



PowerVM Trends and Directions

Jay Kruemcke

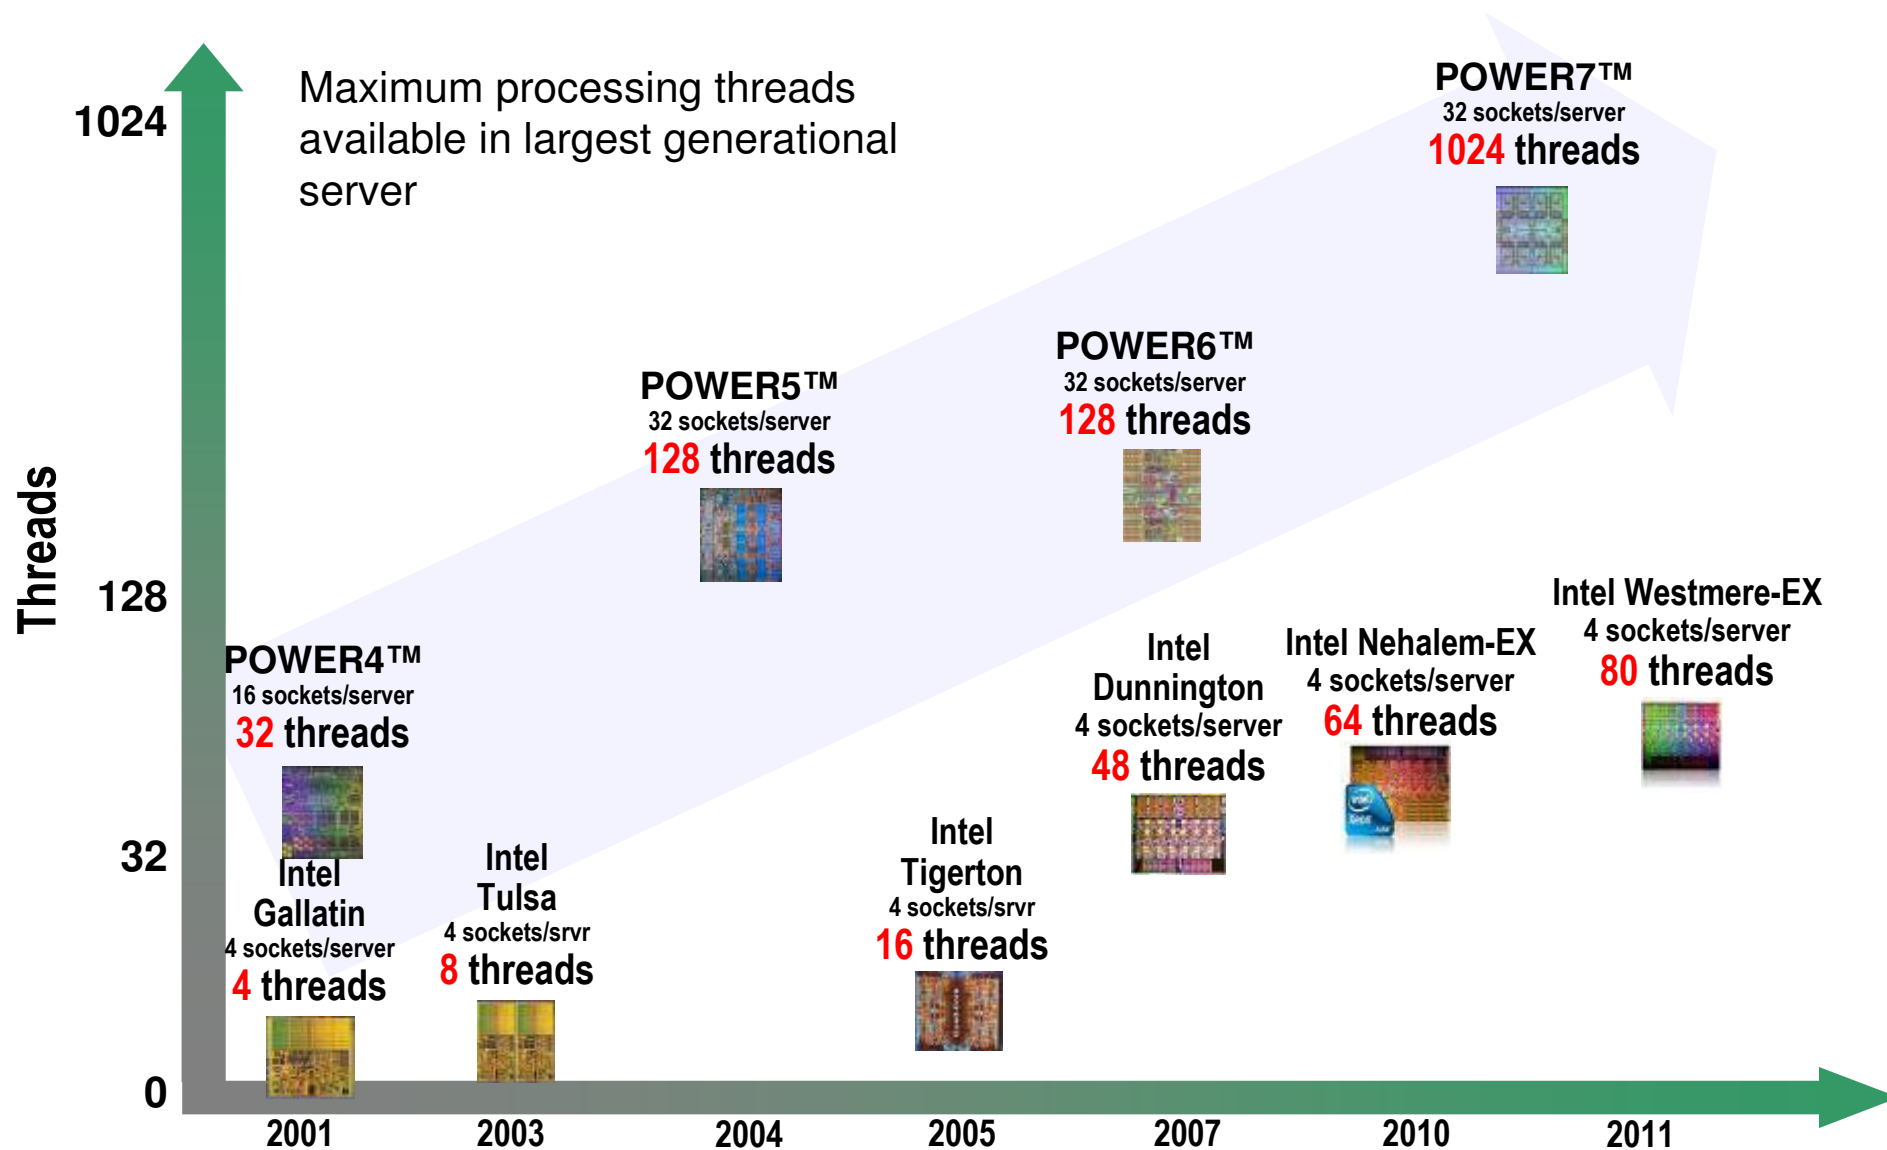
Twitter: [@chromeaix](#)

jayk@us.ibm.com

Kruemcke.com

30th de Mayo, 2013, IBM Argentina

More Processing Capacity In Power Servers



Power Systems Scalability and Value Growth



POWER4 P690

(2002)

List Price ~\$3M

32 Processors

rPerf – 60.6

Weight 1000kg

POWER 710

(2010)

List Price ~\$11K

6 Processors

rPerf 76.69

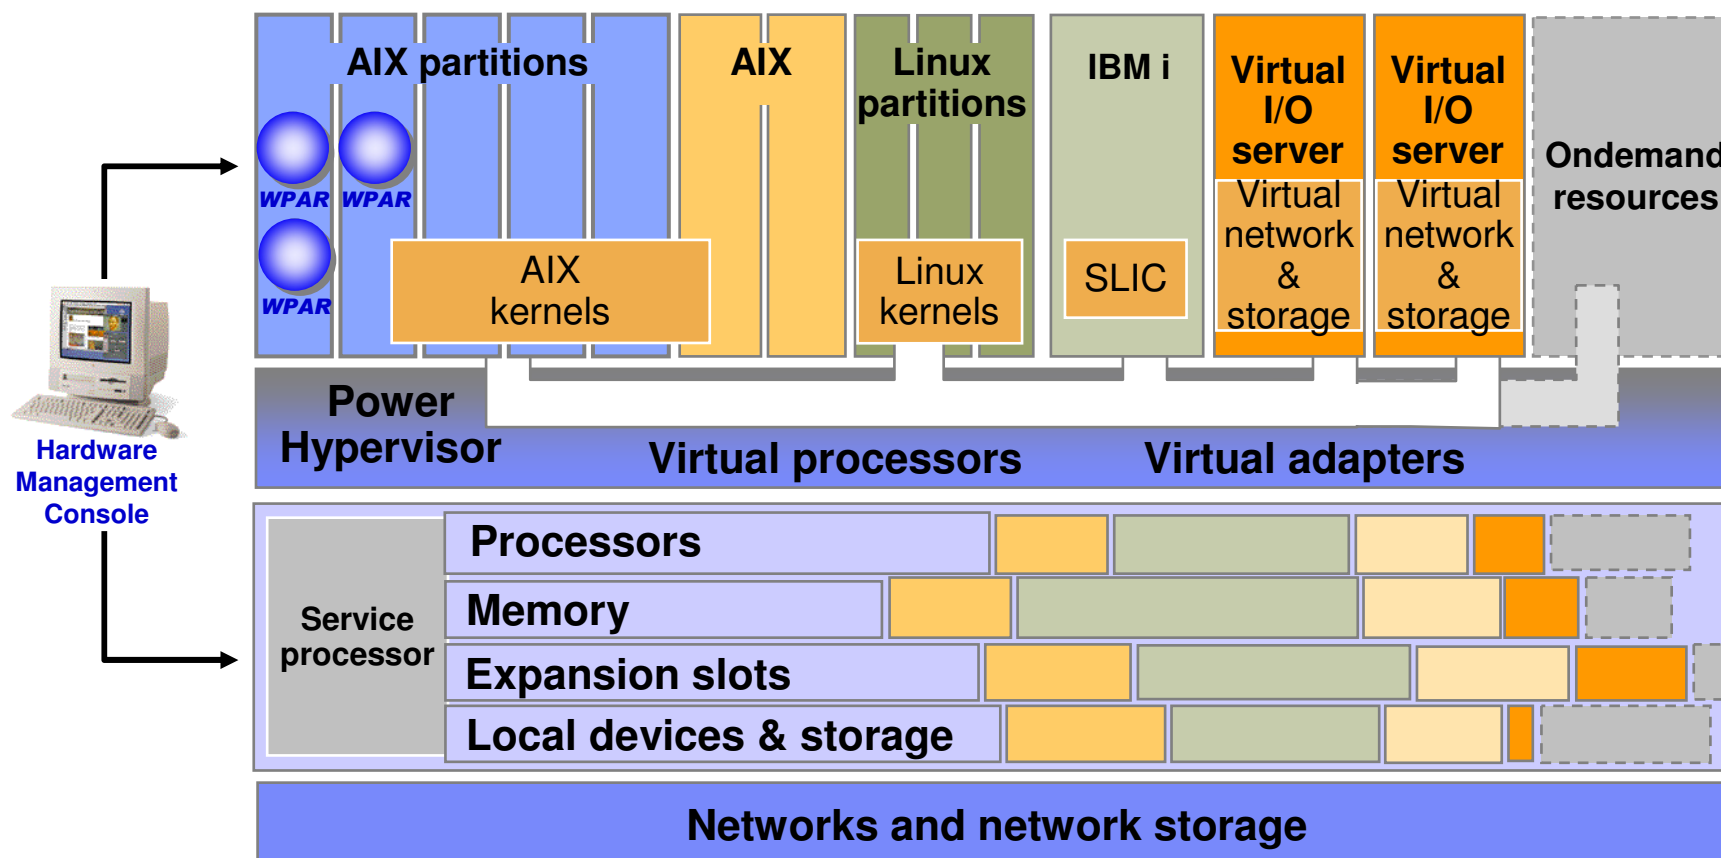
Weight 28kg

PowerVM™ Virtualization Architecture



PowerVM key design points

- Designed for high **efficiency** to provide high overall performance
- Designed for high **scalability** – linear from 1/20 to 256 cores
- Designed for maximum resource **granularity** to reduce wasting resources
- Designed for **isolation** to provide security and “no compromise” consolidation



The value of Power Systems is linked to effective IT

If you do these:

- Virtualization
- Consolidation
- Resource Sharing

You get these:

- Lower IT costs
- Flexibility
- Effectiveness

If you don't use these:

- Virtualization
- Consolidation
- Resource Sharing

You get these:

- Server spawl
- Wasted resources
- Even more budget pressure


Study: Only 1 in 5 can allocate more than half their IT budget to innovation

Least efficient data centers

Use of new technology:

- 43% first and fast technology adoption
- 1% move virtual machines to meet desired outcomes
- 21% use storage virtualization
- 3% use a storage service catalog (tiered storage)

Results:




Most efficient data centers

Use of new technology:

- 86% first and fast technology adoption
- 58% move virtual machines to meet desired outcomes
- 93% use storage virtualization
- 87% use a storage service catalog (tiered storage)

Results




Source: 2012 IBM Data Center Study: www.ibm.com/data-center/study (<http://www.ibm.com/data-center/study>)

Typical small server utilization

Typical server running a single UNIX operating environment is less than 20% utilized

- Configuration planned for peaks (50% unused?)
- Configuration planned for growth (20% unused?)
- System waits for I/O and memory access even when it is working (20% unused?)



- Single workload on a single system
 - Average Utilization: 20.7%
 - Peak: 79%


Result is that 80% of the hardware, software, maintenance, floor space, and energy that you pay for, is wasted

Scenario with Virtualized Server Consolidation

8 to 1 Systems Consolidation
(16 CPUs)


- 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)**




Utilization increases
while peak actually
decreases


Single Application Server
(4 CPUs)




Single Application Server
(4 CPUs)




Single Application Server
(4 CPUs)




Single Application Server
(4 CPUs)




Single Application Server
(4 CPUs)




Single Application Server
(4 CPUs)



Single Application Server
(4 CPUs)



Single Application Server
(4 CPUs)



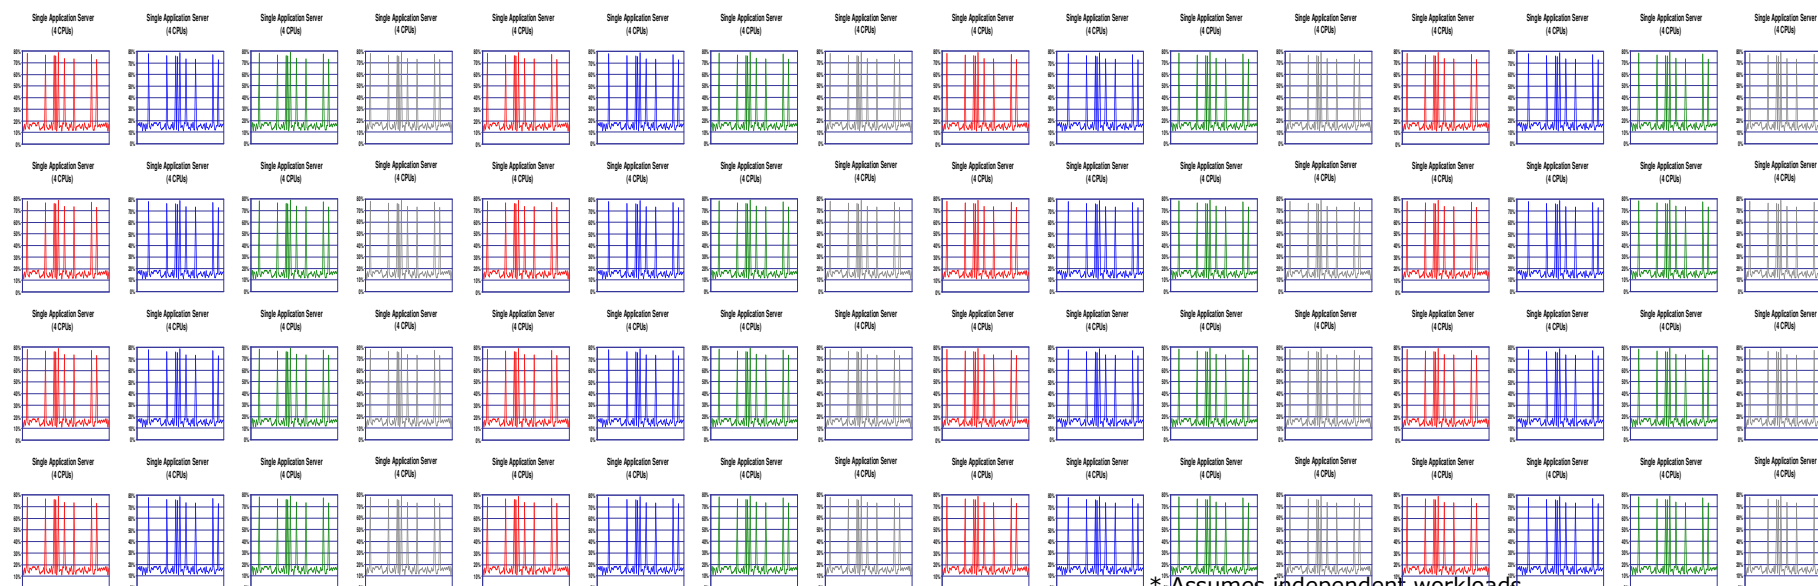

* Assumes independent workloads

Very Large Scenario with Virtualized Server Consolidation

64 to 1 Systems Consolidation
(72 CPUs)

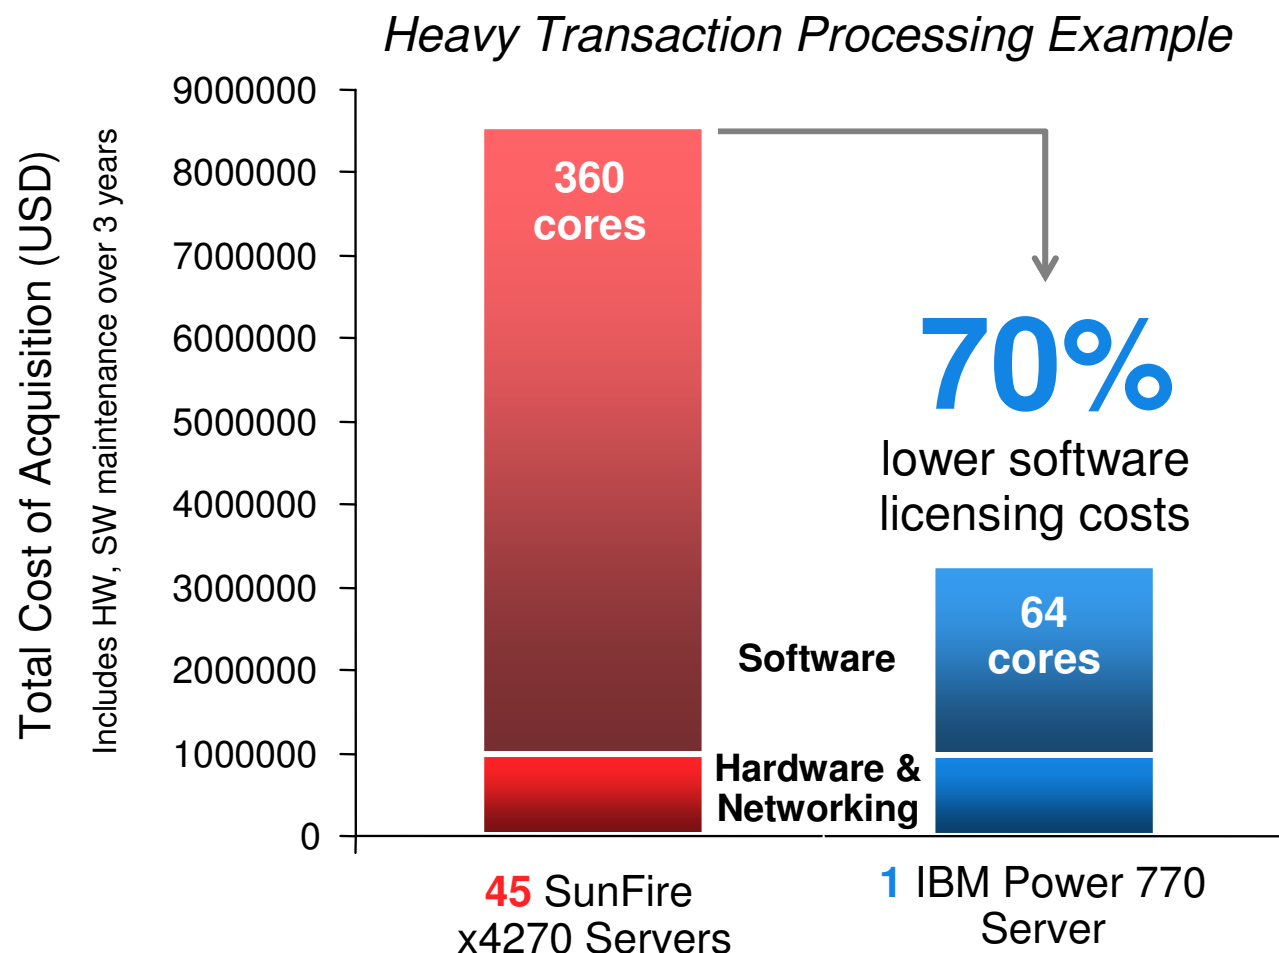
- 64 separate workloads on 64 identical systems
 - Average utilization is 20.7%
 - Peak is 79%
- 64 separate workloads on one system*
 - Average utilization is 61%
 - Peak is 78%

*** 256 CPUs reduced to 72 CPUs (3.5 to 1)**




* Assumes independent workloads

Consolidation Helps Drive Down Software Licensing Costs




IBM PowerVM Hypervisor is More Efficient than Competitive Hypervisors



Source: IBM Software Group Competitive Study

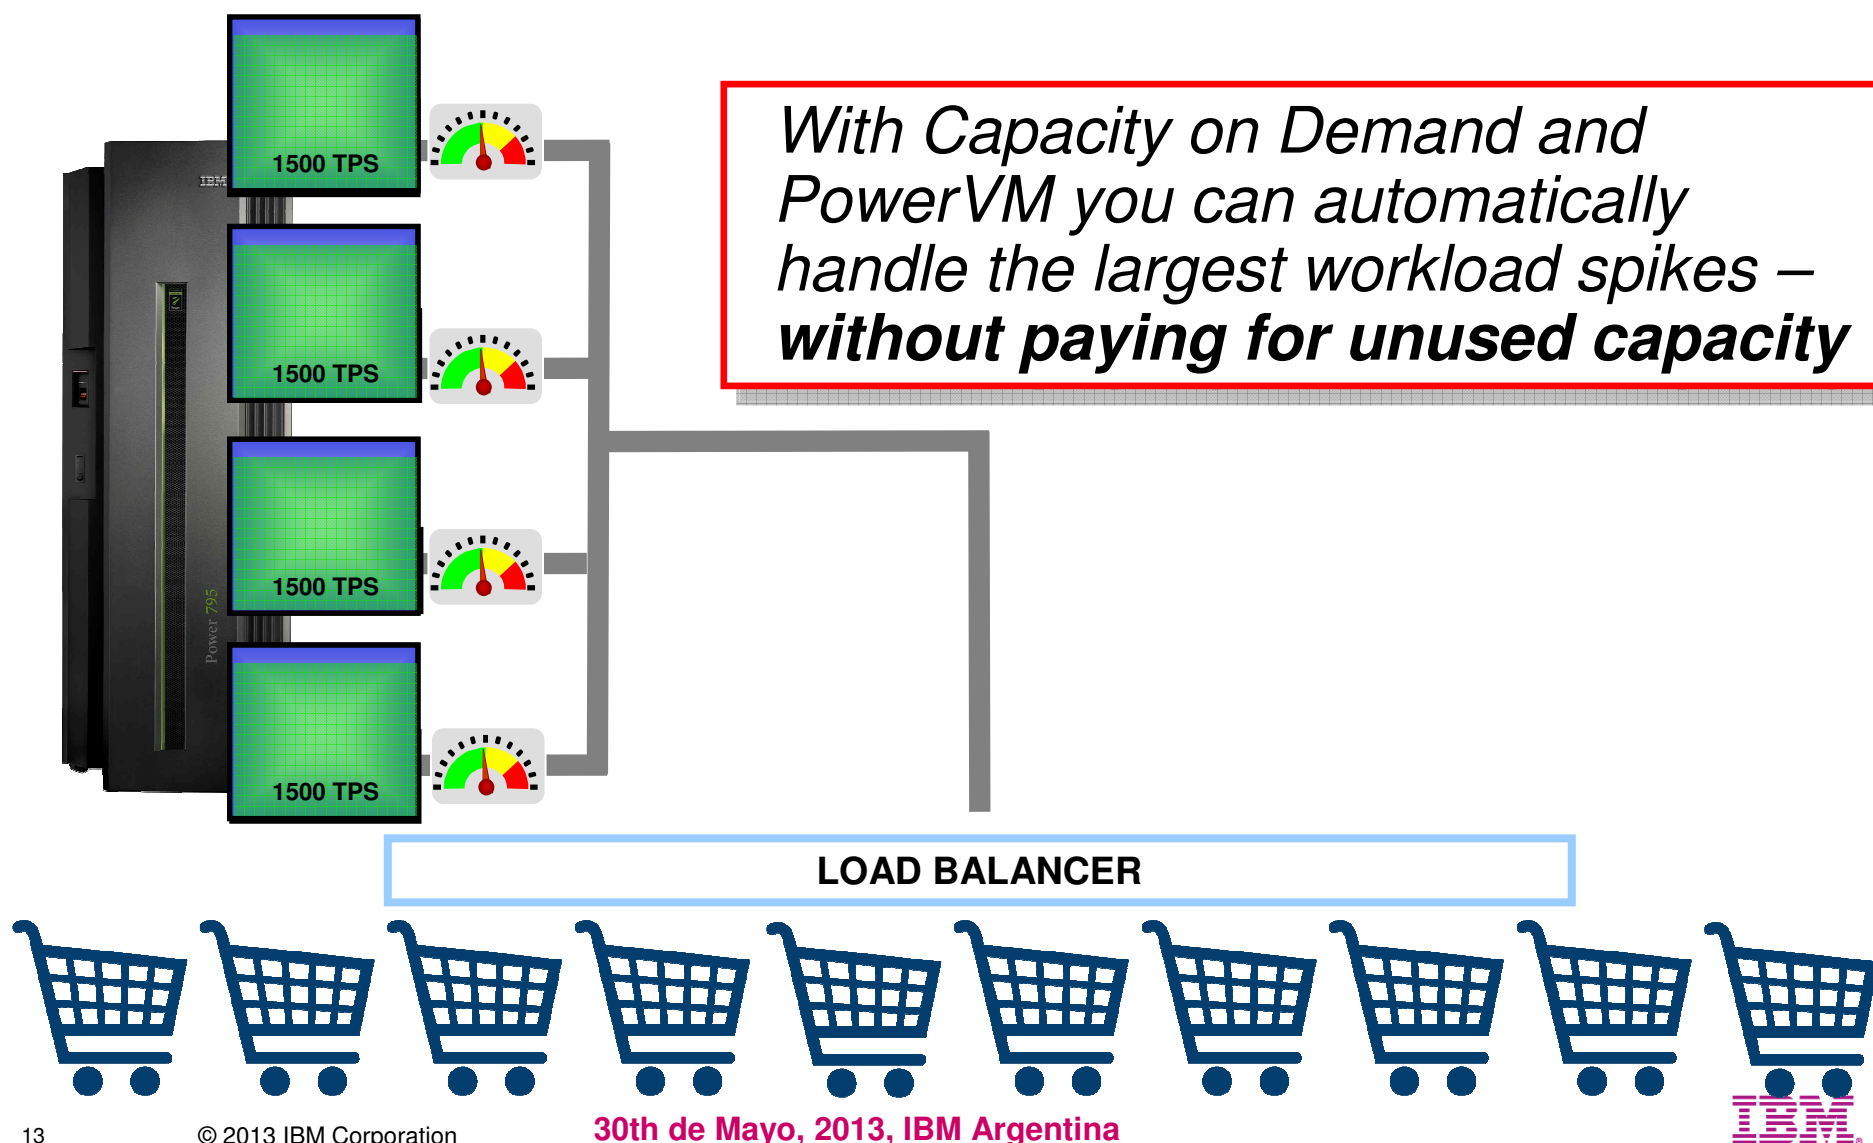
PowerVM delivers superior flexibility to optimize IT resource utilization and improve responsiveness



<i>Flexibility Factors</i>	VMware vSphere 4 & 5	PowerVM
Dynamic virtual CPU changes in VM	Add (but not Remove)	Yes
Dynamic memory changes in VM	Add (but not Remove)	Yes
Dynamic I/O device changes in VM	Some	Yes
Direct access to I/O devices from within VM	Some (with VT-d enabled)	Yes
Integrated LPAR and WPAR support	No	Yes

Source: <http://www.vmware.com/files/pdf/products/vsphere/vmware-what-is-new-vsphere5.pdf>

The best of everything: PowerVM and Elastic Capacity on Demand



PowerVM delivers firmware-based security



- Unlike x86-based virtualization products, the PowerVM hypervisor is secure by design. IBM is the only vendor that has designed the virtualized environment from 'bare metal' through the hypervisor.
- PowerVM hypervisor is part of the digitally-signed firmware with strong cryptography which makes it very difficult to attack.
- There are zero vulnerabilities reported against the PowerVM hypervisor by US CERT or by MITRE Corporation
- PowerVM is certified at a Common Criteria Evaluated Assurance Level 4+

Remember, zero is a number too ...
a very good number in the Security domain.



PowerVM Agenda

2012 Enhancements Review

- *VIOS Performance Advisor*
- *Shared Storage Pools*
- *Dynamic Platform Optimizer*

Future Enhancements


- *Virtualization Management Futures*
- *SR-IOV Future*
- *Remote Restart*
- *Shared Storage Pools*

Fall 2012 PowerVM and Virtualization Features

- Live Partition Mobility for IBM i (1H12)
- Allow 1/20 core minimum entitlement
- LPM Concurrency Improvements up to 16 concurrent
- LPM Performance Improvements (up to 3x)
- VIOS Performance Advisor
- Shared Storage Pools Improvements



VIOS Performance Advisor



- Makes assessments and recommendations based on the expertise and experience available within the IBM systems performance group
- Runs within the VIOS for a user specified amount of time (hours)
- Polls and collects key performance metrics
- Analyzes results
- Provides a health check report and proposed changes to the environment

New!! VIOS Performance Advisor

- ✓ Proactive VIOS Health Check
- ✓ Provides Advise on how to Tune the VIO Server Before Problems Occur
- ✓ Shipped with PowerVM 2.2.2
- ✓ Reports are viewed with a web browser
- ✓ Can be use as a historical measure of performance

CPU Section Snapshot

VIOS - CPU							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
✓	CPU Capacity	1.0 ent	-	01/12 14:01:29	-	n/a	n/a
🔒	CPU Consumption	avg:14.1% (cores:0.2) high:14.3% (cores:0.2)	-	-	-	n/a	n/a
🔒	Processing Mode	Shared CPU, (UnCapped)	-	01/12 14:01:29	-	n/a	n/a
⚠️	Variable Capacity Weight	128	129-255	01/12 14:01:29	-	1	5
✗	Virtual Processors	6 vCPUs	3 vCPUs or less	01/12 14:01:29	-	n/a	n/a
✓	SMT Mode	SMT2	-	01/12 14:01:29	-	n/a	n/a

SYSTEM - SHARED PROCESSING POOL							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
✓	Shared Pool Monitoring	enabled	-	01/12 14:01:29	-	n/a	n/a
🔒	Shared Processing Pool Capacity	3.0 ent.	-	01/12 14:01:29	-	n/a	n/a
✓	Free CPU Capacity	avg_free:2.9 ent. lowest_free:2.8 ent.	-	-	-	n/a	n/a


PowerVM – Shared Storage Pools (SSP)

Extending Integrated Storage Virtualization Beyond a Single Server

Advantages of Shared Storage

- **Cost savings** through sharing; efficient utilization of physical I/O
- **Facilitates server consolidation** by enabling more LPARs (increased density)
- **Agility**: quick LPAR deployment / teardown
- **Simplify** Live Partition Mobility
- **Reduce** SAN infrastructure and **SAN** management **costs**
- **Reduce** datacenter **footprint**

PowerVM with Shared Storage Pools




SSP enables optimized utilization of SAN resources across many Power Servers

SSP provides space efficient virtual storage that can be rapidly provisioned

Shared Storage Pools

Server based storage virtualizer that is clustered across multiple Power servers

- **Server and storage integration**



■ Clustered Storage

- *Share common pool of storage*
- *Location transparency*
- *Snapshot / rollback*
- *Linked clones - deduplication*
- *Thick or thin provisioned vols*
- *Live storage mobility support*

■ Benefits


- *Agility, automation, elasticity*
- *Utilization*
- *Simplifies virtual disk setup*
- *Faster time to value*

Fall 2012 VIOS Shared Storage Pools

- Increase scaling to 16 nodes in a cluster
- Scalability Improvements
- Storage Utilization Statistics and Reporting
- Cluster Level RAS Improvements



Shared Storage Pools Enhancements in 2012



Scalability Comparisons 2011 / 2012

- Non-disruptive Rolling Upgrade
- Repository resiliency
- Thin provisioning management and reporting
- IPv6
- Multiple network interface support
- Tagged VLAN support

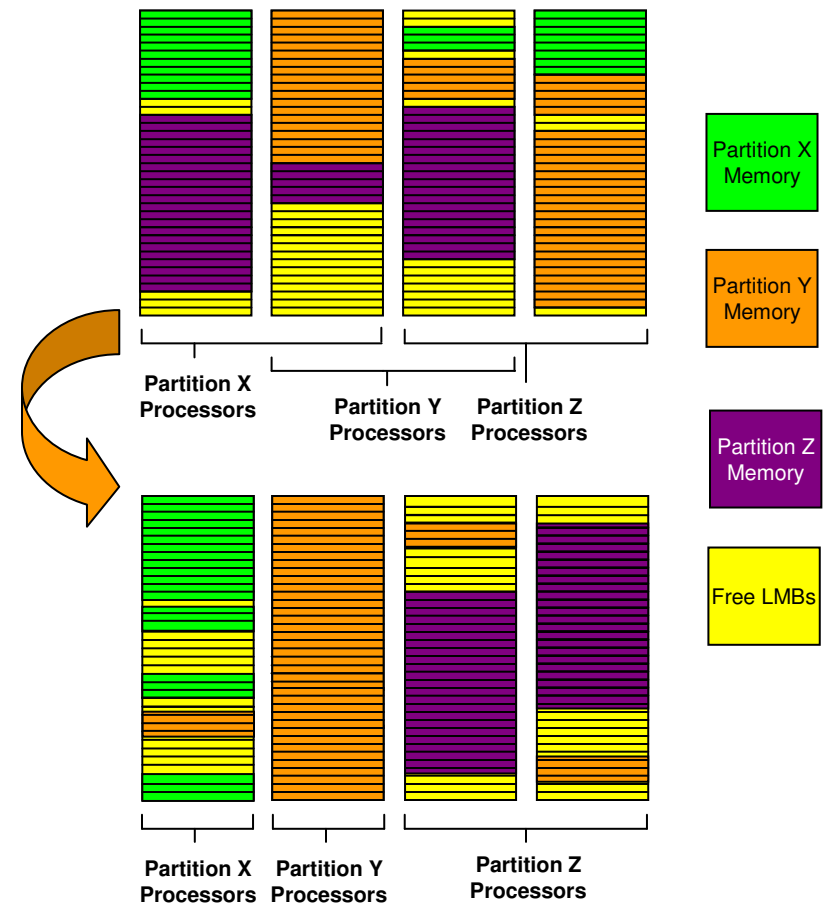
➤ Max Number of Nodes in Cluster	4 / 16
➤ Max Number of Physical Disks in Pool	256 / 1024
➤ Max Number of Virtual Disks	1024 / 8192
➤ Max Number of Client LPARs per VIOS	40 / 125
➤ Max Capacity of Physical Disks in Pool	4TB / 4TB
➤ Min/Max Storage Capacity of Storage Pool	128TB / 512TB
➤ Max Capacity of a Virtual Disk (LU) in Pool	4TB / 4TB

PowerVM Dynamic Platform Optimizer (DPO)

Dynamic Platform Optimizer (DPO)

Dynamically optimize partition placement to mitigate NUMA effects

- **Motivation:** Dynamic operations such as VM creation/deletion, dynamic reconfiguration, partition mobility can cause partition placement to be sub-optimal
- **Project Details**
 - Partition placement (memory, CPUs) optimized dynamically to improve affinity
 - OS adjusts to new affinity properties after optimization operation
 - Free of charge
 - 4Q/12, Release 760
- **Customer/IBM Benefits**
 - Improved performance in a cloud environment, p6->p7 migration, mobility
 - Improves P6 to P7 migration story
 - VM adapts to new topology after mobility operation



User Interface Summary

- Control is via HMC command-line interface only
- Start/Stop Optimization (optmem)
 - *“optmem” controls starting and aborting an operation*
- Requested/protected partition lists
 - *Sets of partitions can be prioritized or protected (untouched) for the DPO operation*
- Optimizer status (lsmemopt)
 - *Shows progress percentage if currently running*
 - *Result (success, failure, etc) of last operation*
 - *Impacted partitions*
- Notion of current and potential “affinity score”
 - *Enables system administrator to make decisions about value of running optimizer*
 - *System-wide score from 0-100*

GRACIAS

