

# IBM Dynamic Infrastructure FORUM

## Effective Workload Consolidation on Power





## Globally, systems and infrastructure are reaching a breaking point.

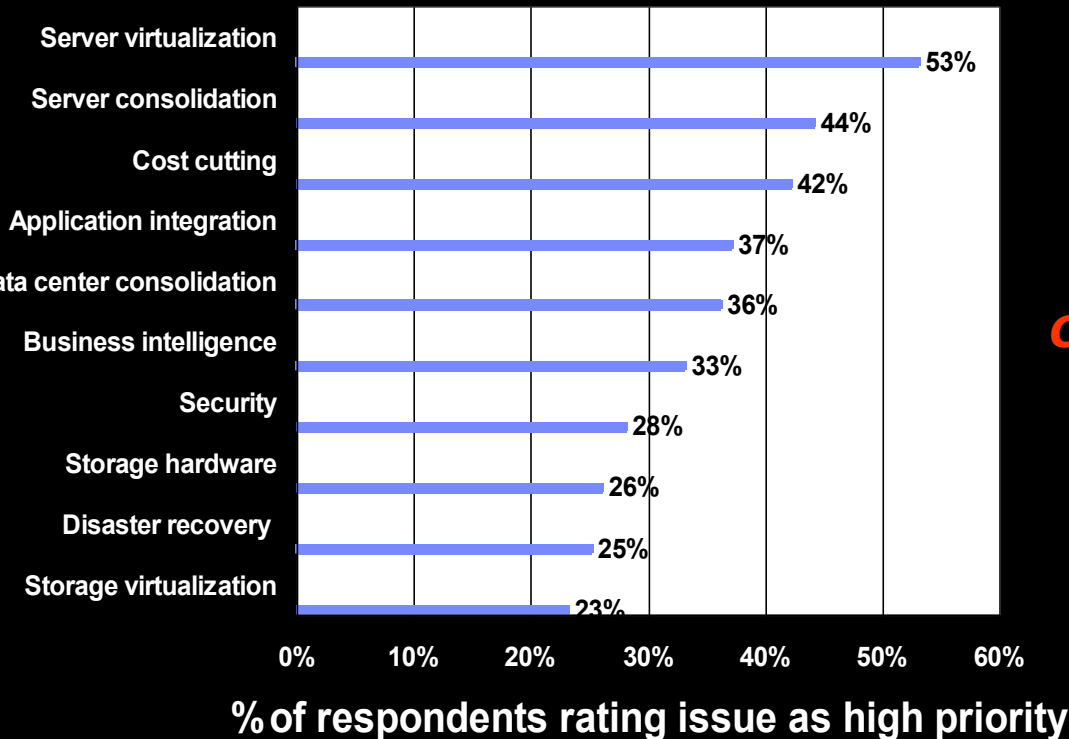
- Proliferation of servers and networking devices
- Excessive energy usage and heating problems
- Inadequate power and cooling infrastructure
- Data silos and data synchronization
- Expectations that “everything” is connected
- Linear staffing costs
- Skyrocketing software costs
- Unexplained outages



*Meanwhile, customer expectations, competitive pressures, regulatory requirements and fiscal pressures are increasing.*

# IT Leaders Need to do More With Less

*Operational issues can inhibit business growth and innovation*



**Overburdened IT staff**



**Complex service delivery**



**Strained IT budgets**



Source: Goldman Sachs Group IT Spending Survey, July 2008



# A dynamic infrastructure addresses today's challenges... *and tomorrow's opportunities*



## REDUCE COST

---

Not just containing operational cost and complexity, but achieving *breakthrough* productivity gains through virtualization, optimization, energy stewardship, and flexible sourcing.

## IMPROVE SERVICE

---

Not only ensuring high availability and quality of existing services, but also meeting customer expectations for real-time, dynamic access to innovative *new* services.

## MANAGE RISK

---

Not only addressing today's security, resiliency, and compliance challenges, but also preparing for the new risks posed by an even more *connected* and *collaborative* world.

# Consolidating with Enterprise Power Systems *enables clients to...*

## REDUCE COST

Server consolidation with shared resources enables high system utilization, which lowers the cost of ownership by reducing networking, energy, floor space, and software costs.

## IMPROVE SERVICE

Server consolidation improves service to clients by delivering flexible performance, dynamic provisioning and enabling clients to avoid disruption

## MANAGE RISK

Server consolidation manages IT risk by improving security, increasing business resiliency and simplifying operations.



***Power™ Systems and PowerVM™ are designed to deliver effective consolidation in the most demanding data centers***

## REDUCE COST by consolidating with Power Systems

### ❖ Resource sharing

- Sharing system resources through virtualized consolidation **reduces unused system overhead**
- Virtualized consolidated systems are evidenced by **high utilization rates**
- High utilization means **less hardware**

### ❖ Environmentally friendly

- **Less power and cooling** is required
- **Less floor space** is required

### ❖ Fiscally responsible

- Fewer processor cores drives **less software costs**
- Newer systems are **more reliable** and less costly to maintain than older systems
- Fewer systems translates to **reduced people costs**




## Typical small server utilization

***Typical UNIX or x86 serving running a single operating environment is 10 - 20% utilized***

Configuration planned for growth  
(20% unused?)

Configuration planned for peaks  
(50% unused?)

System waits for I/O and memory  
access even when it is working  
(20% unused?)



**Result is that 80% of the hardware, software, maintenance, floor space, and energy**

# Typical Scale-out Approach

Single workload on a single system

Assumed average utilization of 20%

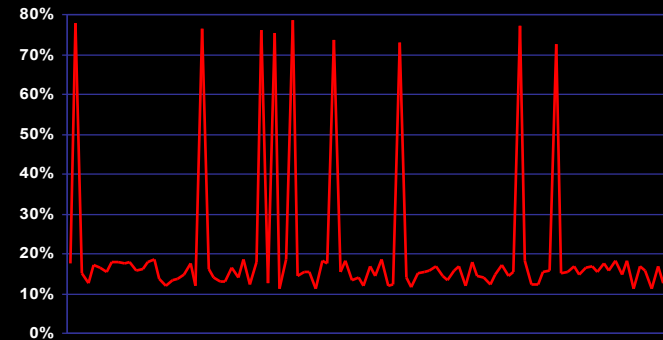
Assumed peak of 4X

Peaks are assumed to be random


Eight separate workloads on eight identical systems

Same assumptions


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, and floor space**



# Same Scenario with Physical Server Consolidation

Single workload on a single system


Average Utilization: 20%

8 separate workloads on one partitioned system

Average utilization is still 20%


8 times the hardware in one physical system

Single Application Server



Utilization remains the same due to no resource sharing

Partitioned Application Server with eight fixed partitions



**Result is 8X the hardware and software**

# Same Scenario with Virtualized Server Consolidation (Shared resources)

Eight separate workloads on eight identical systems


Average utilization is 20%

Eight separate workloads on one system\*

Average utilization is 39%


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

8 to 1 Systems Consolidation  
(16 cores)




Utilization increases


Single Application Server  
(4 cores)




Single Application Server  
(4 cores)




Single Application Server  
(4 cores)




Single Application Server  
(4 cores)




Single Application Server  
(4 cores)




Single Application Server  
(4 cores)



Single Application Server  
(4 cores)



Single Application Server  
(4 cores)



## The effect of sharing system resources

Average utilization per system almost doubles

20% to 39%

Peak usage actually drops

79% to 76%

Any single application now has access to more resource


Previous peaks, capped at 4 cores, can now go to 16 cores

Critical workloads can now be prioritized and enabled to run faster

Batch jobs, for instance, now run faster, due to the ability to access more capacity



**\*** 16 cores of unused capacity eliminated and applications run faster



# Larger Scenario with Virtualized Server Consolidation (Shared resources)

Sixteen separate workloads on sixteen identical systems

Average utilization is 20%

Peak is 79%



Sixteen separate workloads on one system\*

Average utilization is 48%

Peak is 78%

**\*** 64 cores reduced to 24 cores (2.65 to 1)

**16 to 1 Systems Consolidation  
(24 cores)**



\* Assumes independent workloads

© 2009 IBM Corporation

# Very Large Scenario with Virtualized Server Consolidation (Shared resources)

64 separate workloads on 64 identical systems

Average utilization is 20%

Peak is 79%

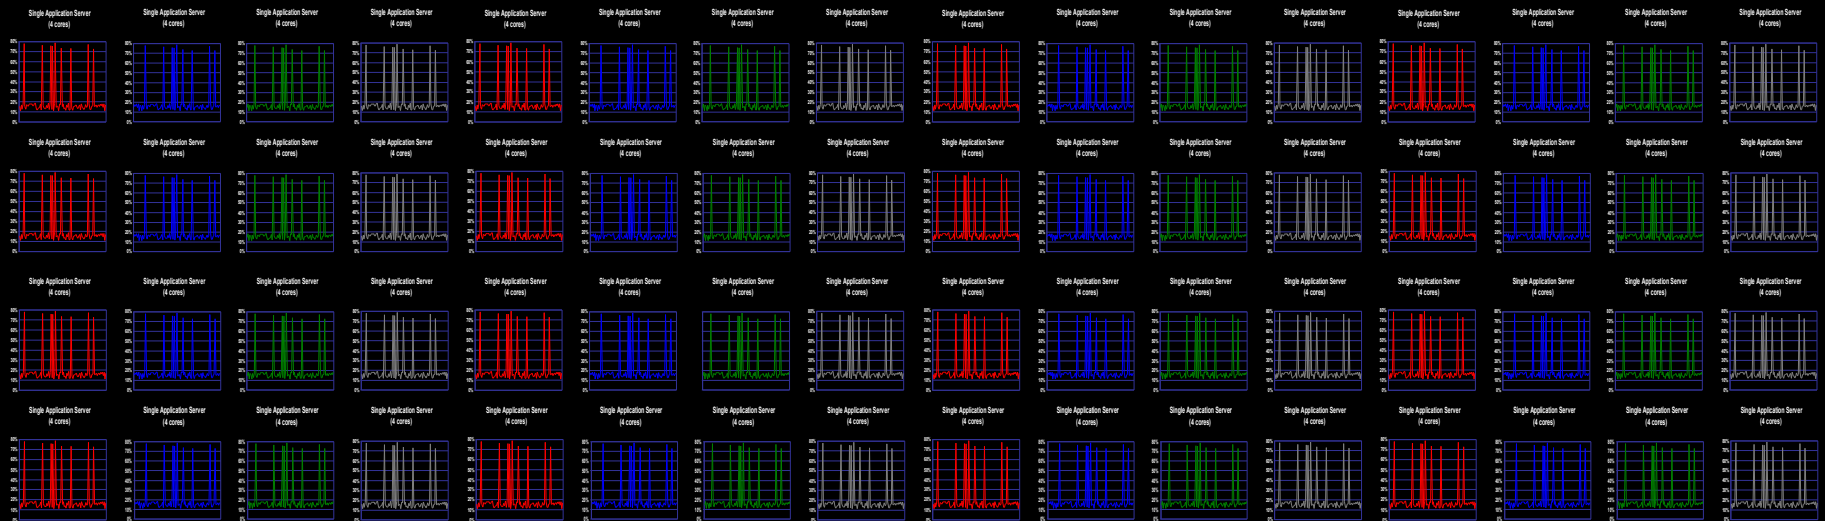

64 separate workloads on one system\*

Average utilization is 61%

Peak is 78%

**\*** 256 cores reduced to 72 cores (3.5 to 1)

**64 to 1 Systems Consolidation  
(72 cores)**






# Virtualization enables higher system usage through consolidation and workload smoothing

The amount of “leverage” increases as the number of workloads are added

The potential for savings increases with the amount of consolidation



## Server Consolidation Concepts


- ❖ Sharing system resources dramatically improves system utilization and reduces required capacity
- ❖ Degree of compaction can be affected by:
  - Mixture of workloads (Different application characteristics)
  - Amount of consolidation (8 workloads or 80 workloads)
  - System design (Ability to feed processor)
  - Sophistication of scheduling mechanism (PowerVM)
  - Degree of priority (Ability to assign various priorities)
  - Granularity of work allocation (1/100<sup>th</sup> of a processor)
- ❖ Server consolidation experience with SAP workloads on IBM Power Systems and System z<sup>TM</sup> support these conclusions

## Power Systems Impact on Server Consolidation Concepts

- ❖ 3X reduction in number of cores is possible with POWER6™ and PowerVM in a virtualized consolidation environment

**assuming all of the cores have equal performance**

- ❖ What if the new cores have better performance?



- ❖ IBM POWER6 technology with PowerVM has the potential to reduce the number of cores by **six fold or more**

# Very Large Scenario with Virtualized Server Consolidation

*with POWER6 technology*

64 separate workloads on 64 identical systems

Average utilization is 20%

Peak is 79%

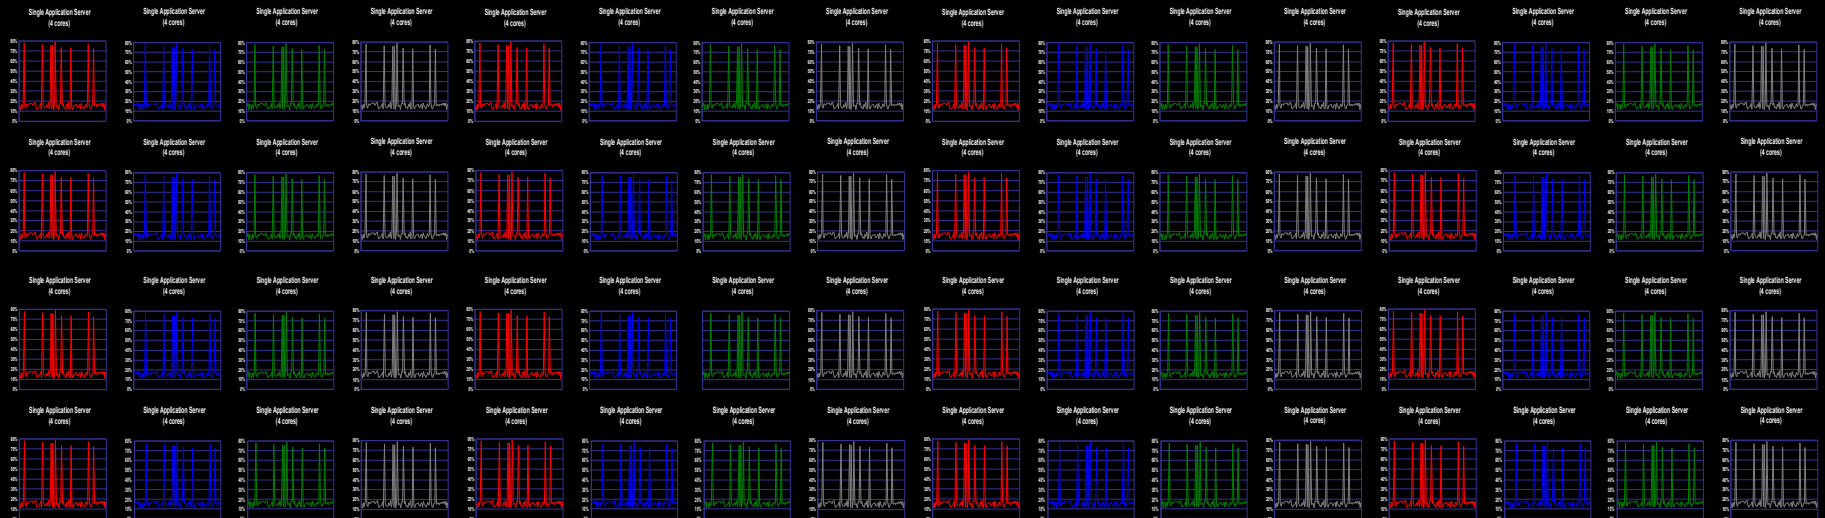
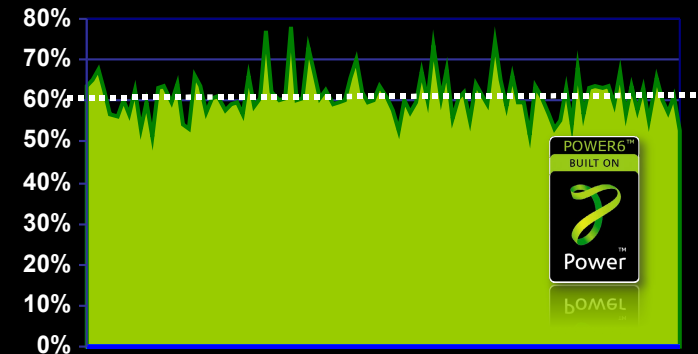
64 separate workloads on one system\*

Average utilization is 61%

Peak is 78%

**\*256 cores reduced to 36 cores (7 to 1)\***

64 to 1 Systems Consolidation  
(36 cores)



# Consolidation Template Example

4	Processor cores per system	
16	Quantity of servers	
64	Total processor cores	
— 2.65*	Consolidation factor	Range 2.0 to 6.0
— 2	Performance ratio	
12	Required cores with consolidation	

Mixture of workloads  
 Amount of consolidation  
 Average utilization  
 System design  
 Scheduling sophistication  
 Degree of priority  
 Granularity of work allocation

Performance ratio =  $\frac{\text{Performance per core of consolidated solution}}{\text{Performance per core of existing solution}}$   
 (Range 1.0 to 4.0)

**\* 5.3 consolidation ratio in processor cores (64/12)**  
 (Range 2.0 to 24.0)



# POWER6 and PowerVM

## Made for each other



- ❖ POWER6 systems are designed with scalable memory, memory bandwidth, and I/O bandwidth to handle high utilization rates
- ❖ POWER6 and PowerVM enable sharing of:




- POWER6 processors (Shared processors, Micro-Partitioning™)
- POWER6 memory (Active Memory Sharing™)
- POWER6 I/O (Virtual I/O Sharing)

- ❖ PowerVM dynamically manages and adjusts system resources
- ❖ PowerVM manages priorities with Partition Availability Priority
- ❖ PowerVM allocates resources down to 1/100<sup>th</sup> of a processor
- ❖ POWER6 and PowerVM support virtual servers up to 64 cores

# IBM Power 595

## Designed with the capacity for consolidation

Capacity per core relative to the Power 595



**3.5X** memory  
per core

**4X** memory bandwidth  
per core

**4X** I/O bandwidth  
per core

**8X** L2 + L3 cache  
per core

	<u>IBM Power 595</u>	<u>HP DL370 G6</u>	<u>Sun T5240</u>
Cores	64	8	16
Memory (GB)	4,096	144	128
Memory Bandwidth (GB/s)	1,376	38.4	63.4
I/O Bandwidth (GB/s)	640	19.2	8
L2 + L3 cache	1,280	18	8
Memory (GB) per core	64	18	8
Memory Bandwidth (GB/s) per core	21.5	4.8	3.9625
I/O Bandwidth (GB/s) per core	10	2.4	0.5
L2 + L3 cache per core	20	2.25	0.5

*You can use the tremendous capacity of the IBM Power™ 595 to run challenging applications in every virtual server.*

# Reduce cost with Active Memory Sharing

## *PowerVM dynamically optimizes shared memory*

Dynamically adjusts memory available on a physical system for multiple virtual images based on their workload activity levels:



- 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


Blends Power Systems hardware, firmware and software enhancements to optimize resources

Supports over-commitment of logical memory



## Reduce cost and complexity with Virtual LAN

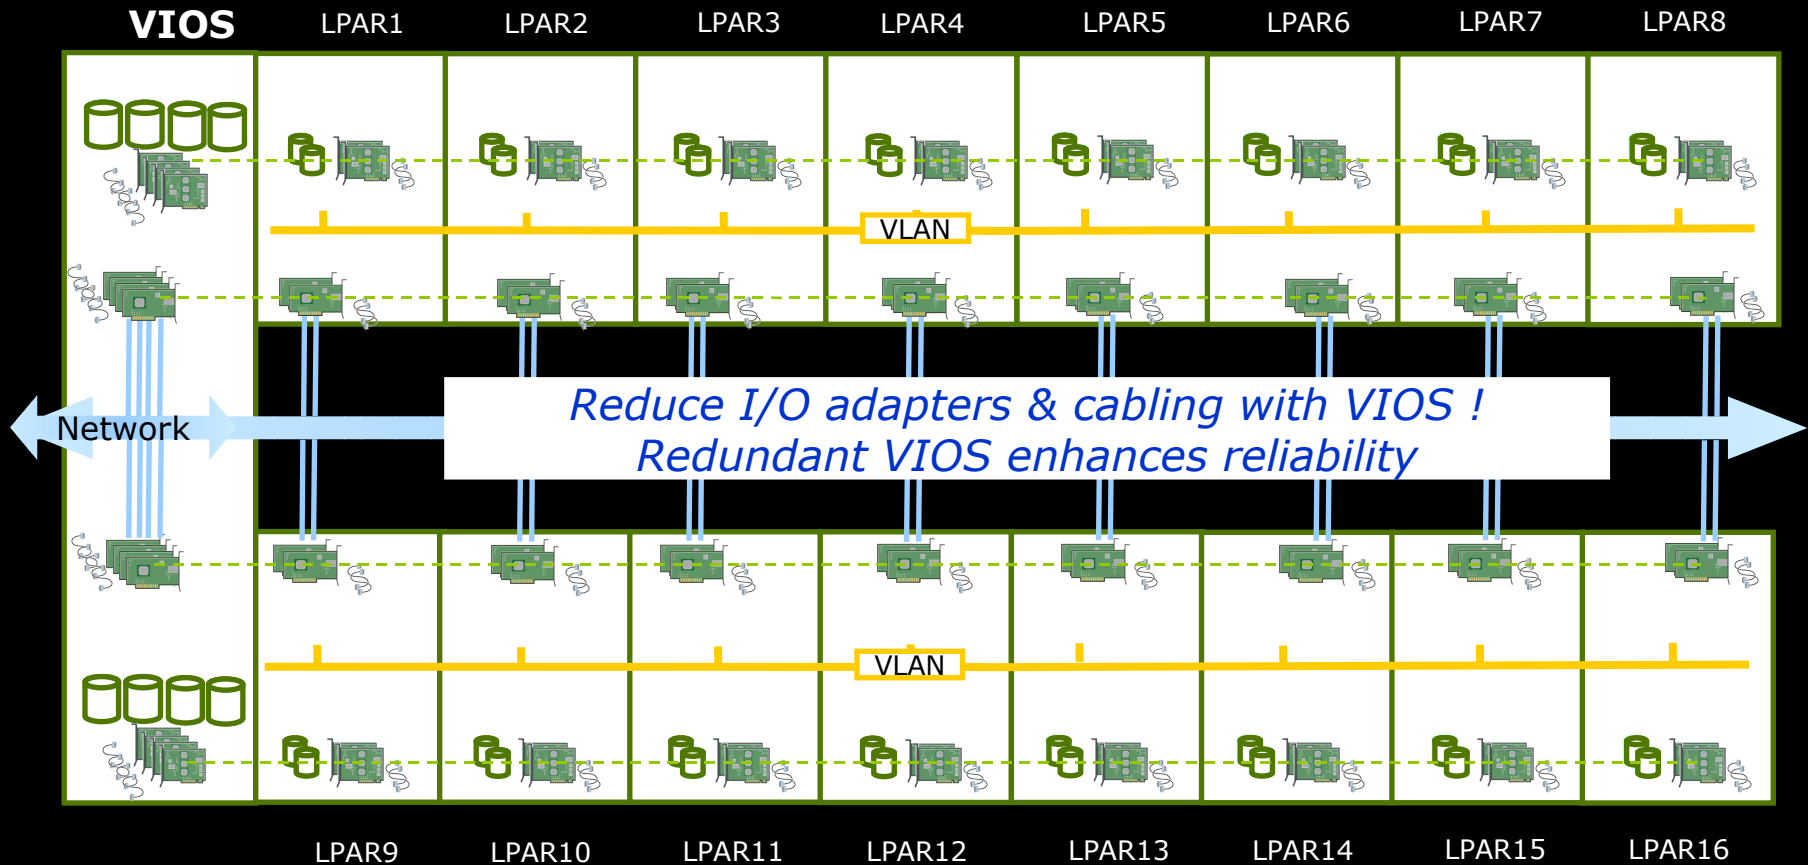
- Virtual networking for server-to-server communication can sustain performance while reducing network utilization and physical points-of-failure
- Removes the need for network adapters and switches for communication between LPARs
- Can drastically improve network recovery in the event of a disaster – no physical cables to re-connect or diagnose



*Reduce network latency*  
*Increase server network throughput up to 3X!*

## Reduce Cost with Virtual I/O Server

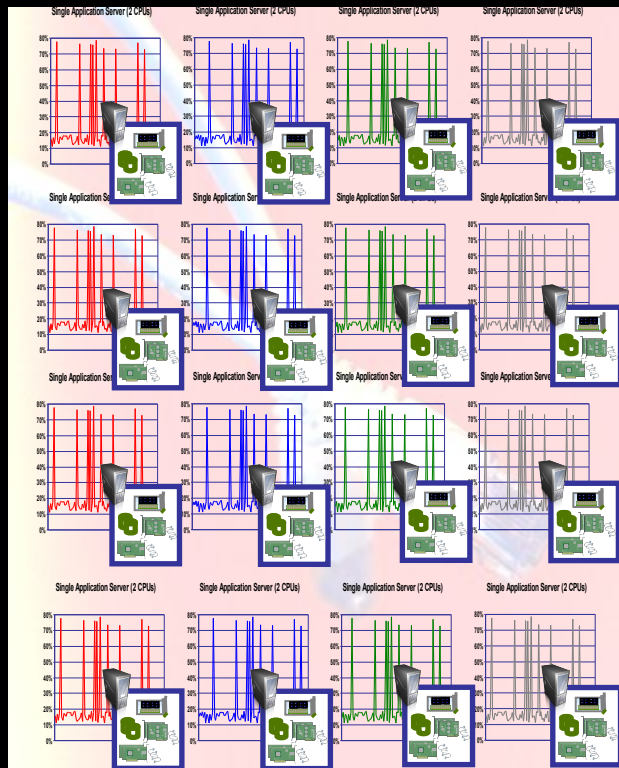
- VIOS can *dramatically* reduce physical resources and associated costs through more effective resource sharing






# Reduce cost by eliminating hardware and points-of-failure

- Eliminate underutilized adapters
- Reduce switches and cost of network operations and maintenance



*Virtualize & Consolidate*

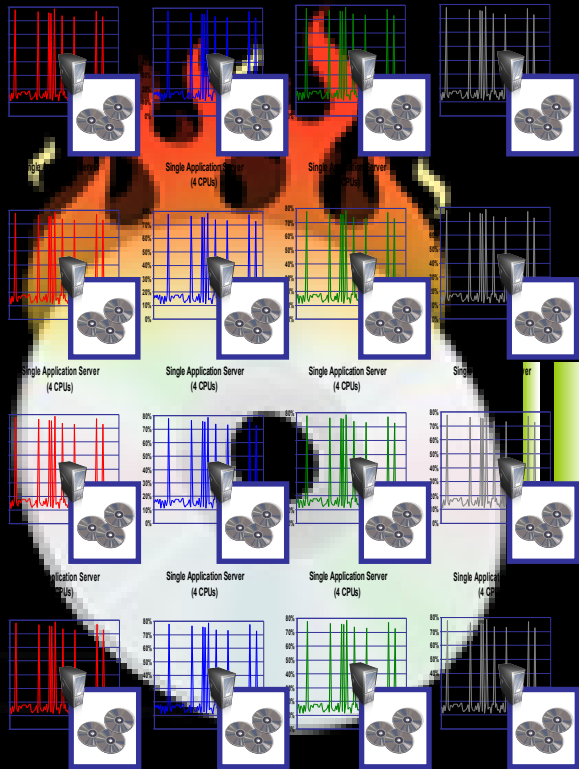


- *Reduce 32 disk drives to 8!*
- *Reduce 32 fibre adapters to 4!*
- *Reduce 32 LAN adapters to 4!*
- *Reduce 128 cables to 16!*


# Reduce cost of software licensing

*Virtualization can generate significant software savings*

- *Reduce software costs by up to 80%!*
- *Reduce software support costs*
- *Reduce software subscription costs*
- *Costs are ongoing operational costs*



Database software  
\$40k per core!



## Reduce cost with environmental efficiency

*By consolidating 16 x16-core Sun V890s into ONE Power 570 system*

- Save up to 92% in annual energy costs!*
- Reduce floor space required by over 85%*
- Reduce processing cores by over 87%*


**16 Sun V890s**  
(@ 20% utilization)

- **256 total cores** @ 2.1 GHz
- over 51 sq. ft. floor space
- up to 678 MWh annual energy



**One Power 570**  
(@ 60% utilization)

- **32 total cores** @ 4.2 GHz
- Only 7.6 sq. ft. floor space
- Save up to 629 MWh annually
- up to **\$94k in energy savings per year!**



## IMPROVE SERVICE by consolidating with Power Systems

- ❖ Delivering ***flexible performance*** to tailor capacity based on service level agreements and business priorities
- ❖ Supporting ***dynamic provisioning*** of virtual servers and resources to rapidly deploy new applications and infrastructure
- ❖ Enabling clients to ***avoid disruption*** to deliver growth or perform system maintenance



## Improve service with Flexible Performance



### **Power Systems capacity is fluid and can be tailored to changing requirements**

- A partition size is not fixed and may be adjusted dynamically to match demand; resources added or moved are instantly available
- Capacity on Demand offers insurance to efficiently respond to steady increases as well as unexpected spikes



### **Power Systems virtual server size is not limited by small physical system boundaries**

- Partition sizes can scale from a fraction of a processor to the full size of the system, dynamically
- Unused capacity in one partition can be used by other workloads when required, improving their performance



### **Power Systems scale seamlessly to deliver balanced growth**

- Massive memory and I/O bandwidth enable linear performance as workload and capacity increase
- Modular, Hot-node Add options enable smooth system growth



## Improve service with Dynamic Provisioning



**Virtual servers can be quickly and easily deployed to accommodate new applications and respond to rapidly changing business needs**

- Virtual provisioning takes minutes and can even be done remotely
- Virtual provisioning leverages existing resources and does not require installation of additional hardware
- Workload Partitions can simplify provisioning even further



**New workloads can be tested simply on virtual servers with minimal resources**

- Testing is accomplished in secure, isolated partitions that are created when needed and consume resource only when used



**Virtual servers may be de-commissioned or re-used with minimal effort, easily recycling their resources**

- Resources are infinitely reusable and continuously shareable with all workloads, enabling high utilization and efficiency

## Improve service by Avoiding Disruption



### **Enterprise Power Systems are designed for growth without application outage**

- Available system resources can be automatically increased without disruption to the application through PowerVM
- New resources can be physically added to the server without downtime through Capacity on Demand and Hot-node Add
- Applications can be moved to other systems without disruption using Live Partition Mobility and Live Application Mobility



### **Enterprise Power Systems are designed for concurrent service to avoid costly application downtime**

- System resources can be repaired while operations continue using Hot-node Repair
- Operating systems and firmware can often be updated without disruption to operations



### **Pervasive redundancy, dynamic failover and extensive self-healing functions convert unplanned component failures into autonomic events**

- Failed processors and memory DIMMS can be automatically replaced by the system, maintaining application availability



## Power Systems avoid disruption for our clients

- Best availability according to Yankee group


**Live Partition Mobility and Live Application Mobility for seamless upgrades and maintenance**

**Power Systems have exclusive features to keep you going even when something goes wrong**

- Bit Steering helps avoid double bit memory errors
- Alternate Processor Retry keeps applications running even when a processor has a hard failure
- I/O EEH to recover from transient PCI bus errors

**New features reduce your risk even more**

- RAID controllers
- Virtual tape support
- Solid State Drives with fewer moving parts than disk



According to a recent Yankee Group study\* of 400 Windows, Linux and UNIX users, AIX was the most reliable server operating system:

*"IBM's AIX achieved the highest level of reliability, with corporate enterprises reporting an average of only 36 minutes of downtime per server in a 12-month period"*

\* Source: "Unix, Linux Uptime and Reliability Increase; Patch Management Woes Plague Windows"  
© 2008 Yankee Group Research, Inc. All rights reserved

# POWER6 RAS – clearly the one to keep you in the GREEN



Feature	POWER6	SPARC	Integrity	Xeon
<b>Application/Partition RAS</b>				
<b>Live Partition Mobility</b>	Yes	No	No	Yes
<b>Live Application Mobility</b>	Yes	No	No	No
<b>Partition Availability Priority</b>	Yes	No	No	No
<b>HW Isolated Partitions</b>	No	Yes <sup>1</sup>	Yes <sup>1</sup>	No
<b>System RAS</b>				
<b>OS independent First Failure Data Capture</b>	Yes	No	No	No
<b>Redundant System Interconnect</b>	No	Yes	No	No
<b>Processor RAS</b>				
<b>Processor Instruction Retry</b>	Yes	Yes	No	No
<b>Alternate Processor Recovery</b>	Yes	No	No	No
<b>Dynamic Processor Deallocation</b>	Yes	Yes	Yes	No
<b>Dynamic Processor Sparring</b>	Yes	Yes <sup>2</sup>	Yes <sup>2</sup>	No
<b>Memory RAS</b>				
<b>Memory Keys</b>	Yes	No	No	No
<b>Chipkill™</b>	Yes	Yes	Yes	Yes
<b>Redundant Memory</b>	Yes <sup>3</sup>	Yes <sup>3</sup>	Yes <sup>3</sup>	Yes <sup>3</sup>
<b>I/O RAS</b>				
<b>Extended Error Handling</b>	Yes	No	No	No


#1,2,3 - See "POWER6 RAS" in backup; See the following URLs for addition details:

<http://www-03.ibm.com/systems/migratetoibm/systems/power/availability.html>

<http://www-03.ibm.com/systems/migratetoibm/systems/power/virtualization.html>

## MANAGE RISK by consolidating with Power Systems

- ❖ *Delivering **Proven Security*** that allows you to stay connected with confidence
- ❖ *Offering **Business Resiliency*** solutions that allow you to weather the unforeseen
- ❖ *Enabling **Effective Systems Management*** to allow you to simplify your operations



## Manage risk with Proven Security



### **Power Systems security offers assurance to aggressively pursue opportunities in today's interconnected world**

- IBM is a trusted, stable IT vendor with decades of experience with virtualization & consolidation Built as a completely integrated solution, not a collection of products from three or more vendors
- World-class operating system security with AIX and i
- Fewer avenues of entry



### **Recognized security certification for the whole platform, including PowerVM virtualization**

- EAL4 certification in place for POWER6 hardware and software




### **It is easier to secure & administer one enterprise server with virtualization than a collection of smaller servers**


- Fewer operating systems to install security patches on
- Fewer systems to monitor for intrusion and compliance
- Eliminates need to synchronize security policies across multiple systems
- Reduced number of administrators and administrative touch points
- Eliminates domino effect" of a single system with weak security falling and being used to leverage escalation of privileges and node hopping

# Securely run multiple applications on a single system

Power servers with PowerVM are certified for Common Criteria EAL4+



National Information Assurance Partnership  
**Common Criteria Certificate**  
*is awarded to*  
**International Business Machines Corporation**



The IT product identified in this certificate has been evaluated at an accredited testing laboratory using the Common Methodology for IT Security Evaluation (Version 2.3) for conformance to the Common Criteria for IT Security Evaluation (Version 2.3). This certificate applies only to the specific version and release of the product in its evaluated configuration. The product's functional and assurance security specifications are contained in its security target. The evaluation has been conducted in accordance with the provisions of the NIAP Common Criteria Evaluation and Validation Scheme and the conclusions of the testing laboratory in the evaluation technical report are consistent with the evidence adduced. This certificate is not an endorsement of the IT product by any agency of the U.S. Government and no warranty of the IT product is either expressed or implied.

**Product Name:** IBM Logical Partition Architecture for Power Systems  
**Evaluation Platform:** IBM Power 595 (9119-FHA), IBM Power 570 (9117-MMA) and IBM Power 550 (8204-ESA)  
**Protection Profile Identifier:** None

**CCTL:** Science Applications International Corporation  
**Validation Report Number:** CCEVS-VR-VID10299-2008  
**Date Issued:** 26 November 2008  
**Assurance Level:** EAL 4 Augmented ALC\_FLR.2

**Original Signed By**  
*Director, Common Criteria Evaluation and Validation Scheme*  
National Information Assurance Partnership

**Original Signed By**  
*Information Assurance Director*  
National Security Agency



## Manage risk with Business Resiliency



### **Enterprise Power Systems are designed to minimize the risk of an unplanned outage**

- Enterprise Power Systems leverage superior quality components whenever feasible to minimize component failure
- Enterprise Power Systems endure extensive system testing that includes PowerVM, AIX, and IBM i to ensure solid performance under pressure
- Enterprise Power Systems utilize extensive predictive failure technologies to avoid impending component failures
- Enterprise Power Systems enable memory keys to avoid inadvertent memory overwrites that lead to unexplained system crashes



### **Enterprise Power Systems inspire confidence to know that regardless what happens, your business will continue to operate**



- Disaster recovery becomes part of the plan, not just an afterthought
- System failover becomes manageable with fewer pieces to manage
- Supports regular and sustained role swap operations with PowerHA

# Manage risk with Effective Systems Management



## **Simplify your operations by consolidating to one system**

- One system, one set of firmware, one set of virtualization software
- Fewer cables, switches, routers to order, install, track, and maintain
- Share operating system images with AIX Workload Partitions
- Support multiple levels of AIX, Linux and IBM i if required



## **Simplify your operations with one consistent management tool**

- Let System Director be your view into your system and across your partitions, including energy management
- Systems Director features are integrated with POWER6 and PowerVM and work across operating systems
- Systems Director can even manage Windows, Linux, zLinux servers, and storage systems in the same consistent interface



## **Automate your IT environment by enabling your system to diagnosis problems**

- When components are internal to a single system, troubleshooting and repair times are reduced
- Fewer networking and storage device interconnections simplifies monitoring and control of service offerings



## The Advantage of Full Service

- ❖ POWER6, Power Systems, PowerVM, AIX, IBM i
- ❖ Technology and solutions that are:


***Designed for each other***

- **Architected together**
- **Built together**
- **Tested together**
- **Delivered together**
- **Supported together**





## Manage risk with security, control and simplicity



*16 production entry servers, limited RAS*  
*16 additional servers for HA/DR*  
*Multiple failover solutions*  
*2X networking & increased operations costs*  
*Complex recovery – days or weeks*

**OR**

*Two enterprise Power Systems servers,  
extensive RAS*  
*Single, high-speed network connection*  
*Simple failover with PowerHA*  
*Regular role-swap ensures proven and  
rapid recovery of all applications*




# Consolidating with Enterprise Power Systems enables clients to...

## REDUCE COST

*Resource sharing drives higher utilization rates and lowers costs*

*Resource sharing reduces networking costs*

*Resource sharing dramatically reduces energy consumption, floor space and software costs*



## IMPROVE SERVICE

*Dynamically adjust partition resources to provide the performance your workloads require, when they need it*

*Respond quickly and easily to changing business requirements*

*Avoid costly application outages and keep your business moving*

## MANAGE RISK

*Stay connected with confidence in today's virtual world*

*Withstand the unforeseen, even in a turbulent time*

*Simplify your operations and your life*

# Thank You!