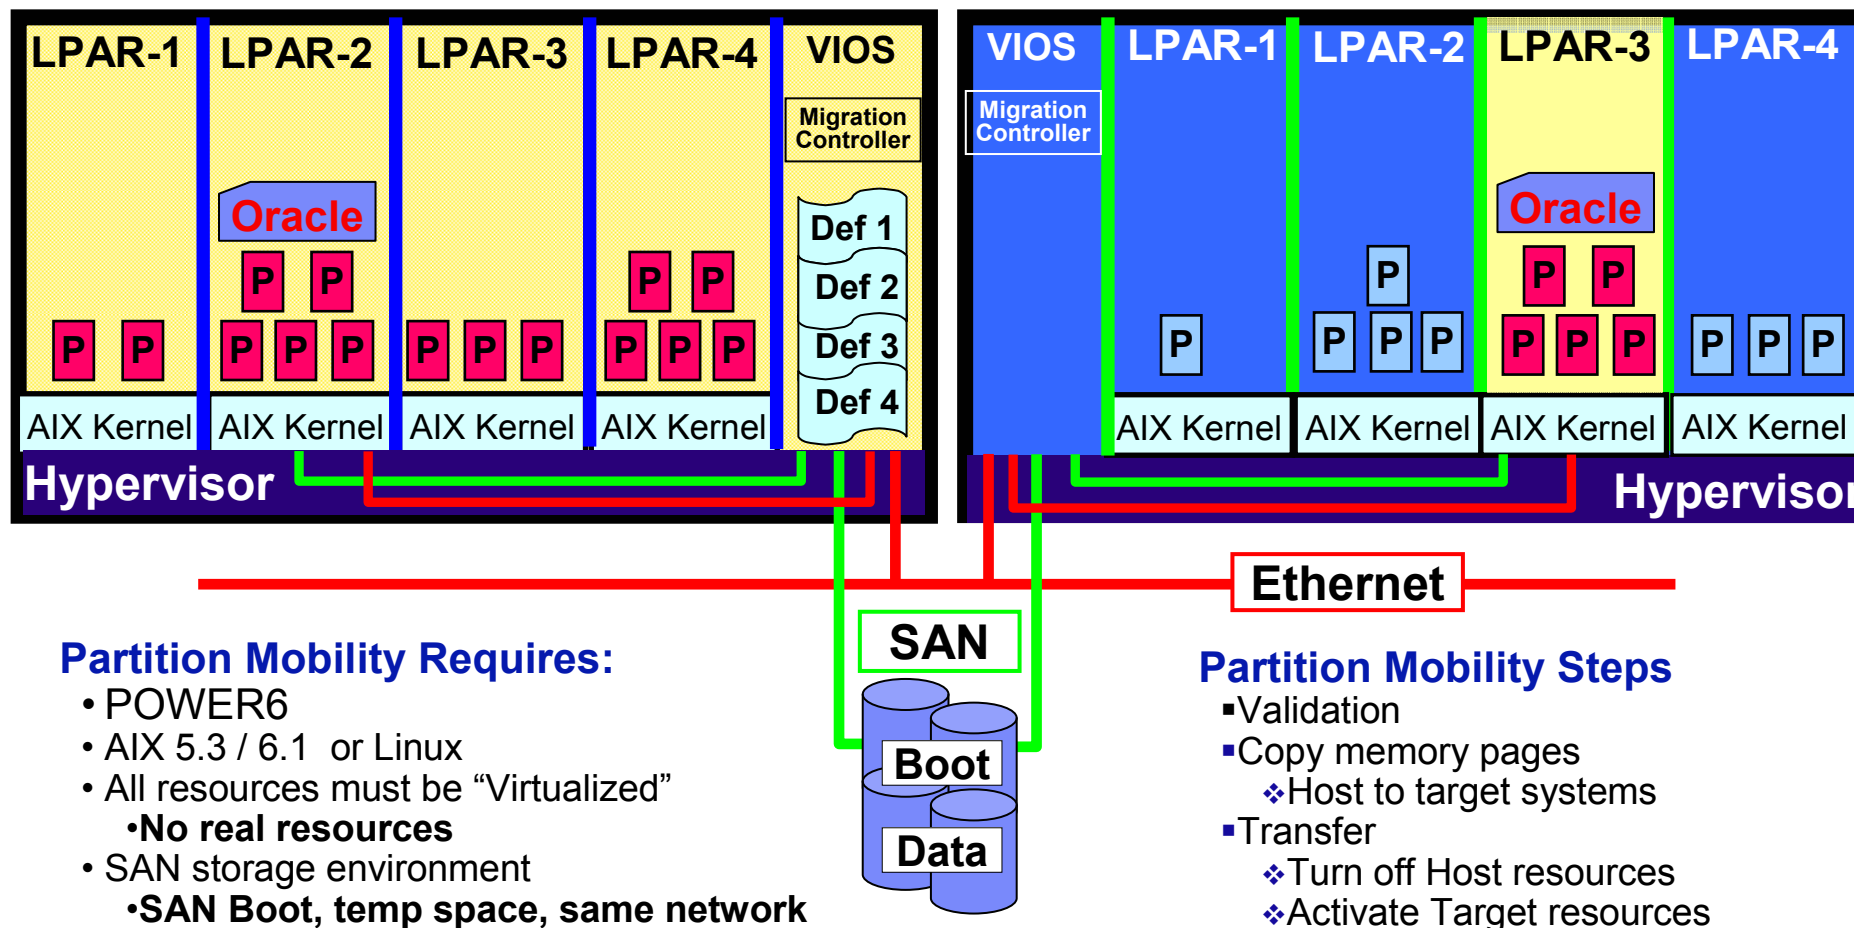
# Reliability by Design

- **Mainframe RAS Functionality for Unplanned Outages**
  - ▶ Extensive error detection and correction
  - ▶ Dynamic de-allocation of processors and cache
  - ▶ Hot plug, power, fans, adapters, disk, nodes
  - ▶ Concurrent maintenance of I/O drawers
  - ▶ **Concurrent firmware upgrades**
  - ▶ **Processor Retry and Recovery**
  - ▶ **Check Point Restart**
  - ▶ **Memory Protection Keys**
- **Elimination of Planned Outages**
  - ▶ **Hot Add - Cold Replace of Processors and Memory**
  - ▶ **Live Partition Mobility**
  - ▶ **Live Application Mobility**



# Live Partition Mobility – migrate active LPAR between systems

Reduce impact of planned outages, relocate workloads to enable growth, provision new technology with no disruption to service



The number of Oracle licenses needed does not change before and after the migration

## Summary of New Features for POWER6 and AIX 6.1

Legend

P6	AIX 6.1	P6 + AIX 6.1
----	---------	--------------

Processor	POWER4		POWER5		POWER6	
OS Level	AIX 5.3	AIX 6.1	AIX 5.3	AIX 6.1	AIX 5.3	AIX 6.1
Multiple Shared Processor Pools (up to 64)	No	No	No	No	Yes	Yes
Partition Mobility	No	No	No	No	Yes	Yes
Share Excess Dedicated Processor Cycles	No	No	No	No	Yes	Yes
Integrated Virtualized Ethernet	No	No	No	No	Yes	Yes
Workload Partitions and Application Mobility (WPARs)	No	Yes	No	Yes	No	Yes
Virtual Memory	No	No	No	No	No	Yes

# IBM Power Systems

## Consistency

- Binary compatibility
- Mainframe-inspired reliability
- Support for virtualisation
- AIX, Linux and IBM i OS

Complete flexibility for workload deployment

**Express**

JS12, JS22, JS23, JS43

520

550

560

**Enterprise**

570

595

575

Linux

IBM i for Business

AIX

# Not all Architectures are Equal

## Power Advantages...

- **More Performance per core and per system**
  - ▶ Results in smaller hardware configurations that deliver the same performance – hardware and facility cost savings
  - ▶ Saves on software costs when software is licensed by the core
- **Virtualization technology**
  - ▶ Allows one physical server to run multiple logical/virtual servers
  - ▶ Allows customer to chose when and how to use vertical and horizontal scaling
    - Without IBM Virtualization, the only choices were horizontal scaling or buying large systems that were under utilized in order to allow growth or to absorb capacity spikes
- **Increased Utilization with High Level of Availability**
  - ▶ Do more work with fewer processors
  - ▶ Reduce frequency of planned and unplanned system down time.
  - ▶ Saves on hardware and software costs
- **Roadmap**
  - ▶ We own our entire chip and system roadmap, development, and production
  - ▶ Our chips and systems are designed synergistically – development of each is done with the needs and requirements of the other taken into consideration
  - ▶ We deliver a vertically integrated solution that provides industry leading business value to our customers.

Q&A