

# International Technical Support Organization

## POWER8 Technology and System Technical Deep Dive

WRZ13G

Alexandre Bicas Caldeira

[abicas@br.ibm.com](mailto:abicas@br.ibm.com)

ATS Power Systems

Brazil



# Agenda

- POWER8 processor technology and architecture
- POWER8 processor-based systems design points and advanced features
- Cache architecture
- Memory infrastructure
- PCIe Gen3
- CAPI
- Power S822, S814, S824, S812L, S822L systems
- Speeds and feeds, form factors, feature options
- Performance information
- POWER8 Solutions overview
  - Traditional Enterprise workloads
- Big Data/Analytics
- Cloud
- HMC version 8
- Requirements and system support
- New features in systems management and monitoring
- PowerVM new capabilities
- AIX exploitation of POWER8
- New and improved features
- IBM i exploitation of POWER8
- New and improved features
- POWER8 and Linux
- OpenPOWER Foundation
- Linux support - SuSE/redHat/Ubuntu
- PowerKVM

## The Power Systems ITSO team

- Scott Vetter, ITSO Systems Team Leader
- Dino Quintero – ITSO Project Leader - BAO and Infrastructure solutions
- Chrys Rayns – ITSO Manager - Systems and Storage redbooks

## The Power Systems/AIX Extended ITSO team

- Alexandre Caldeira - Consulting IT Specialist, ATS Power Systems, Latin America
- Andrei Socoliuc - IT Specialist, AIX/Power Systems, IBM Romania
- César Diniz Maciel - Executive IT Specialist, Power Systems, Global Techline, US
- Daniel Martin-Corben - Technical Solution Designer, IBM UK
- Marcos Quezada - Consulting IT Specialist, ATS Power Systems, Latin America
- Octavian Lascu - Complex IT Project Team Leader, IBM Romania
- Tony Steel - Senior IT Specialist, ATS Power Systems, IBM Australia
  
- ... and some others that contribute teaching workshops around the world.



## Your instructor

- Alexandre Bicas Caldeira – abicas@br.ibm.com
- I work for ATS Power LA – We provide pre-sales consulting for IBM and Business Partners
- Background info
  - Worked 10 years as a business partner with IBM solutions
  - Started with AIX 4.3.2, RS/6000 Model F50
  - Joined IBM in 2010 – Power System Pre-sales
- BIG ITSO fan!! Workshops and residencies
- 3 residencies so far
  - 6 redbooks
  - 1 workshop
- Started with Power Virtualization back in 2004 (POWER4 – Regatta)

# Power Systems

POWER8 is ready for the next generation of enterprise applications

## Designed for Big Data

First server processor generation optimized for big data & analytics with POWER8 innovative design

## Superior cloud economics

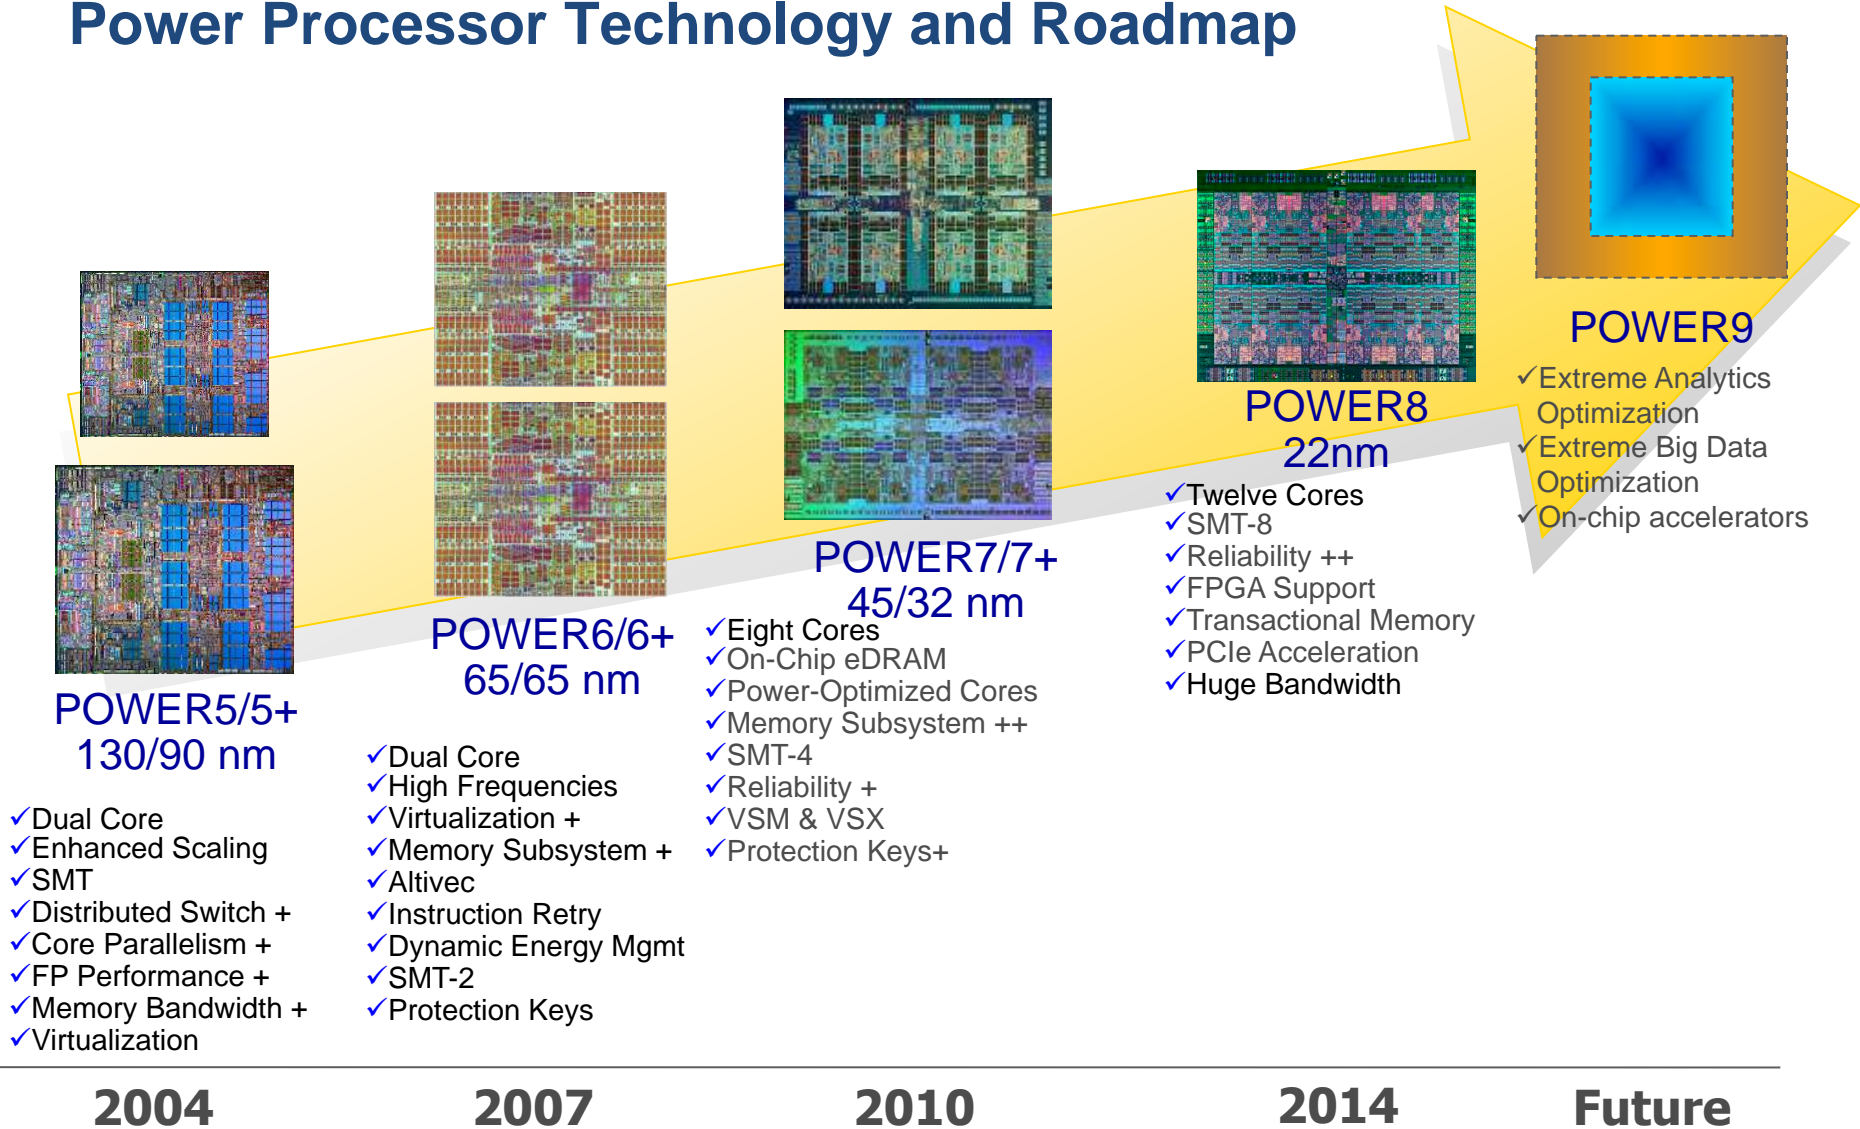
Superior cloud price/performance advantages & security to move data-centric applications to the cloud

## Open Innovation Platform

Delivering the world's first open server ecosystem revolutionizing the way IT is developed & delivered

**'IBM Introduces the first processor designed for big data'**

# Power Processor Technology and Roadmap



# POWER8 Highlights

**Announced 2013 at Hot Chips Conference**

**12 Cores per Socket/Chip**

**Significantly Strengthened Cores**

8 threads per core (SMT8)

Wider fetch / dispatch/issue of instructions

(8 fetch / dispatch, 10 issue)

Doubled highly utilized execution units

**Larger Caches per core/region**

L1: 64K data Cache, 32K instruction

Cache per core

L2: 512K private per core

L3: 8M per core (up to 96MB per chip)

L4: external to the chip

**2 Integrated Memory Controllers w/ Improved Latency & BW**

~ 25% memory latency improvement via on-chip fastpath interconnect

16MB mem cache / buffer chip

**Integrated SMP Interconnect w/**

improved “Flatness”

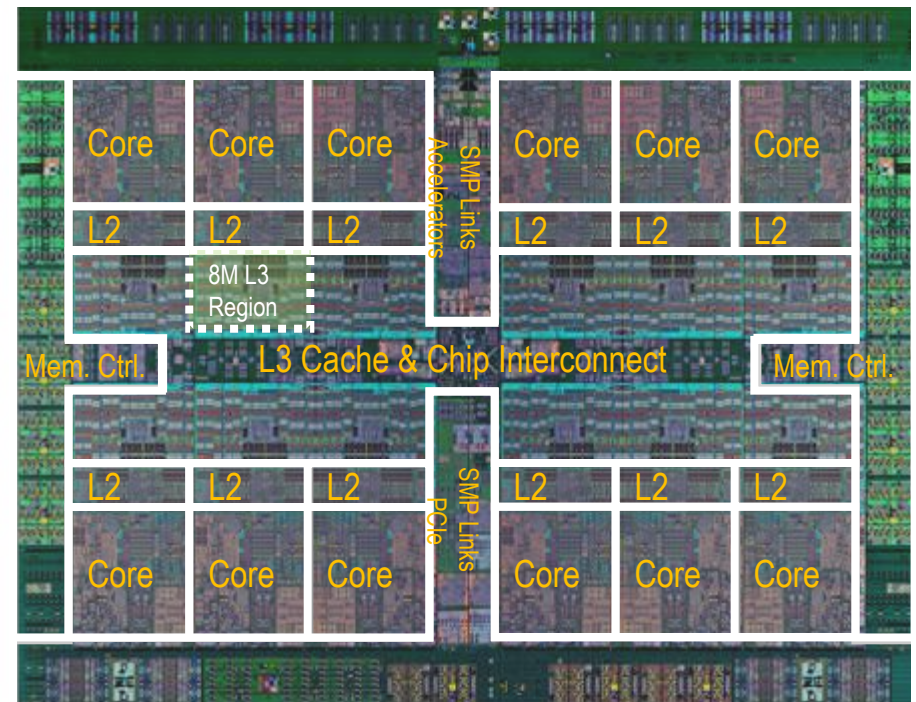
2-Hop fabric topology

**Integrated IO Subsystem**

On Chip PCIe Controller

**Fine Grained Power Management**

On Chip Power Management Controller & Power gating



# Scale-out Systems - DCMs and POWER8 Chips

## 1S & 2S servers use DCM (Dual Chip Module)

- 1 DCM fills 1 socket
- 1 DCM has two Scale-out POWER8 chips
- 1 DCM can provide 6-core, 8-core, 10-core or 12-core sockets


### 6-core Processor Chip

two 3-cores for 6-core DCM

two 4-cores for 8-core DCM

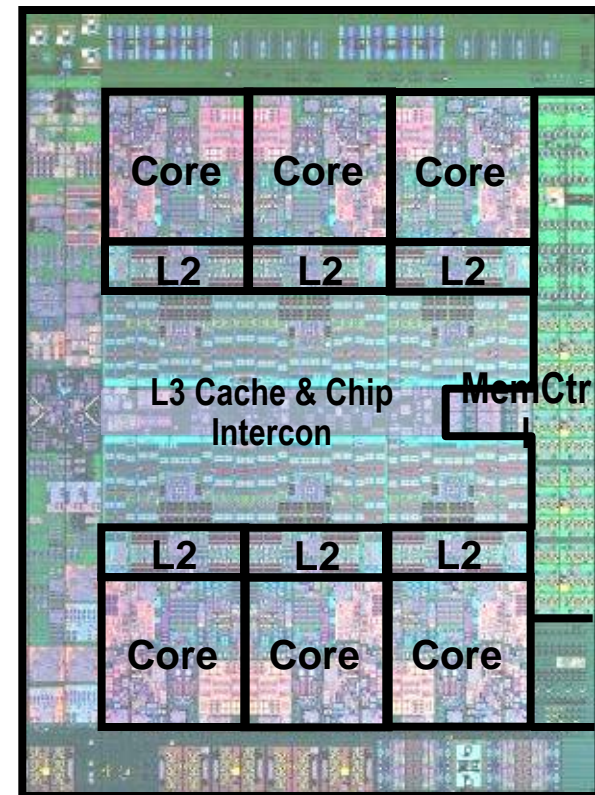
two 5-cores for 10-core DCM

two 6-cores for 12-core DCM



Max 8 Threads per Core


Excellent I/O bandwidth per socket






# POWER8 core

- Execution Improvement (vs POWER7)
  - SMT4 -> SMT8
  - 16 execution units:
    - 2 FXU, 2 LSU, 2 LU, 4 FPU, 2 VMX,  
1 Crypto, 1 DFU, 1 CR, 1 BR
- Larger Caching Structures (vs POWER7)
  - 2x L1 Data Cache (64 KB)
  - 2x Outstanding Data Cache misses
  - 4x Translation Cache
  - 4x Table walk machines
- Per core power/thermal management



# POWER8 core

- Execution Improvement (vs POWER7)
  - SMT4 -> SMT8
  - 16 execution units:
    - 2 FXU, 2 LSU, 2 LU, 4 FPU, 2 VMX, 1 Crypto, 1 DFU, 1 CR, 1 BR




	POWER7	POWER8
Max Cores	8	12
SMT	4	8
L1D	32 K	64 K
L1I	32 K	32 K
L2	256 K	512 K
L3	8 MB	8 MB
L4 (off-chip)	-	128 MB
Memory Controllers	2	2

# eDRAM technology

IBM's eDRAM technology benefits: Greater density, Less power requirements, Fewer soft errors, and Better performance

eDRAM enables on POWER8 96MB on-chip L3 cache

eDRAM enables 128MB off-chip L4 cache (on the DIMM)



eDRAM Cell


L3 Cache critical to balanced design / performance:

- 6:1 Latency improvement for L3 accesses vs external L3
- 2X Bandwidth improvement with on chip interconnect. 32B busses to and from each core
- No off chip driver or receivers in L3 access path.

eDRAM is nearly as fast as conventional SRAM but requires far less space

- 1/3 the space of conventional 6T SRAM memory cell implementation
- 1/5 the standby power
- Soft Error Rate 250x lower than SRAM ( Better availability )

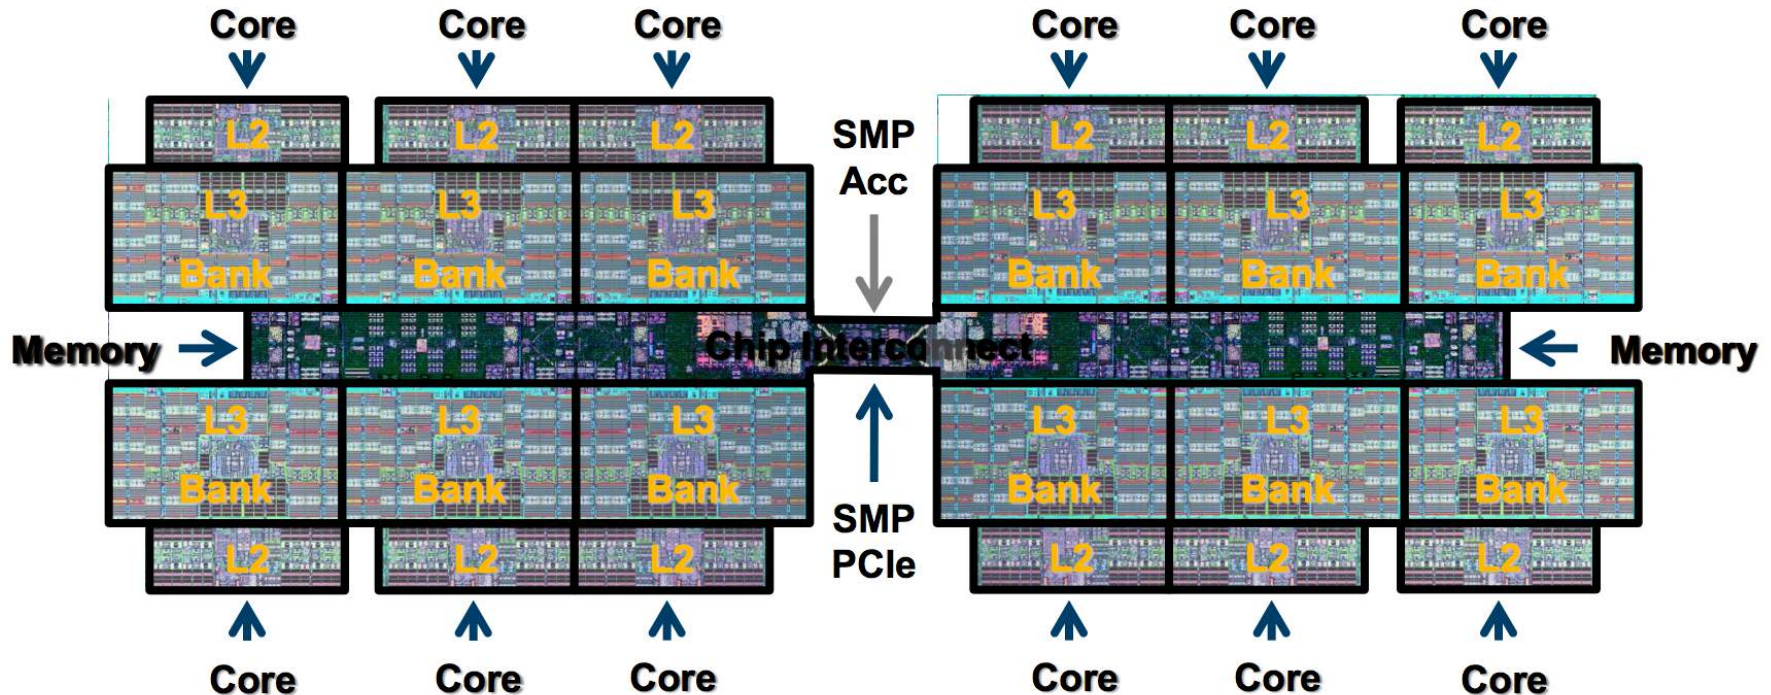
IBM is effectively doubling microprocessor performance beyond what classical scaling alone can achieve," said Dr. Subramanian Iyer, IBM Fellow






# POWER8 Chip Caches

- L2: 512 KB 8 way per core
- L3: 96 MB (12 x 8 MB 8 way Bank)
- “NUCA” Cache policy (Non-Uniform Cache Architecture)
  - Scalable bandwidth and latency
  - Migrate “hot” lines to local L2, then local L3 (replicate L2 contained footprint)
- Chip Interconnect: 150 GB/sec x 12 segments per direction = 3.6 TB/sec




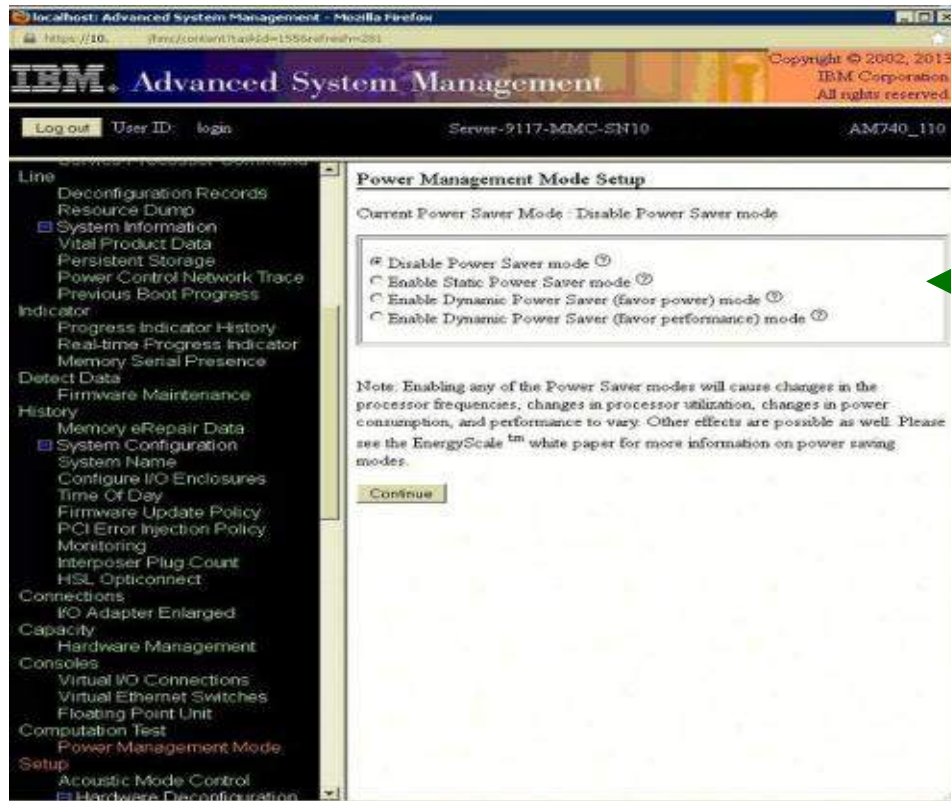
## POWER8 Multi-threading Options

- **ST**: Largest unit of execution work
- **SMT2**: Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT4**: Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT8**: Smallest unit of work, but provides the maximum amount of execution work per cycle
- Can dynamical shift between modes as required: ST / SMT2 / SMT4 / SMT8
- Mixed SMT modes supported within same partition
  - Requires use of “Resource Groups”



# POWER8 “Turbo”

12 Core Sockets




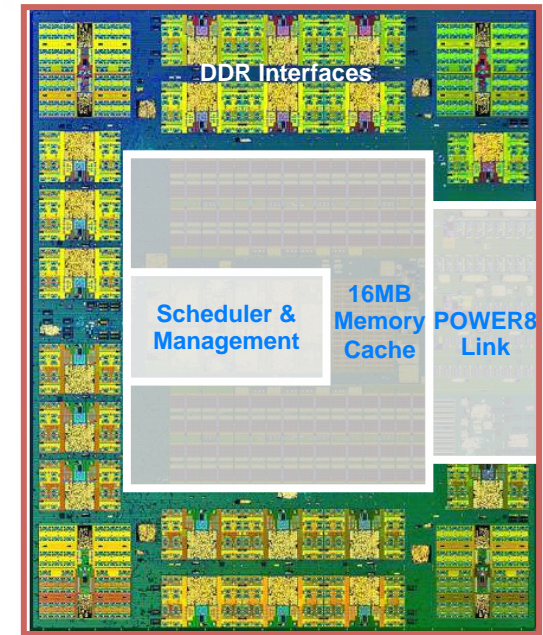
- Configure via ASMI menu
- Potential increase in processor frequency: ~ 11+%
- Requires firmware 810 ( POWER8 support)
- CPW & rPerf measured using nominal
- Note: fans run faster increasing noise level

This is not POWER7 Turbo-core !!!  
But the same option exists for POWER7 systems




# POWER8 Memory Chip

- IBM Memory DIMM
  - DDR3 RAM 1.600 MHz
- Intelligent Memory
  - Tasks performed by cores now done by DIMM
    - RAS decision point
    - Energy Management
    - Scheduling Logic
    - Writes Optimization
- L4 Memory Buffer
  - 16MB eDRAM / DIMM
  - Write Cache
  - Partial Updates
  - Reduces Latency
- 2 DIMM types
  - Low Profile
    - 2U & 4U Servers
  - High Profile – 4U Servers  
and future Enterprise servers




## POWER8 DIMM

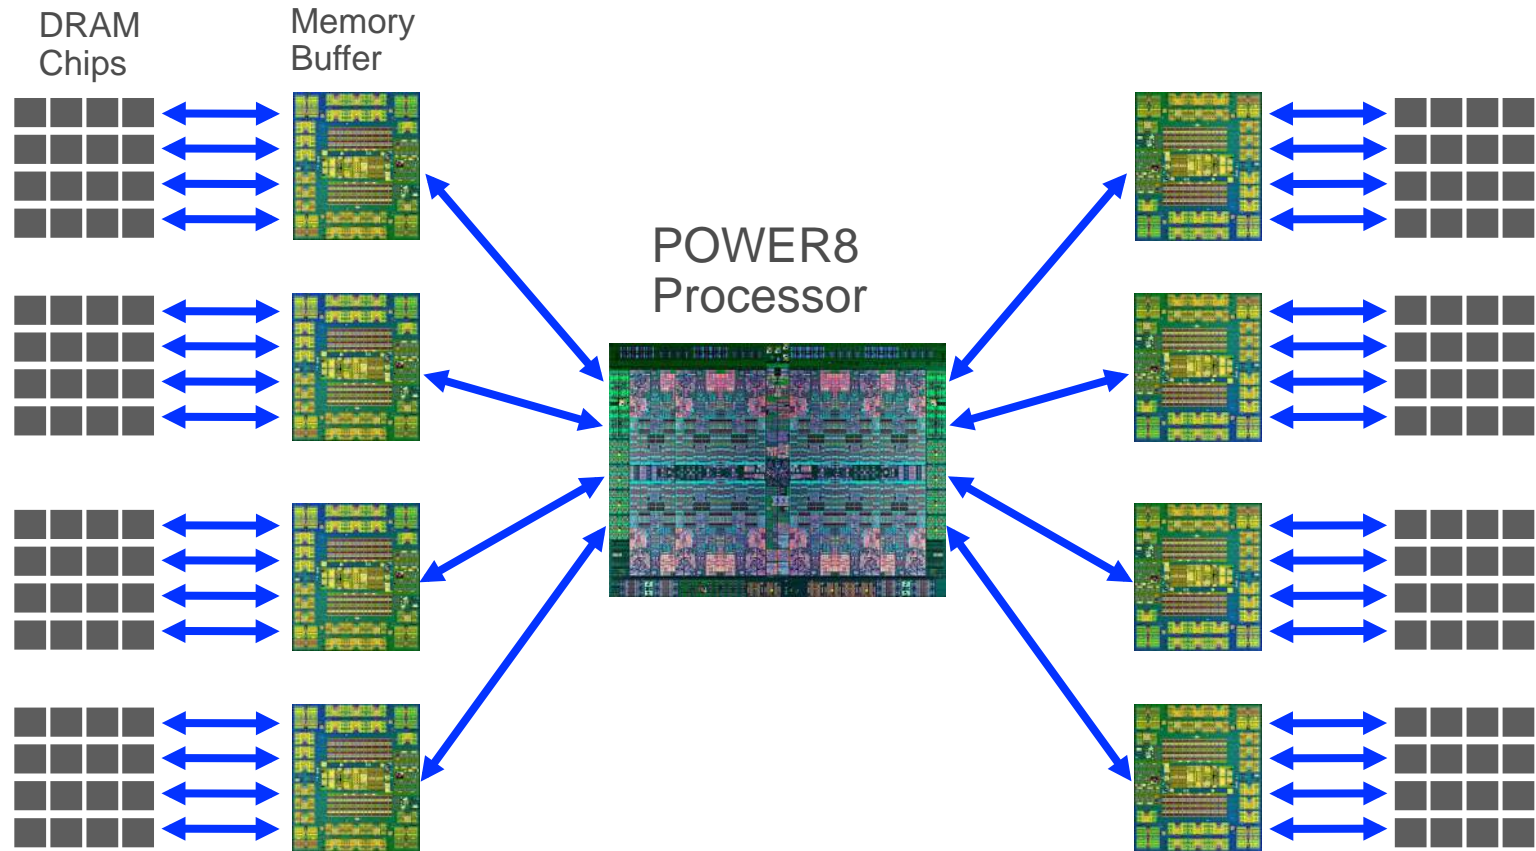
High-Profile – planned for  
4U Servers and Enterprise  
servers  
128GB/DIMM and above



Low-Profile  
2U & 4U Servers  
Up to 64GB/DIMM

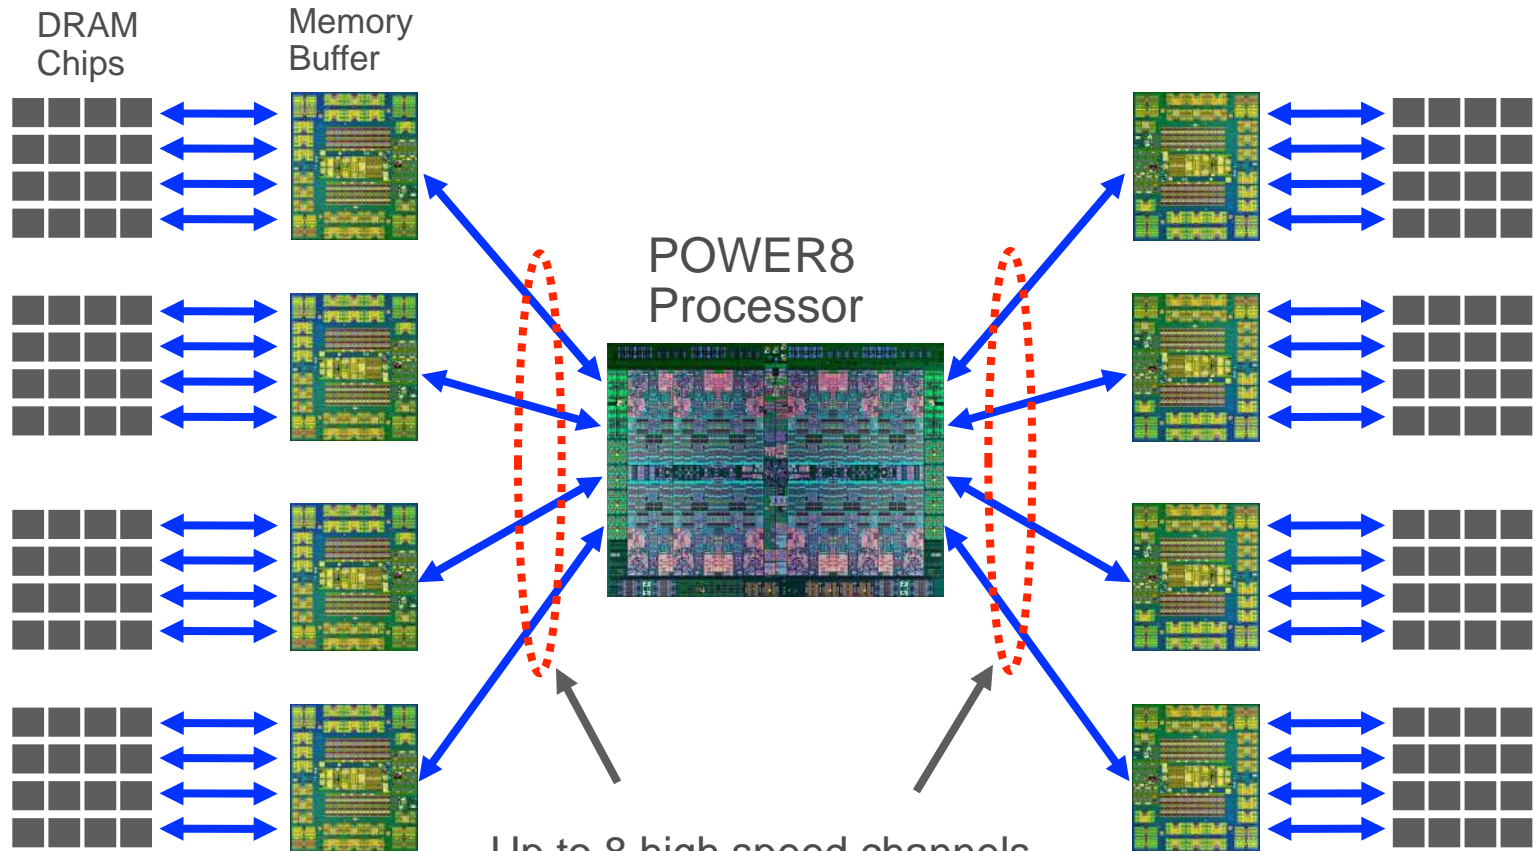


# POWER8 Memory



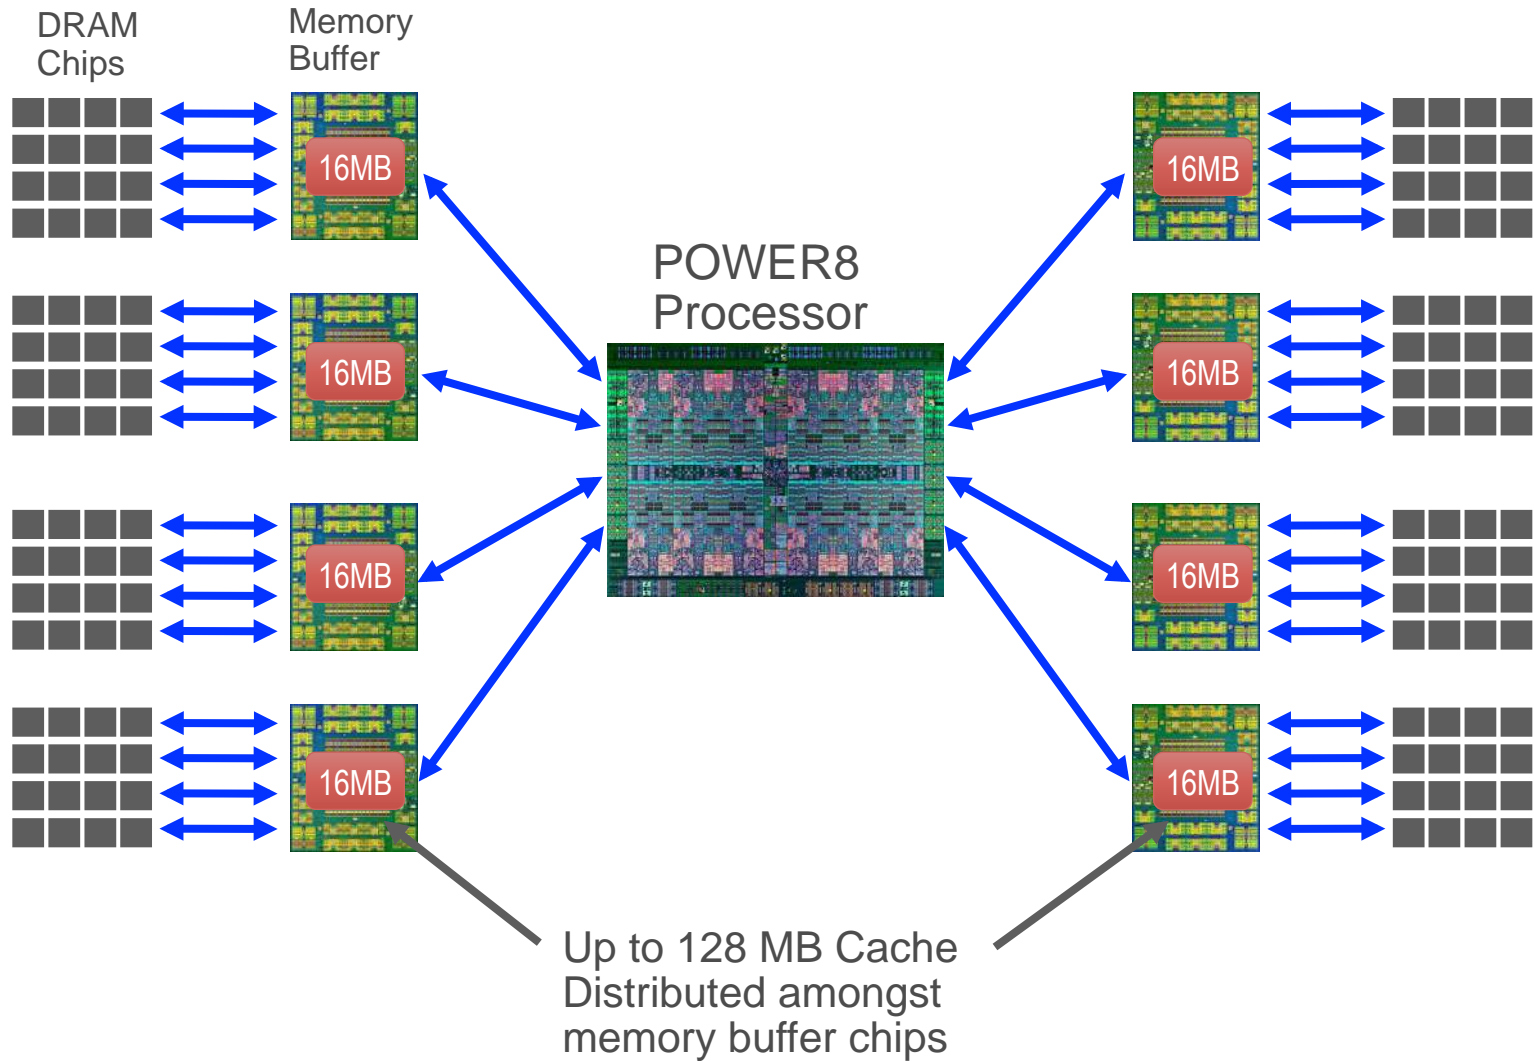


# POWER8 Memory



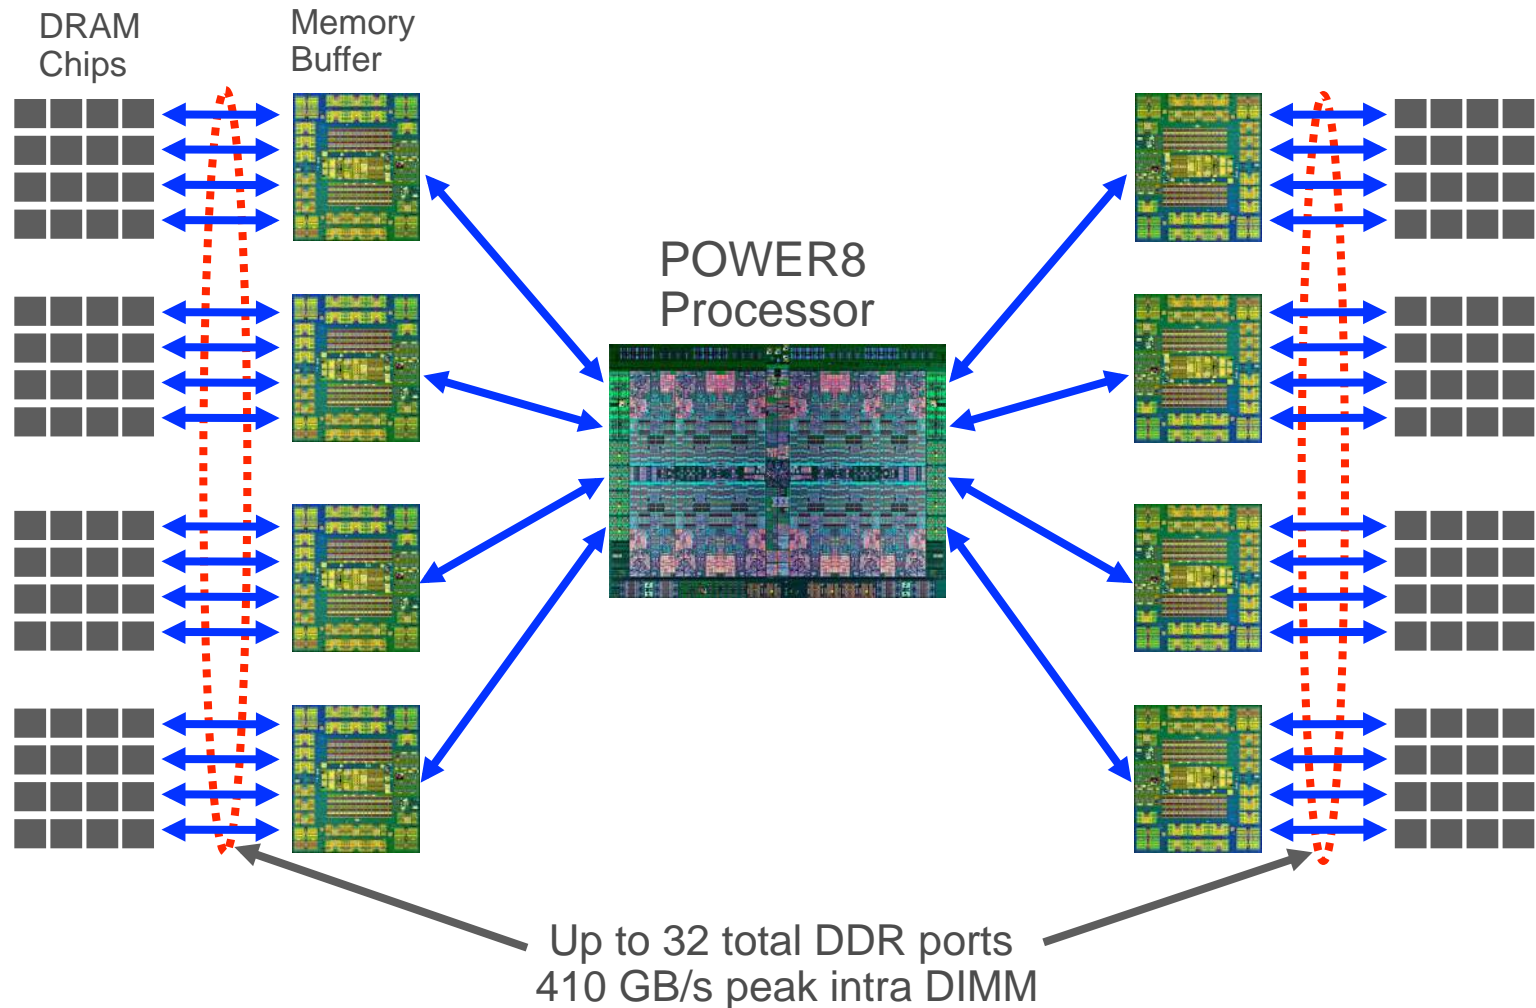
Up to 8 high speed channels,  
Up to 9.6 GB/s per channel  
(2B read / 1B write + cmd)  
for up to 230 GB/s sustained

# POWER8 Memory

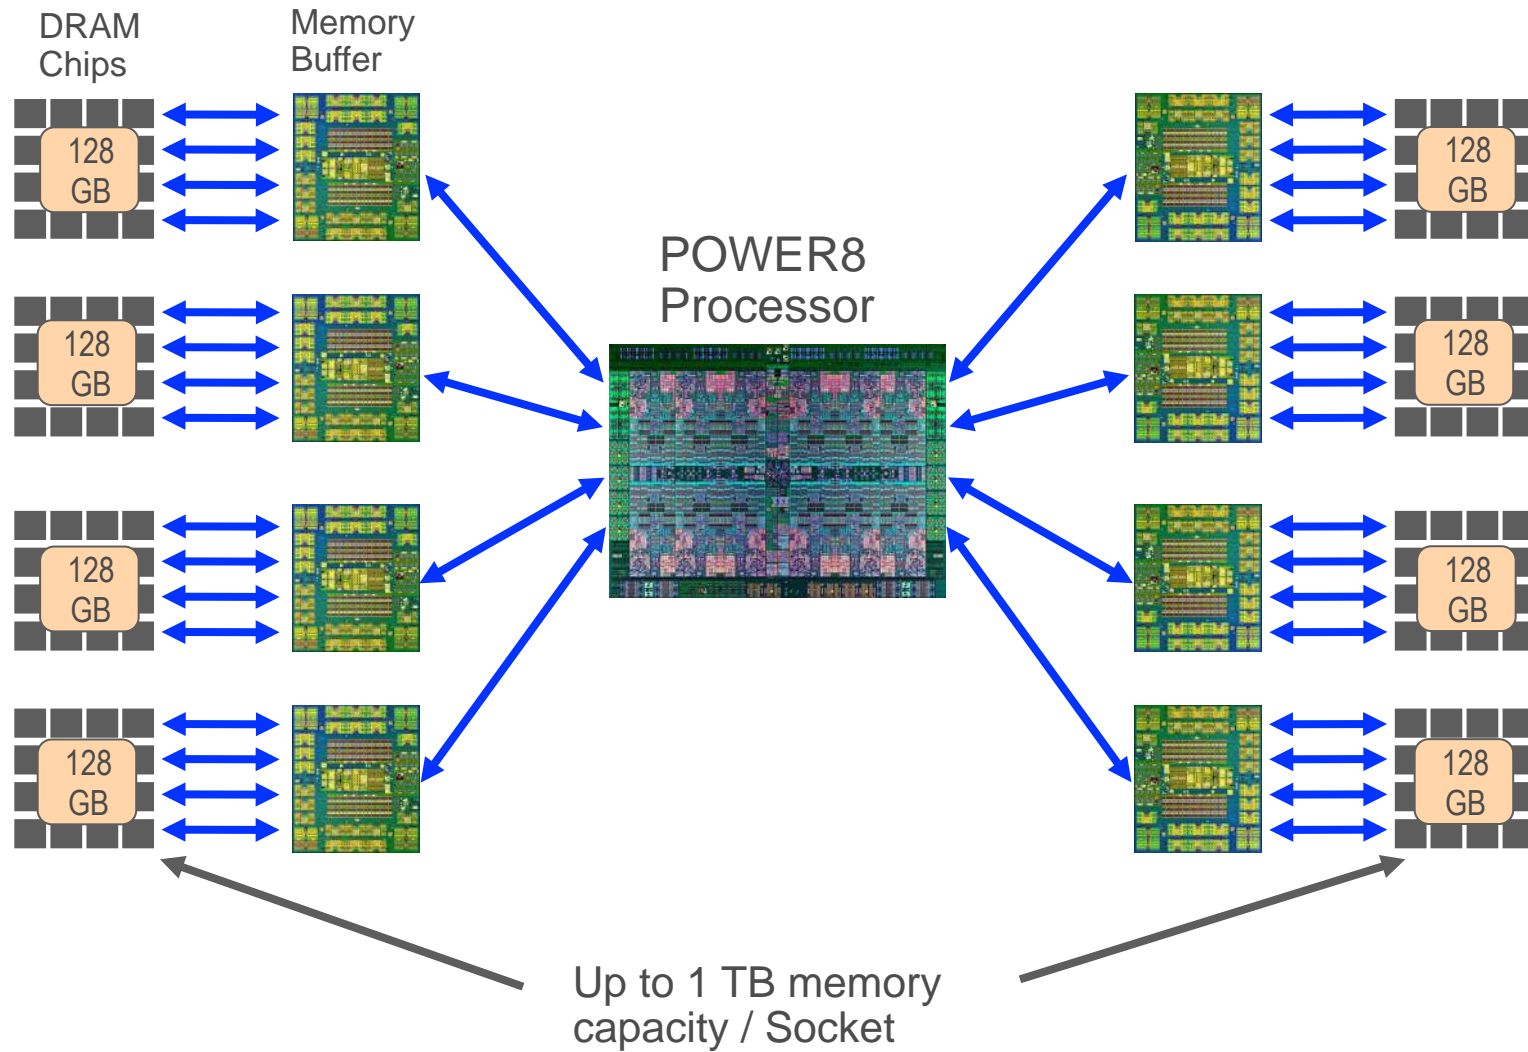





# POWER8 Memory



# POWER8 Memory

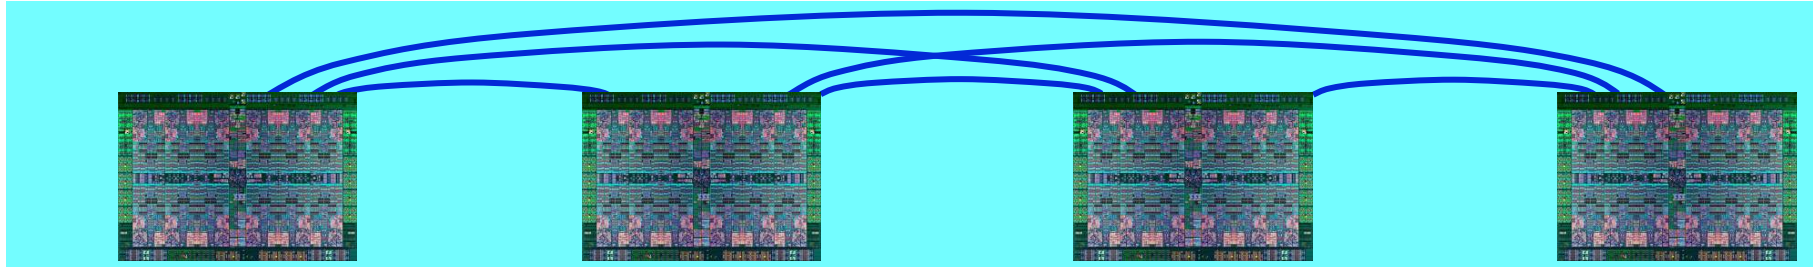



## POWER8 2 Sockets SMP Interconnect



3 12.8 GB/s channels = 38.4GB/s  
+ Spare Lane


# POWER8 4 Sockets SMP Interconnect - Future




 38.4 GB/s

# POWER8 Max SMP Interconnect – Future


48-way Drawer/ Node




192-way SMP system

 38.4 GB/s  12.8 GB/s

# IO Bandwidth




# IO Bandwidth



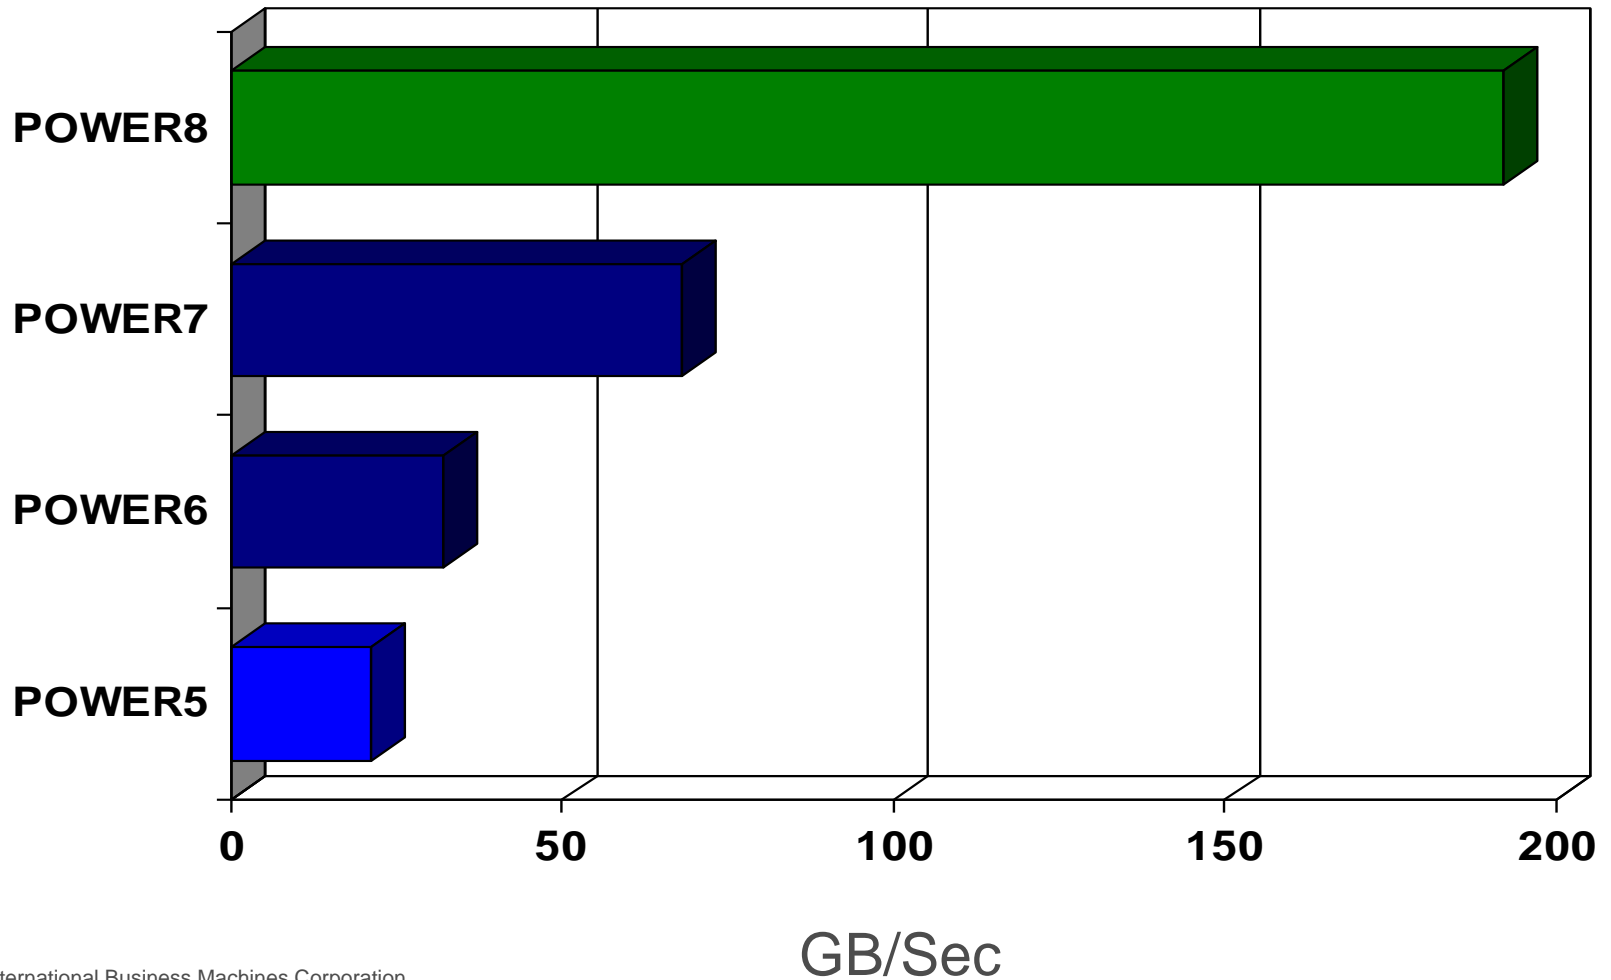
GB/Sec

## Memory Bandwidth per Socket






# Memory Bandwidth per Socket



# POWER8 Integrated PCI Gen 3

## POWER7



GX Bus

I/O Bridge

PCIe G2

PCI Devices


## Native PCIe Gen 3 Support

- Direct processor integration
- Replaces proprietary GX/Bridge
- Low latency
- Gen3 x16 bandwidth (32 GB/s)

## Transport Layer for CAPI Protocol


- Coherently Attach Devices connect to processor via PCIe
- Protocol encapsulated in PCIe

## POWER8



PCIe G3

PCI Device



# POWER8 CAPI (Coherent Accelerator Processor Interface)

## Virtual Addressing

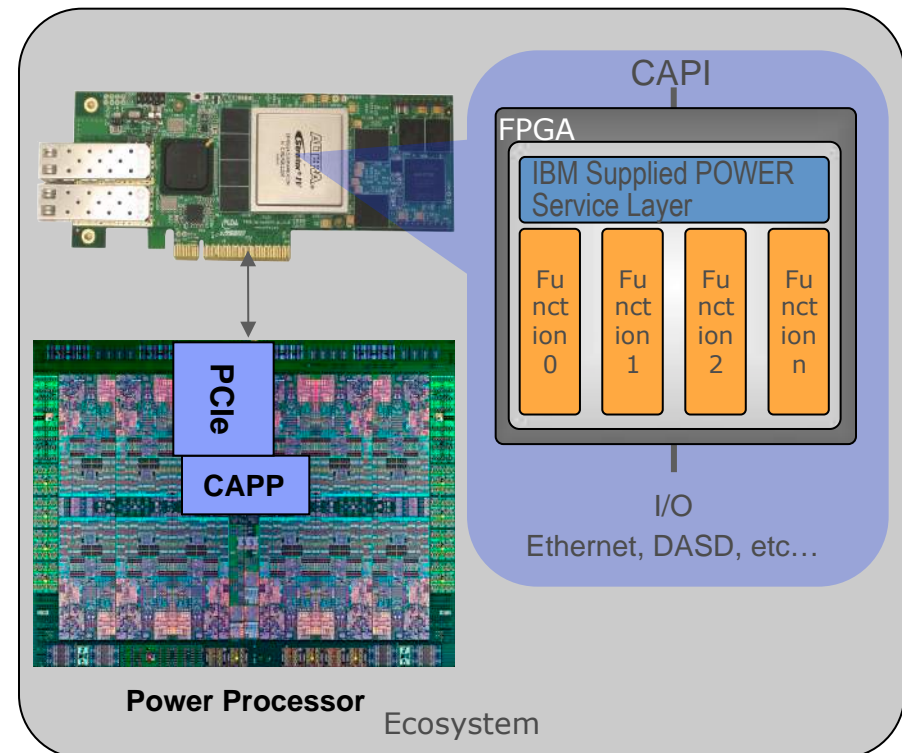
- Accelerator can work with same memory addresses that the processors use
- Pointers de-referenced same as the host application
- Removes OS & device driver overhead

## Hardware Managed Cache Coherence

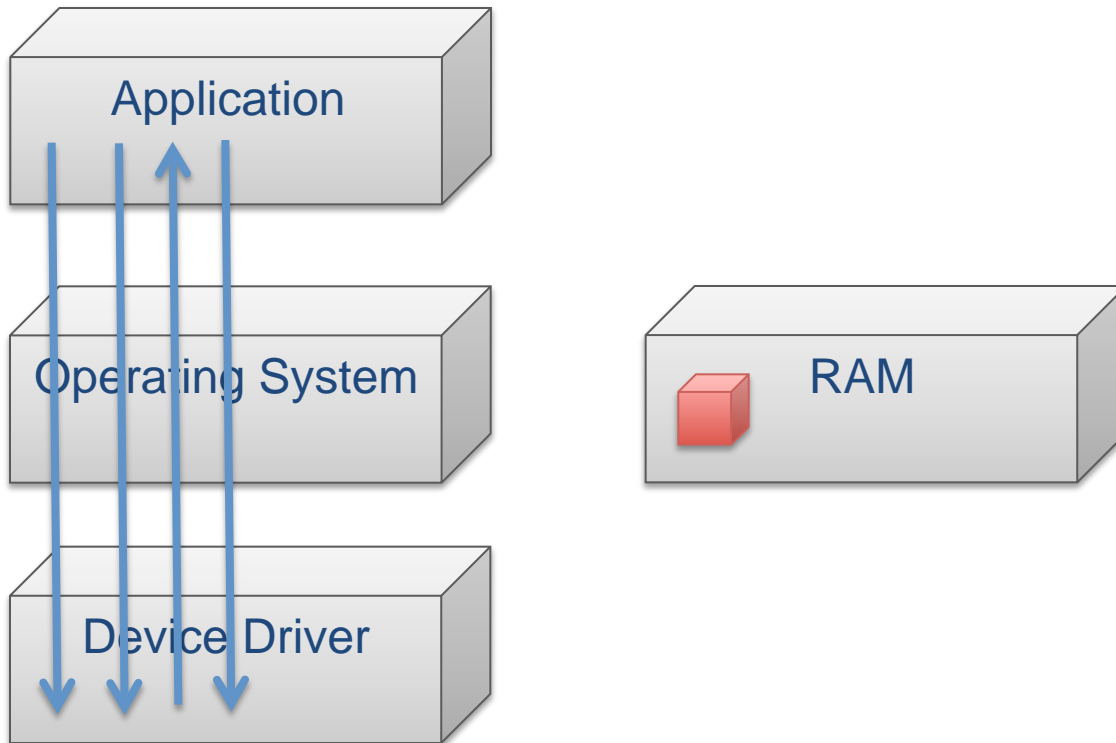
- Enables the accelerator to participate in “Locks” as a normal thread
- Lowers Latency over IO communication model

## Customizable Hardware Application Accelerator

- Specific system SW, middleware, or user application
- Written to durable interface provided by PSL



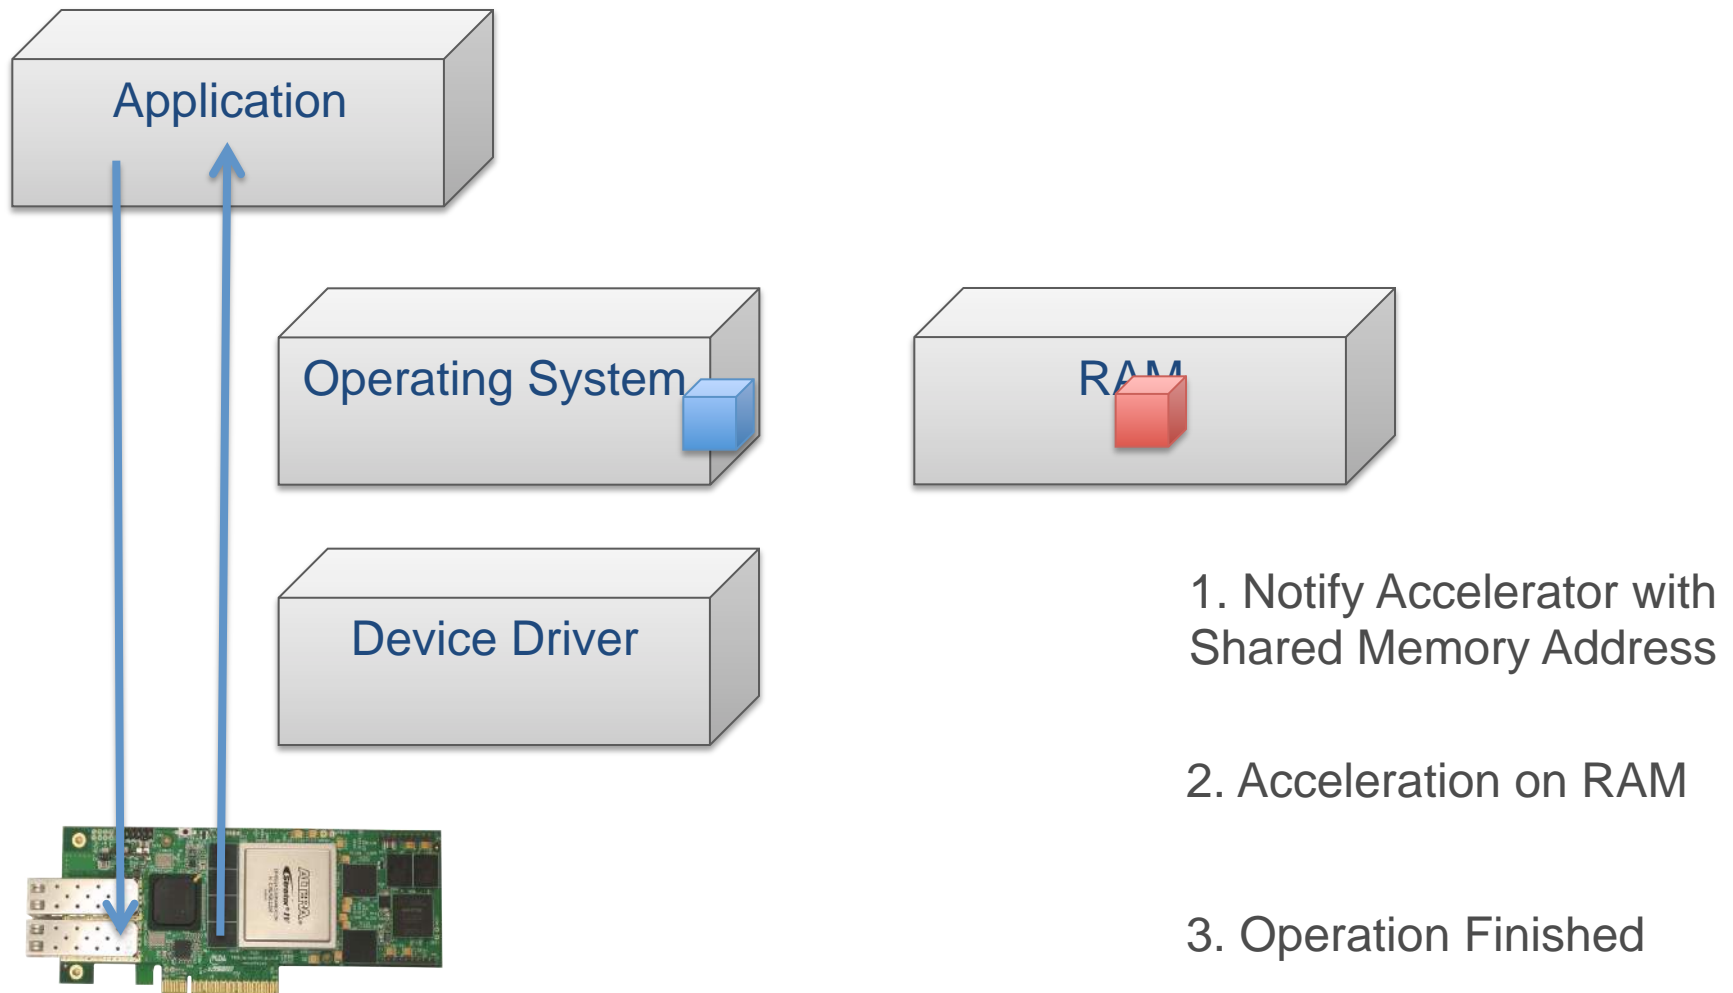
## Traditional Accelerator Behavior




1. Device driver Call
2. Move data to Accelerator
3. Operation Start Interrupt
4. Acceleration
5. Operation Finish Interrupt
6. Get data from Accelerator
7. Device Driver Close




# Coherent Accelerator Behavior




# Coherent Accelerator Processor Interface (CAPI) Overview



## Typical I/O Model Flow



## Flow with a Coherent Model



## Advantages of Coherent Attachment Over I/O Attachment

- **Virtual Addressing & Data Caching**
  - Shared Memory
  - Lower latency for highly referenced data
- **Easier, More Natural Programming Model**
  - Traditional thread level programming
  - Long latency of I/O typically requires restructuring of application
- **Enables Applications Not Possible on I/O**
  - Pointer chasing, etc...

# Advantages of Coherent Connected Accelerator

## ▪ Virtual Addressing

- Removes the requirement for pinning system memory for PCIe transfers
  - Eliminates the copying of data into and out of the pinned DMA buffers
  - Eliminates the operating system call overhead to pin memory for DMA
- Accelerator can work with same addresses that the processors use
  - Pointers can be de-referenced same as the host application
    - Example: Enables the ability to traverse data structures

## ▪ Coherent Caching of Data

- Enables an accelerator to cache data structures
- Enables Cache to Cache transfers between accelerator and processor
- Enables the accelerator to participate in “Locks” as a normal thread

## ▪ Elimination of Device Driver

- Direct communication with Application
- No requirement to call an OS device driver or Hypervisor function for mainline processing

## ▪ Enables Accelerator Features not possible with PCIe

- Enables efficient Hybrid Applications
  - Applications partially implemented in the accelerator and partially on the host CPU
- Visibility to full system memory
- Simpler programming model for Application Modules

# PCIe Java Gzip Compression FPGA

## Transparent zlib HW acceleration integrated into JAVA 7.1

### Corsa 5 A7

- Altera Stratix V PCIE Gen 2/3 card
- Currently supported as PCIe adapter - CAPI support possible in the future
- Initial offering with gzip use case for Java
- Dual SFP+ cages for networking - not used
- 8 GB DDR3 1600 SDRAM
- Full Card Level RAS

### Attributes

- FPGA base accelerator card
- PCIe gen2 x8, ranges between 25W & 40W depending on application
  - ~21W for gzip application
- Follows Altera Stratix roadmap

### PCIe base I/O adaptor

- JAVA 7.1 gzip acceleration use case
- Supports multi-threaded environment



### OS Support

- AIX 7.1
- RHEL 7.0 with patches. Full distro support on RHEL 7.1
- SLES 12 later in 2014 for LE support

### Programming model

- Altera Quartus 13.1
  - OpenCL Power run time support

### RAS

- Follows same model as Z-EDC use case
- Full error detection scheme
- This application not reprogrammable by customers



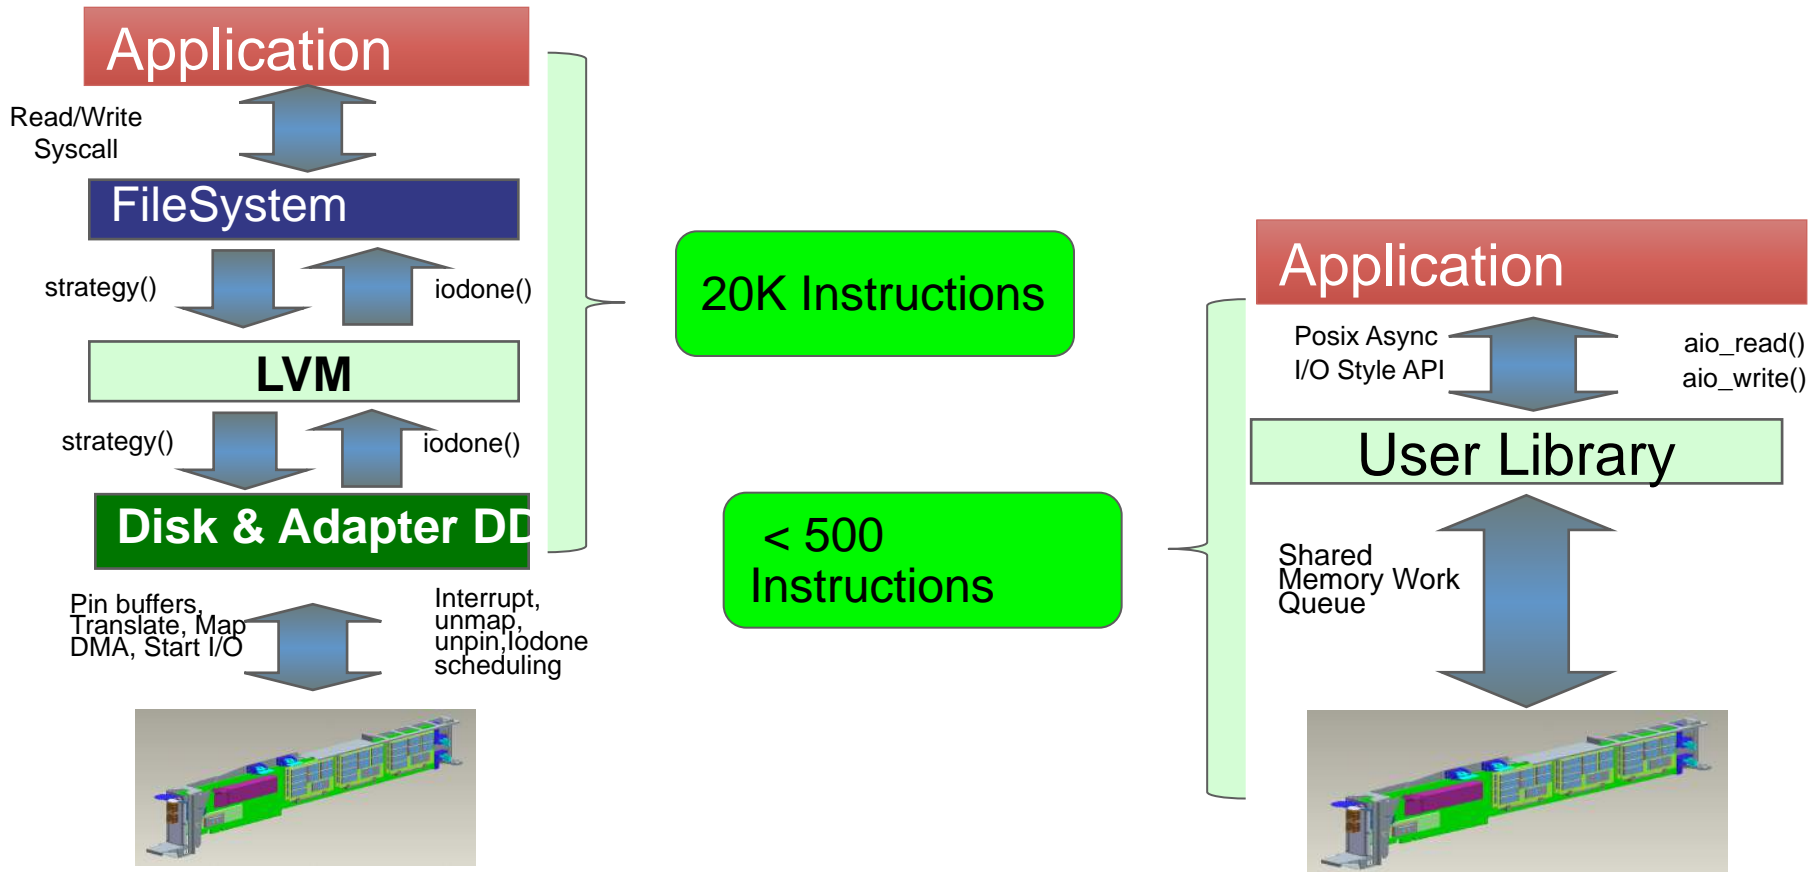
## PCIe Java GZIP FPGA

- Compression tests


File	Original File Size	Compressed File Size	Compress Ratio	zlib Inflate Time (ms)	FPGA Inflate Time (ms)	FPGA Throughput	FPGA Speedup
asyoulik.txt	125179	59320	2.11	0.59	0.040	3.16 GB/s	14.99X
alice29.txt	152089	65188	2.33	0.67	0.047	3.21 GB/s	14.06X
lcet10.txt	426754	172770	2.47	1.77	0.129	3.31 GB/s	13.75X
plravn12.txt	481861	241161	2.0	2.33	0.147	3.28 GB/s	15.87X
ptt5	513216	67529	7.6	1.69	0.150	3.42 GB/s	11.27X
pi.txt	1000000	662248	1.51	4.67	0.334	2.99 GB/s	13.98X
world192.txt	2473400	840481	2.94	9.01	0.739	3.35 GB/s	12.20X
bible.txt	4047392	1421406	2.85	15.17	1.205	3.36 GB/s	12.58X
e.coli	4638690	1891365	2.45	16.76	1.391	3.33 GB/s	12.04X

# CAPI Attached Flash Optimization - planned


- Attach FlashSystem to POWER8 via CAPI coherent Attach
- Issues Read/Write Commands from applications to eliminate 97% of instruction path length
  - CAPI Flash controller Operates in User Space
- Saves 10 Cores per 1M IOPs




# CAPI Flash vs Fibre Channel Flash



IOPS per HW thread →



Latency [us] →



**Processor and Memory**

- Up to 12-core POWER8 processor card
- Up to 8 Simultaneous Multi Thread per core
- Transactional Memory
- Java Code Optimization w/HW Assist

**I/O innovation**


- PCIe Gen3
- Coherent Accelerator Processor Interface

**RAS**

- RAID 0, 5, 6, 10 in the base
- Concurrent maintenance PCIe Gen3 slots
- Standard redundant power/cooling

**POWER Scale-out Systems - Created for New Workloads & Enhanced Core Values**

- ✓ Provide efficient consolidation of business workloads
- ✓ Designed for demanding environments where uptime is critical and downtime is costly
- ✓ Deliver a higher value, more secure, and more reliable platform than competitors




**Power Systems S812L**

- 1-socket, 2U
- Up to 12 cores POWER 8
- 512 GB memory
- 6 PCI Gen3 slot
- Linux only
- Up to One CAPI support
- PowerVM or PowerKVM




**Power Systems S822L**

- 2-socket, 2U
- Up to 24 cores POWER8
- 1 TB memory
- 9 PCI Gen3 slot
- Linux only
- Up to 2 CAPI support
- PowerVM or PowerKVM




**Power Systems S822**

- 2-socket, 2U
- Up to 20 cores POWER8
- 1 TB memory
- 9 PCI Gen3
- AIX & Linux
- Up to 2 CAPI support
- PowerVM



**Power Systems S814**

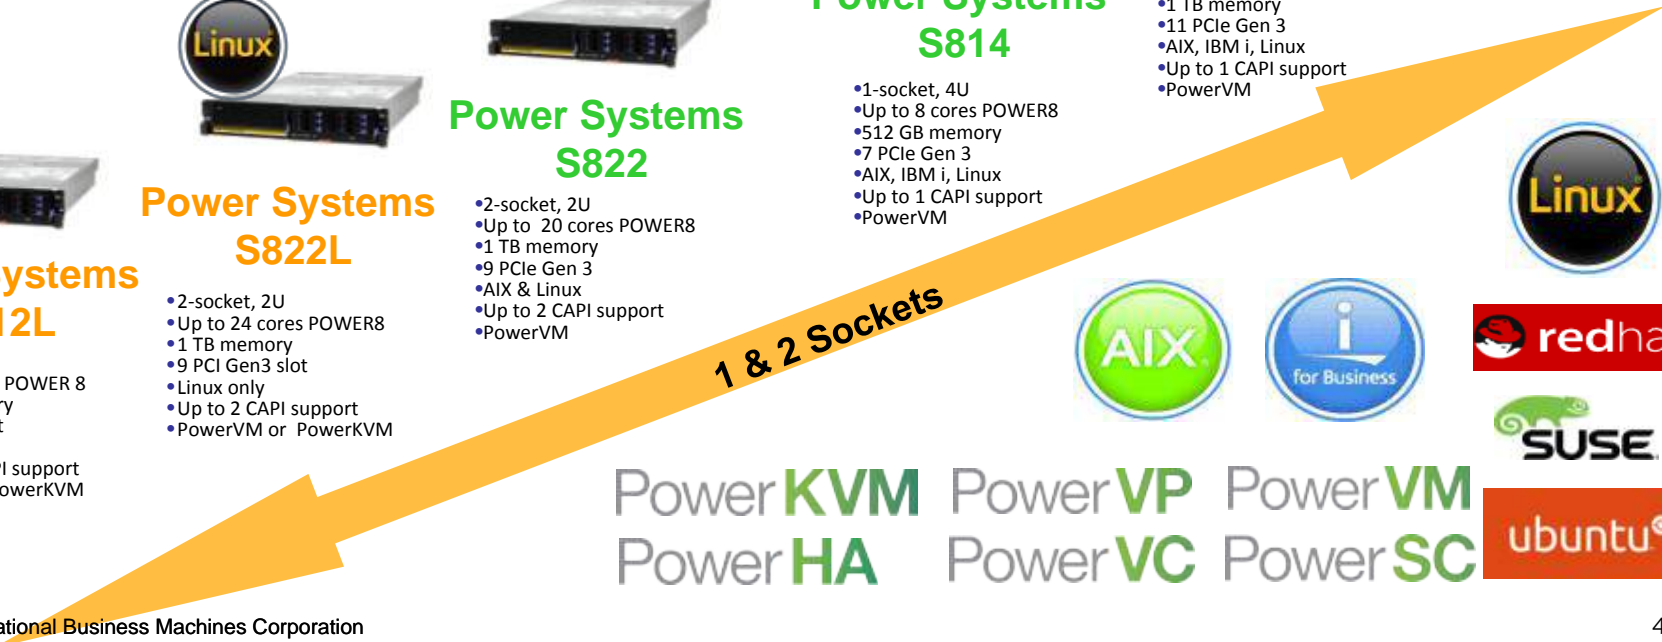
- 1-socket, 4U
- Up to 8 cores POWER8
- 512 GB memory
- 7 PCI Gen 3
- AIX, IBM i, Linux
- Up to 1 CAPI support
- PowerVM

**New 4-core**



**Power Systems S824**

- 2-socket, 4U
- Up to 24 cores POWER8
- 1 TB memory
- 11 PCI Gen 3
- AIX, IBM i, Linux
- Up to 1 CAPI support
- PowerVM



Power **KVM** Power **VP** Power **VM**  
 Power **HA** Power **VC** Power **SC**



AIX  
 for Business  
 Linux  
 redhat  
 SUSE  
 ubuntu®



# Servers

2U Servers



## 1&2 Socket Servers

### New Scale-out Servers with POWER8 technology

- 1 socket : 4U S814
- 2 socket: 2U and 4U S822 and S824

### Linux-only Power Systems (Not called “PowerLinux”)

- 1 socket: 2U S812L
- 2 socket: 2U S822L

1 Socket	1 Socket	2 Socket Option 1S	2 Socket Option 1S	2 Socket
S814	S812L	S824	S822	S822L
8286-41A	8247-21L	8286-42A	8284-22A	8247-22L
4U	2U	4U	2U	2U


Order number = “number EIA/U” then “socket”

For example 8286 - 4 2 A




## Power S812L

- Form Factor: 2U
- Sockets: 1S
- Single Socket
  - Cores: 10 / 12
  - Memory: Up to 512 GB / 8 DIMMs
  - Slots: 6 PCI Gen3 LP (Hotplug)
- Ethernet: Quad 1 Gbit / (x8 slot)
- Media: DVD
- Internal Storage
  - Up to 12 SFF Bays
  - Up to 8 SFF Bays & Up to 6 1.8" SSD Bays with Easy Tier
- Hypervisor: PowerVM or PowerKVM




# Power S812L




## POWER S822 & Power S822L

- Sockets: 2S ( 1S upgradeable )
- Single Socket
  - Cores: 6 or 10
  - Memory: Up to 512 GB / 8 DIMMs
  - Slots: 6 PCI Gen3 LP (Hotplug)
- Dual Sockets
  - Cores: 12 or 20 ( 2 x 6 or 10 )
  - Memory: Up to 1 TB / 16 DIMMs
  - Slots: 9 PCIe Gen3 LP (Hotplug)
- SRIOV support (SoD)
- Ethernet: Quad 1 Gbit / (x8 slot)
- Native IO: USB (2), Serial (2), HMC (2)
- Internal Storage
  - DVD
  - Up to 12 SFF Bays
  - Optional: Split Backplane: 6 + 6
  - Optional: Up to 8 SFF Bays & Up to 6 1.8" SSD Bays with Easy Tier
- Power Supplies: 2 ( 200-240 AVC)





# Power S822 - Storage Backplanes





**Backplane FC EJ0T - 12 SFF HDD/SSD Bays SFF-3**

- Optional Split Disk Backplane (FC EJ0V) – 6 + 6 Disks

Op-panel + 2 USB Ports  
DVD

SSD CAGE (FC EJTL)  
(located behind the bezel)




**Backplane FC EJ0U - 8 SFF HDD/SSD Bays SFF-3 + RAID 0,5,6,10,5T2,6T2,10T2**

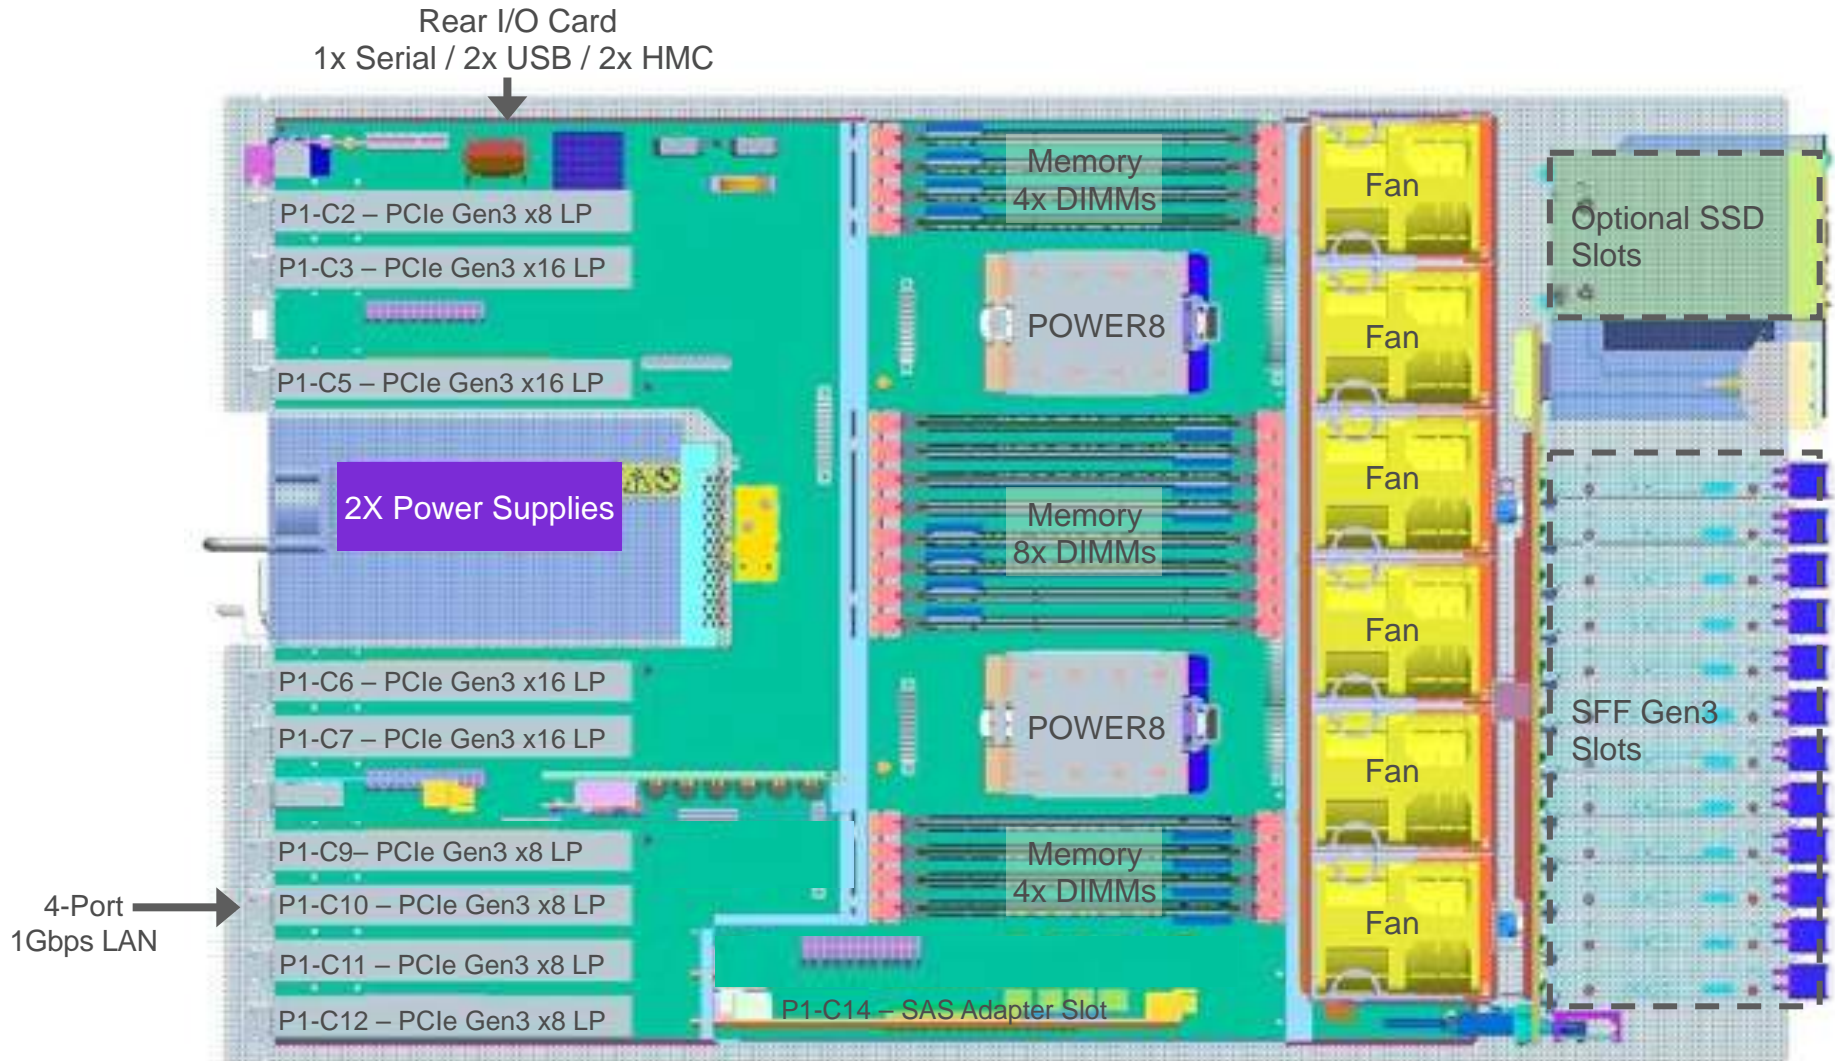
- Optional 6-disk 1.8-inch SSD CAGE (FC EJTL)

Op-panel + 2 USB Ports  
DVD

# Power S822 – SSD Cage



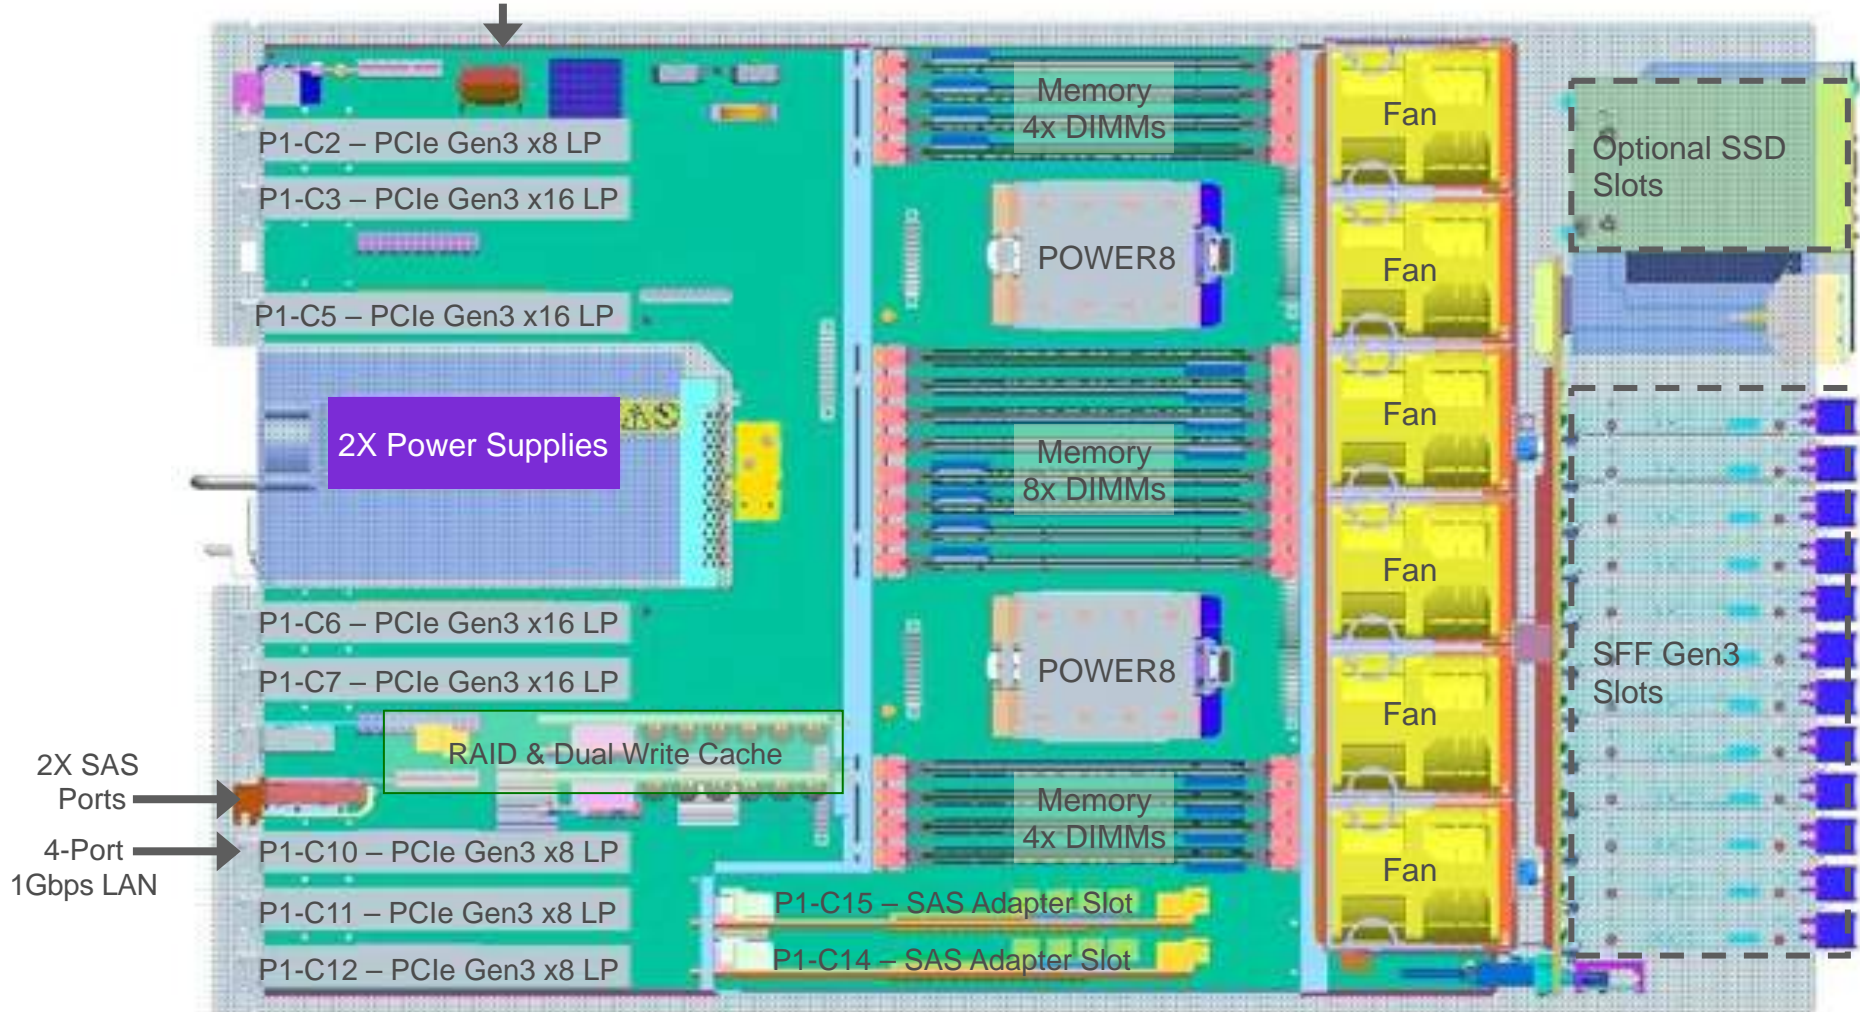
# POWER S822 – Slots vs Standard I/O Backplane






# POWER S822 – Slots vs High-Performance I/O Backplane

Rear I/O Card  
Serial / 2x USB / 2x HMC



# Power S822 – Back view



2X USB  
Serial Port  
HMC Ports

Dual Power Supplies

SAS Port  
or PCIe x8



## Power S822 compared to Power 730

	<i>Power 730</i>	<b>Power System S822</b>
<i>Processor</i>	<i>POWER7+</i>	<i>POWER8</i>
<i>Sockets</i>	<i>2</i>	<i>1 (Upgradeable) / 2</i>
<i>Cores</i>	<i>8 / 12 / 16</i>	<i>6 / 10 or 12 / 20</i>
<i>Maximum Memory</i>	<i>512 MB @ 1066 MHz</i>	<i>256 / 512 GB @ 1600 MHz</i>
<i>Memory Cache</i>	<i>No</i>	<i>Yes</i>
<i>Memory Bandwidth</i>	<i>136 GB/sec</i>	<i>192 / 384 GB/sec</i>
<i>Memory DRAM Spare</i>	<i>No</i>	<i>Yes</i>
<i>IO Expansion Slots</i>	<i>Dual GX++</i>	<i>SoD</i>
<i>PCIe slots</i>	<i>5 PCIe x8 LP</i>	<i>4 / 5 PCIe x8 LP 2 / 4 PCIe x16 LP</i>
<i>CAPI (Capable slots)</i>	<i>N / A</i>	<i>One / Two</i>
<i>PCIe Hot Plug Support</i>	<i>No</i>	<i>Yes</i>
<i>IO bandwidth</i>	<i>40 GB/sec</i>	<i>192 GB/sec</i>
<i>Ethernet ports</i>	<i>Quad 1 Gbt</i>	<i>Quad 1 Gbt</i>
<i>SFF</i>	<i>6</i>	<i>12</i>
<i>Easy Tier Support</i>	<i>No</i>	<i>Yes</i>
<i>Integrated split backplane</i>	<i>No</i>	<i>Yes ( 6 + 6 )</i>
<i>SRIOV Support</i>	<i>No</i>	<i>SoD</i>
<i>Service Processor</i>	<i>Generation 1</i>	<i>Generation 2</i>






# Servers

4U Servers




## Power S814

- Form Factor: 4U or Tower
- Sockets: 1S
- Single Socket
  - Cores: 4, 6 or 8
  - Memory: Up to 512 GB / 8 DIMMs
  - Slots: 6 PCI Gen3 LP (Hotplug)
- SRIOV support (SoD)
- Ethernet: Quad 1 Gbit / (x8 slot)
- Native IO: USB (3), Serial (2), HMC (2)
- Internal Storage
  - DVD
  - Up to 12 SFF Bays
  - Optional: Split Backplane: 6 + 6
  - Optional: Up to 18 SFF Bays & Up to 8 1.8" SSD Bays with Easy Tier
- Power Supplies: 2 ( 110 or 200-240 AVC)
- 4-core specially attractive to IBM i customers because of P05 pricing tier



# Power S814






## 4-Core S814

- Four-core server for AIX, IBM i and Linux.
- Especially interesting to IBM i clients because of its P05 software tier
- 3 GHz cores
- Like the POWER7 4-core offering – no feature code I/O drawers like #5887 and no future PCIe expansion drawer. (Fibre Channel attached drawers/SANs supported.)
- Like the POWER7 4-core – **64GB memory maximum**. 4-core POWER8 uses only 16GB memory features
- Max of ten SFF-3 drives in system unit (SSD or disk or mix) -- POWER7 4-core maximum is eight. POWER8 4-core is 25% more drives as of July.
- Supports 15k 139/146GB and 15k 283/300GB disk drives and supporting 387GB SSD. Not supporting other drive sizes.
- 4-core S814 offers better price/performance than 4-core 720 and higher function, but the absolute price of S814 4-core option higher than 720 4-core

## POWER S824

- Form Factor: 4U
- Sockets: 2S ( 1S upgradeable )
- Single Socket
  - Cores: 6 or 8
  - Memory: Up to 512 GB / 8 DIMMs
  - Slots: 7 PCI Gen3 (Hotplug)
- Dual Sockets
  - Cores: 12, 16, or 24 ( 2 x 6, 8, or 12 )
  - Memory: Up to 1 TB / 16 DIMMs
  - Slots: 11 PCIe Gen3 (Hotplug)
- SRIOV support (SoD)
- Ethernet: Quad 1 Gbit / (x8 slot)
- Native IO: USB (3), Serial (2), HMC (2)
- Internal Storage
  - DVD
  - Up to 12 SFF Bays
  - Optional: Split Backplane: 6 + 6
  - Optional: Up to 18 SFF Bays & Up to 8 1.8" SSD Bays with Easy Tier
- Power Supplies: 4 ( 200-240 AVC)



3 Yr Warranty

# Power S824 - Storage Backplanes

Backplane FC EJ0N - 12 SFF HDD/SSD Bays SFF-3

- Optional Split Disk Backplane (FC EJ0S) – 6 + 6 Disks



P2-D1

P2-D2

P2-D3

P2-D4

P2-D5

P2-D6

P2-D7

P2-D8

P2-D9

P2-D10

P2-D11

P2-D12


Op-panel + 2 USB Ports  
DVD

# Power S824 – Storage Backplanes

**Backplane FC EJ0P** - 18 SFF HDD/SSD Bays SFF-3 + RAID 0,5,6,10,5T2,6T2,10T2

- Optional 8-disk 1.8-inch SSD CAGE (FC EJTM)

SSD CAGE (FC EJTM)  
(located behind the bezel)



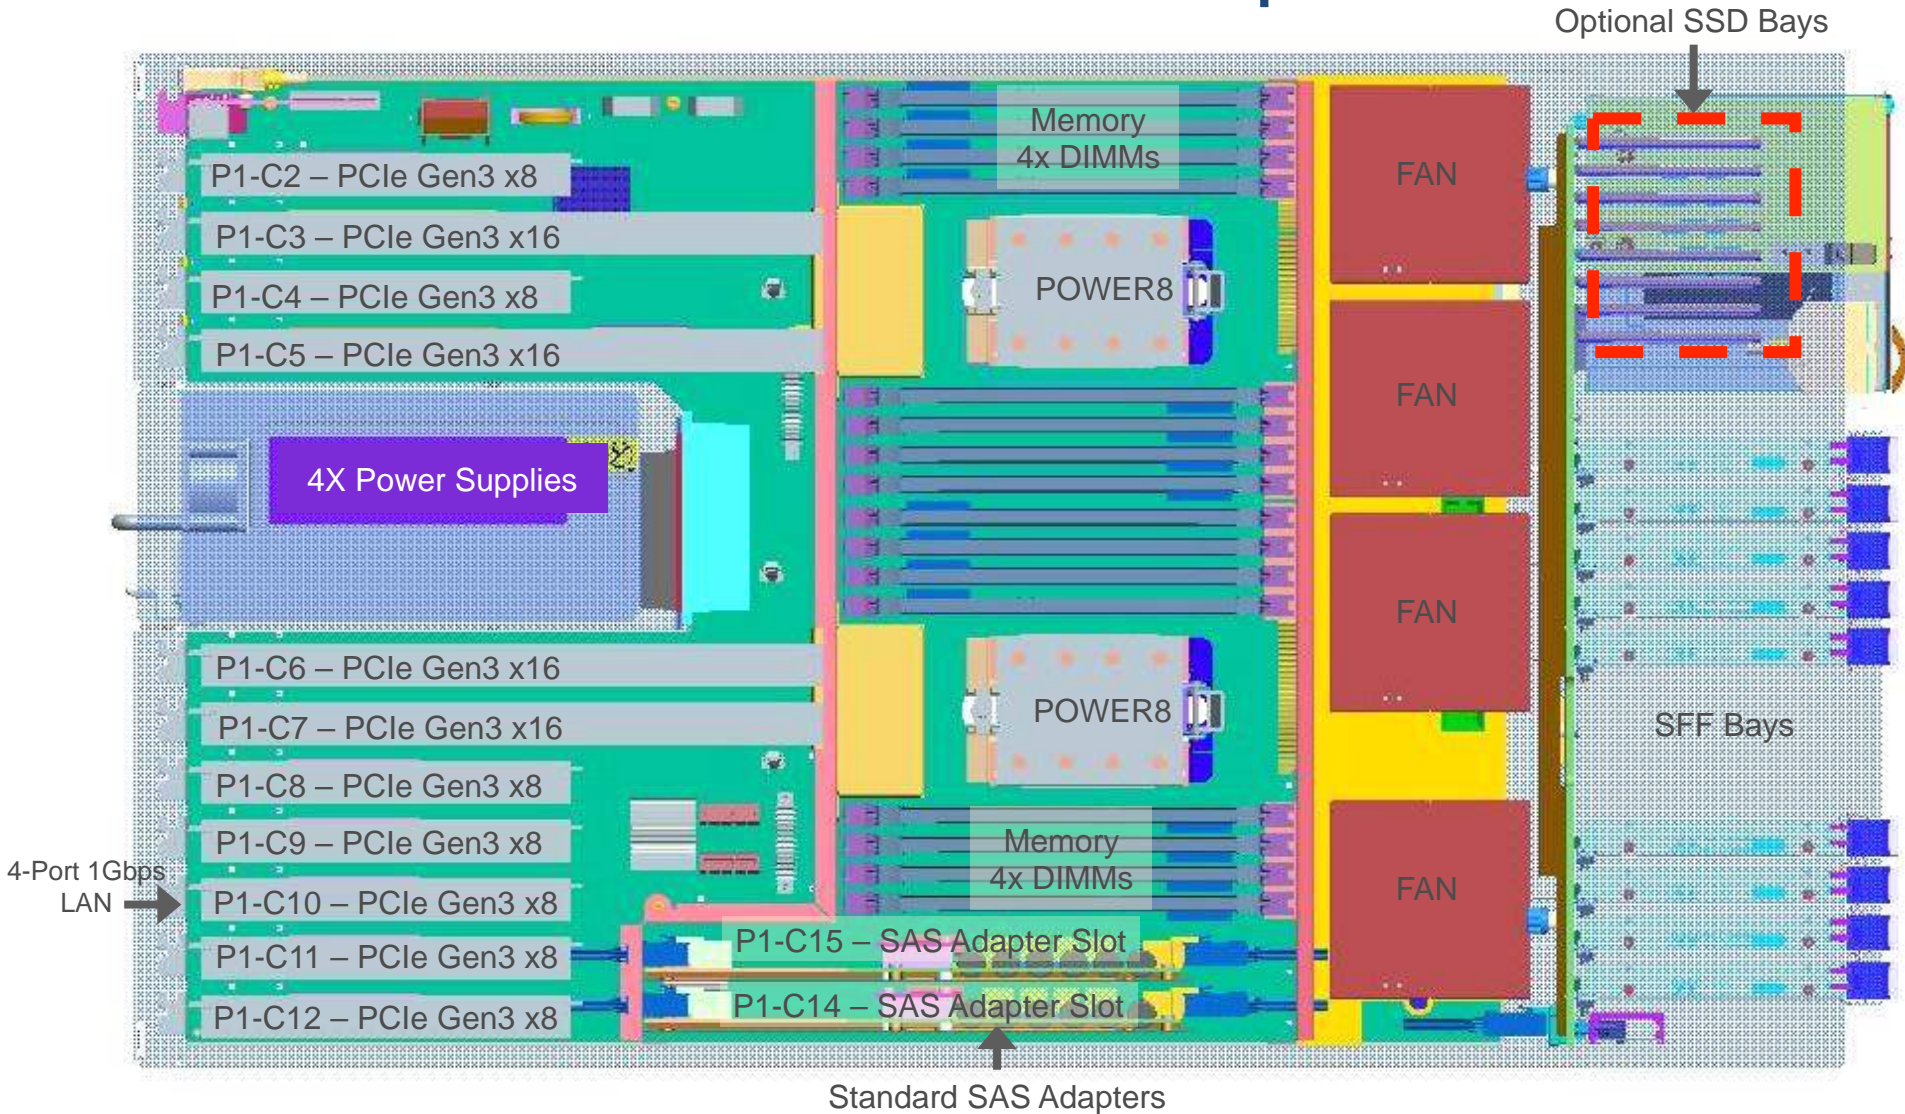
Op-panel + 2 USB Ports  
DVD



## Power S824 – Storage Backplane

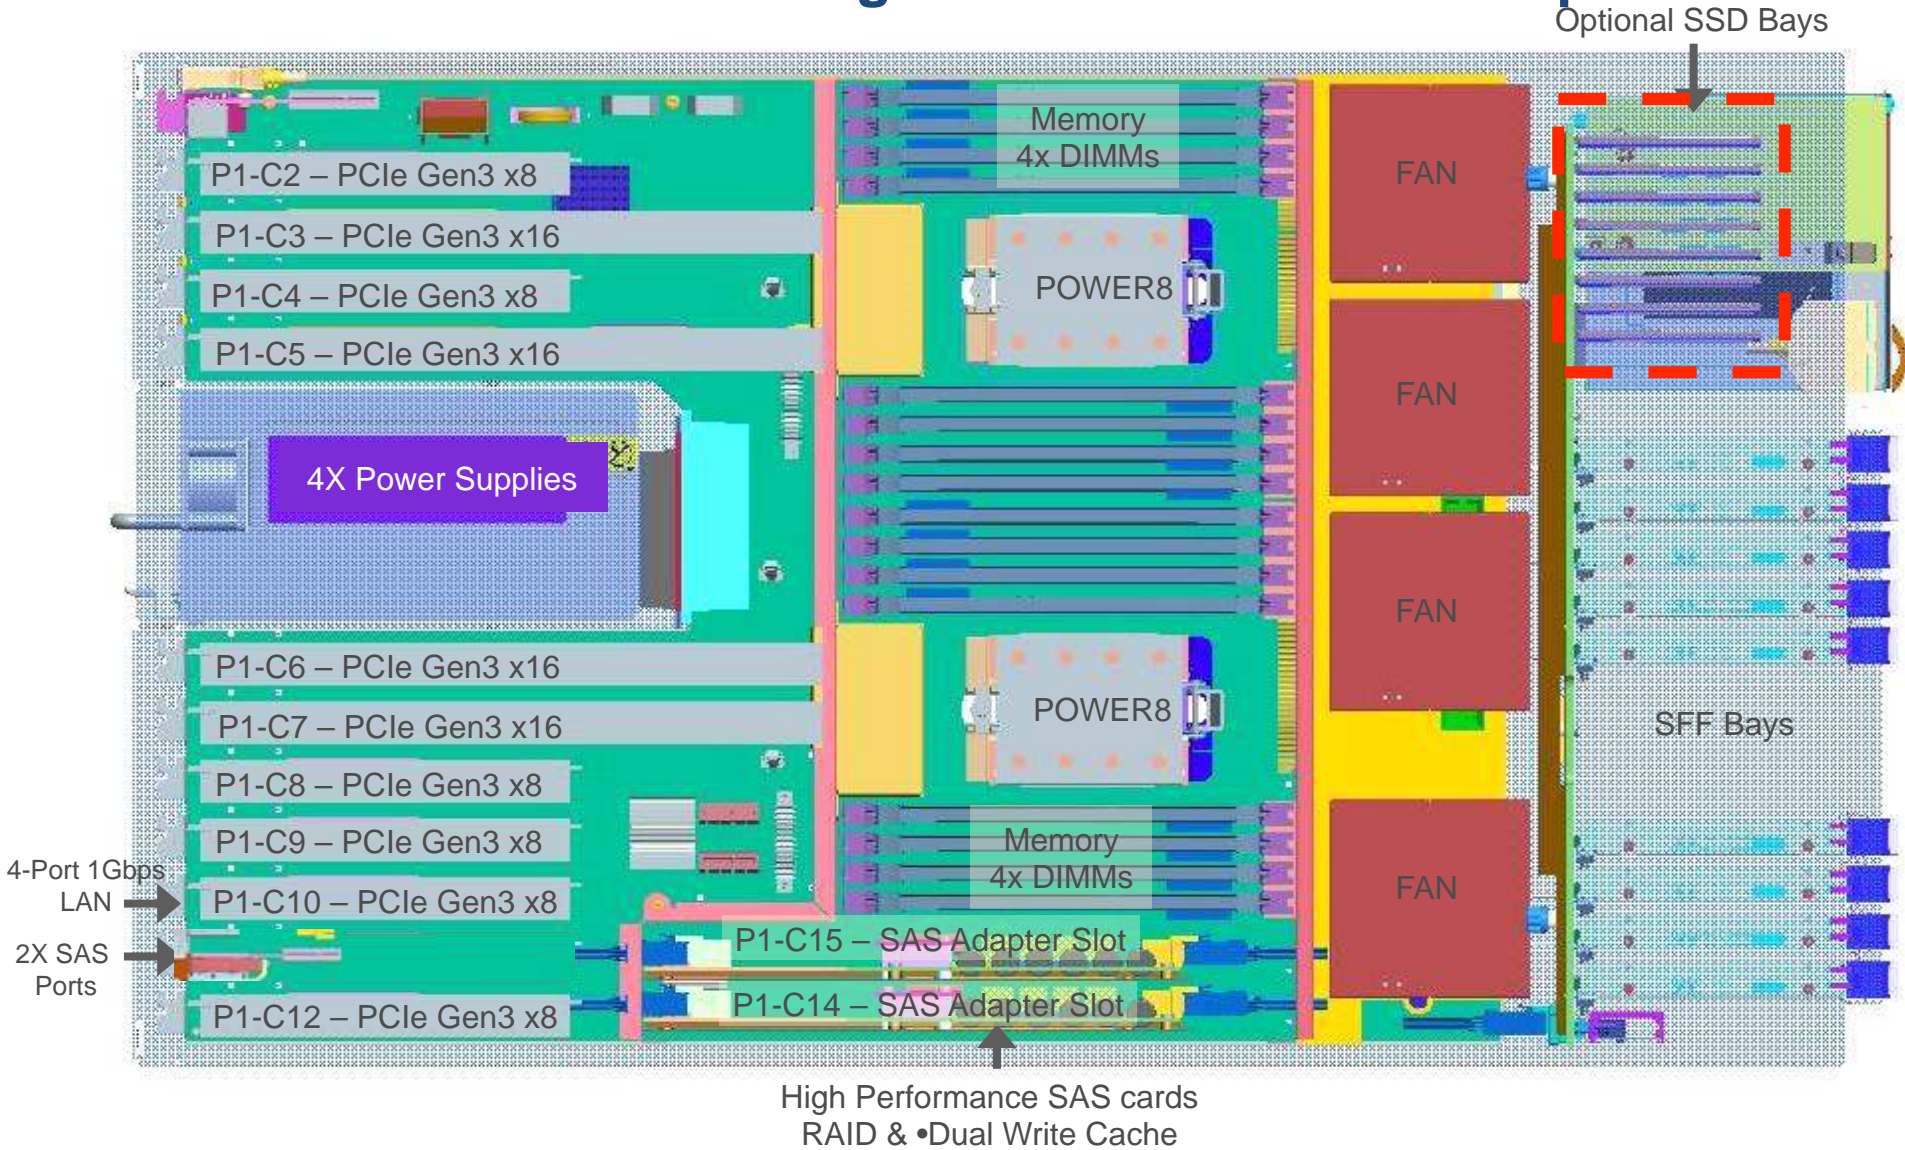


# Power S824 – Slots vs Standard I/O Backplane





# POWER S824 – Slots vs High-Performance I/O Backplane



## Power S814 & Power S824 – Back view



## Power S814 and S824 compared to Power 740

	<b>Power System 740</b>	<b>Power System S814</b>	<b>Power System S824</b>
<b>Processor Offerings</b>	4, 6, or 8 core @ 3.6 GHz  1x or 2x 6-core 4.2 GHz 1x or 2x 8-core 3.6 / 4.2 GHz	6-core 3.02 GHz 8-core 3.69 GHz	6 core @ 3.02 GHz 8 core @ 3.69 GHz 1x or 2x 6 -core@ 3.89 GHz 1x or 2x 8 -core 4.15 GHz 2x 12 -core 3.52 GHz
<b>DDR3 Memory features</b>	8 / 16 / 32 / 64GB 8GB to 256GB / socket	16 / 32 / 64GB 16GB to 512GB	16 / 32 / 64GB 32GB to 1024GB
<b>OS Support</b>	AIX: 5.3 / 6.1 / 7.1 IBM i: 6.1 & 7.1 Linux: RHEL & SLES	AIX: 6.1 & 7.1 IBM i: 7.1 & 7.2 Linux: RHEL & SLES	AIX: 6.1 & 7.1 IBM i: 7.1 & 7.2 Linux: RHEL & SLES
<b>DASD / Bays</b>	Up to 6 HDD or SSD	Up to 12 SFF or SSD	Up to 12 SFF or SSD
<b>PCIe Expansion Slots</b>	Five x8 G2 LP	Two x16 G3 LP Hot Plug Five x8 G3 LP Hot Plug 1x CAPI capable	Four x16 G3 LP Hot Plug Seven x8 G3 LP Hot Plug 2x CAPI capable
<b>Integrated SAS/SATA Cntrl</b>	Standard: RAID 0, 1, & 10 Optional: RAID 5 & 6	Standard RAID 0, 5, 6, 10	Standard RAID 0, 5, 6, 10
<b>Ethernet</b>	Quad 10/100/1000	Quad 10/100/1000	Quad 10/100/1000
<b>Media Bays</b>	1 Slim-line & 1 Half Height ( Optional )	1 Slim-line	1 Slim-line
<b>IO Drawers</b>	0 on 1S , 2 on 2S	N / A	N / A
<b>Power requirement</b>	100V to 240V AC 1S 200 to 240V AC 2S	100V to 240V AC 200V to 240V AC	200V to 240V AC
<b>Red Power &amp; Cooling</b>	Optional / Standard	Optional	Standard
<b>Warranty</b>	3 Years		



# Physical Planning – Tower Acoustic Considerations

S814 Tower considerations: The dB level of the tower meets general business standards and can be installed in an open business environment in addition to a machine room. It may be perceived as noisy to some clients who wish to seat people nearby.

Also note that some PCIe adapters draw more power and require more cooling. The POWER8 server recognizes these adapters and speeds up fans to provide this cooling (airflow). **Faster fans SIGNIFICANTLY increases the noise level.** To avoid potential acoustic situations, these adapters are only supported in the S814 when it is in a rack configuration.

Adapters as of April 2014 requiring more cooling:

- PCIe3 SAS adapters (#EJ0J, #EJ0M, #EL3B, #EJ0L)
- PCIe2 SAS adapters (#ESA1, #ESA2, #ESA3, and #5913)

## Physical Planning – Turbo Mode and Acoustics


If the server is enabled to run Turbo mode, the fans run a bit faster and there is a noticeably increased noise level (5-10 dB estimate). If the server actually enters Turbo mode, the fans speed increase again and the noise increases noticeably again.

Enabling Turbo mode on tower configurations is not recommended for acoustic reasons unless the tower is in a machine room or is some distance away from people in an office environment.

If higher energy PCIe adapters are also present, the fans REALLY significantly increase in speed and in noise level.



# PCIe2 2-port 10/1GbE BaseT RJ45 Adapter



#EN0W (FH) / #EN0X (LP)  
CCIN 2CC4

**Only \$950  
Plus inexpensive cabling**

Supported on POWER8 Systems

Operating System support:

- AIX 6.1 & AIX 7.1 or later
- IBM i via VIOS
- Linux
  - ❖ RHEL 6 or later
  - ❖ SLES 11 or later
- PowerKVM Host
- VIOS

Dual port 10GBase-T RJ45 ports - Ethernet NIC traffic

Ports default to auto negotiate the highest speed:

- 10Gb (10GBaseT), 1Gb (1000BaseT) or 100Mb (100BaseT) full duplex.

Port's configuration are independent of the one other, but both ports owned by one partitions or one VIOS

CAT-6A cabling supported

**NIM/Linux Install support announced July 2014 (planned GA 29 Aug)**

Prices are USA list prices and are subject to change. Reseller prices may vary.

## PCIe3 12GB Cache RAID SAS Adapter Quad-port 6Gb

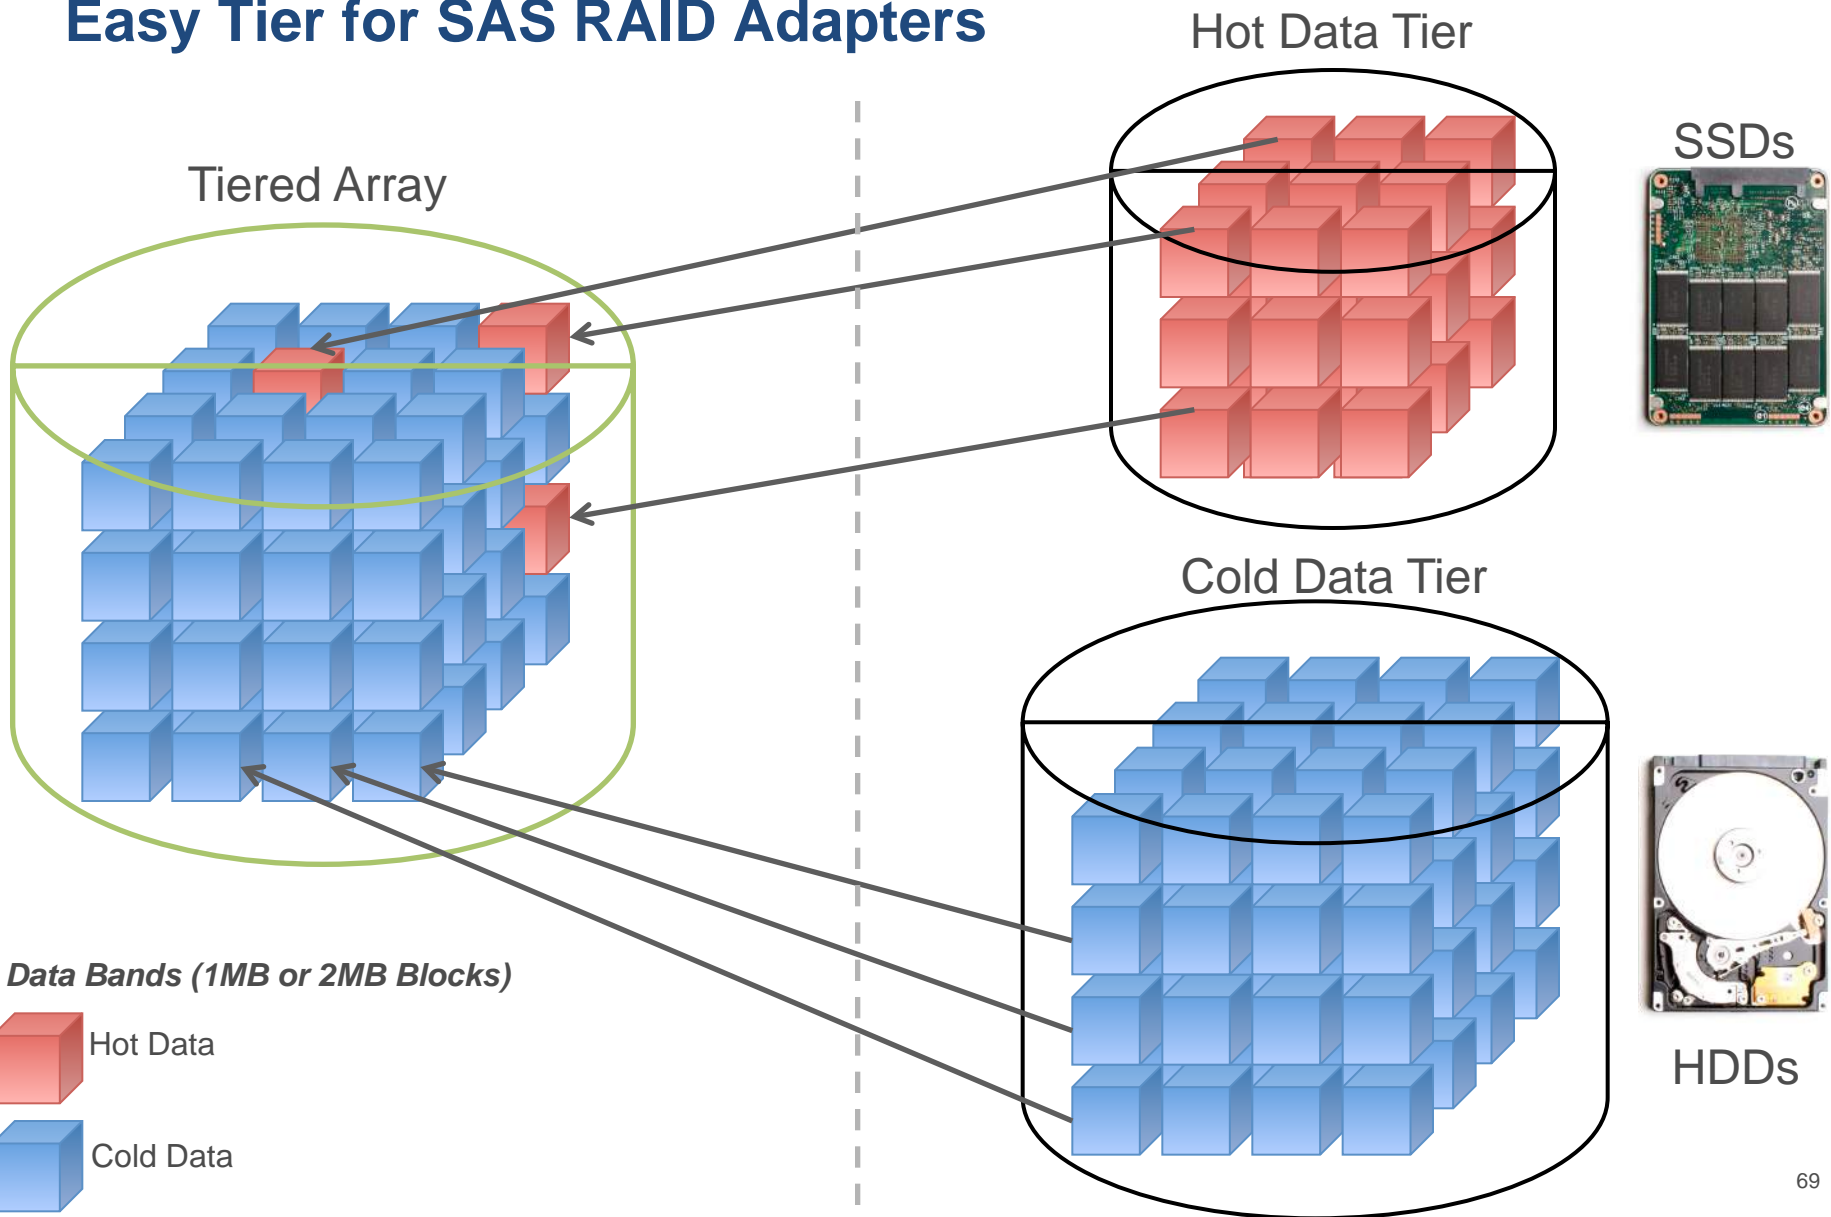
- Dedicated Slots
- 3GB Cache
  - Cache compression
  - Up to 12GB Logical Cache
- Easy Tier
- Redundant Controllers
- External SAS Ports for ONE EXP24S
- RAIDs 0,1,5,6,10 & 5T2, 6T2 & 10T2
- No support for:
  - Split backplane
  - Tape
- Physically capable of more than 1M IOPS on RAID0
  - 900k IOPS on RAID5
- Limited by the number of disks we attach to a server



## Easy Tier for SAS RAID Adapters

- Provides the benefits of Easy Tier and performance advantages of DAS
- Accentuates the Integrated Tier Storage (SSD/HDD) design of POWER8
- The RAID storage subsystem tunes itself based on the workload characteristics of the system to optimize performance
- Adapter dynamically swaps data between tiers to achieve lowest response times and highest throughput
- Swaps happen automatically, without manual user intervention
- Hottest (best) data moved first
- Reacts to workload changes in Real Time, and in seconds/minutes (not hours)
- Improves cost/performance tradeoff between HDDs/SSDs
- GA 6/2014
  - Supported in POWER7+/POWER8 with FC=EJOL / CCIN=57CE (PCIe3 12 GB Cache RAID SAS adapter Quad-port 6 Gb x8)
  - Supported in POWER8 with FC=EJ0P / CCIN=57D8 (PCIe3 x8 Cache SAS RAID internal adapter 6 Gb)
- Supported by AIX, Linux, and VIOS
  - IBM i supported only via VIOS, requires IBM I 6.1 or later and VIOS 2.2.3.2 or later


# Easy Tier for SAS RAID Adapters




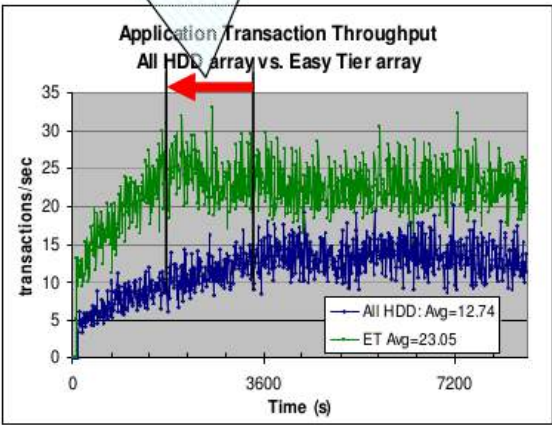
# Easy Tier Performance

AIX Stock Trading Application Performance (RAID 5T2, 3 SSDs, 16 HDDs)

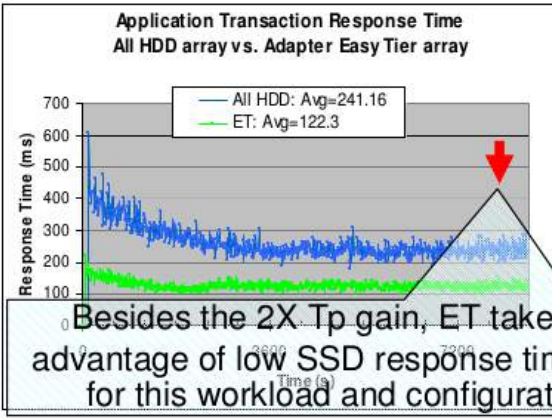
ET array Tp started ~2X higher than all HDD array @ application's start



ET array reached peak throughput ~24 minutes sooner



At steady state ET array Tp was ~2X all HDD Array



Besides the 2X Tp gain, ET takes full advantage of low SSD response times, 2X for this workload and configuration.



# Creating a Tiered Array

```
IBM SAS Disk Array Manager

Move cursor to desired item and press Enter.

List SAS Disk Array Configuration
Create an Array Candidate pdisk and Format to RAID block size
Create a SAS Disk Array
Delete a SAS Disk Array
Add Disks to an Existing SAS Disk Array
Migrate an Existing SAS Disk Array to a New RAID Level
Configure a Defined SAS Disk Array
Change/Show Characteristics of a SAS Disk Array
Manage HA Access Characteristics of a SAS Disk Array
Reconstruct a SAS Disk Array
Change/Show SAS pdisk Status
Diagnostics and Recovery Options

Select a RAID Level

Move cursor to desired item and press Enter.

 0
 5
10
 6
5T2
6T2
10T2

F1=Help      F2=Refresh      F3=Cancel
F8=Image     F10=Exit        Enter=Do
/=Find       n=Find Next

F1=Help
F9=Shel
```

## Tiered arrays on AIX

Name	Resource	State	Description	Size
sissas1	FEFFFFFF	Primary	PCIe3 12GB Cache RAID SAS Adapter Quad-port 6Gb x8	
sissas0	FEFFFFFF	HA Linked	Remote adapter SN 00325001	
hdisk1	FC0000FF	Optimal	RAID 5T2 Array (N/N)	773.5GB
pdisk0	000400FF	Active	Array Member	139.6GB
pdisk1	000401FF	Active	Array Member	139.6GB
pdisk2	000402FF	Active	Array Member	139.6GB
pdisk3	000403FF	Active	Array Member	139.6GB
pdisk7	000407FF	Active	SSD Array Member	177.8GB
pdisk6	000406FF	Active	SSD Array Member	177.8GB
pdisk8	000408FF	Active	SSD Array Member	177.8GB
hdisk2	FC0100FF	Optimal	RAID 6T2 Array (N/N)	1090GB
pdisk10	00040AFF	Active	SSD Array Member	387.9GB
pdisk11	00040BFF	Active	SSD Array Member	387.9GB
pdisk4	000404FF	Active	Array Member	139.6GB
pdisk20	000414FF	Active	SSD Array Member	387.9GB
pdisk21	000415FF	Active	SSD Array Member	387.9GB
pdisk9	000409FF	Active	SSD Array Member	177.8GB
pdisk5	000405FF	Active	Array Member	139.6GB
pdisk12	00040CFF	Active	Array Member	139.6GB
pdisk13	00040DFF	Active	Array Member	139.6GB
pdisk14	00040EFF	Active	Array Member	139.6GB
pdisk15	00040FFF	Active	Array Member	139.6GB
hdisk3	FC0200FF	Optimal	RAID 10T2 Array (0/0)	666.6GB
pdisk22	000416FF	Active	SSD Array Member	387.9GB
pdisk23	000417FF	Active	SSD Array Member	387.9GB
pdisk16	000410FF	Active	Array Member	139.6GB
pdisk17	000411FF	Active	Array Member	139.6GB
pdisk18	000412FF	Active	Array Member	139.6GB
pdisk19	000413FF	Active	Array Member	139.6GB

## PCIe Slots - Memorize This

	<b>4U</b>		<b>2U</b> <small>(no IBM i)</small>	
	<b>1S</b> <b>4U</b>	<b>2S</b> <b>4U</b>	<b>1S</b> <b>2U</b>	<b>2S</b> <b>2U</b>
<b>Total PCIe slots</b>				
<b>All hot swap</b>	<b>7</b>	<b>11</b>	<b>6</b>	<b>9</b>
<b>All Gen3</b>				

- Use 1 of the above slots for a LAN adapter
- Use 1 of the above for expanded function backplane
- PCIe expansion drawers are not yet available.

## PCIe Slots - High Level

	4U		2U (no IBM i)	
	1S 4U	2S 4U	1S 2U	2S 2U
<b>Total PCIe slots (all hot swap)</b>	<b>7</b>	<b>11</b>	<b>6</b>	<b>9</b>
<b>Required* LAN adapter (available for client use)</b>	<b>1</b>	<b>1</b>	<b>1</b>	<b>1</b>
<b>PCIe slots after required* LAN adapter</b>	<b>6</b>	<b>10</b>	<b>5</b>	<b>8</b>
<b>However if use high performance, expanded function backplane</b>	<b>-1</b>	<b>-1</b>	<b>-1</b>	<b>-1</b>
<b>PCIe slots after required* LAN and if using high performance backplane</b>	<b>5</b>	<b>9</b>	<b>4</b>	<b>7</b>


\* required for IBM Manufacturing

- PCIe slots are all Gen3 slots
- 2U are all low profile and 4U are all full high
- There is no PCI expansion drawer announced. There is an SOD.

# PCIe Gen3



- Though these cards physically look the same ... and fit in the same slots
  - Gen3 cards/slots have up to 2X more **bandwidth** than Gen2 cards/slots
  - Gen3 cards/slots have up to 4X more **bandwidth** than Gen1 cards/slots
    - More virtualization
    - More consolidation
    - More ports per adapter
- } saving PCI slots and I/O drawers



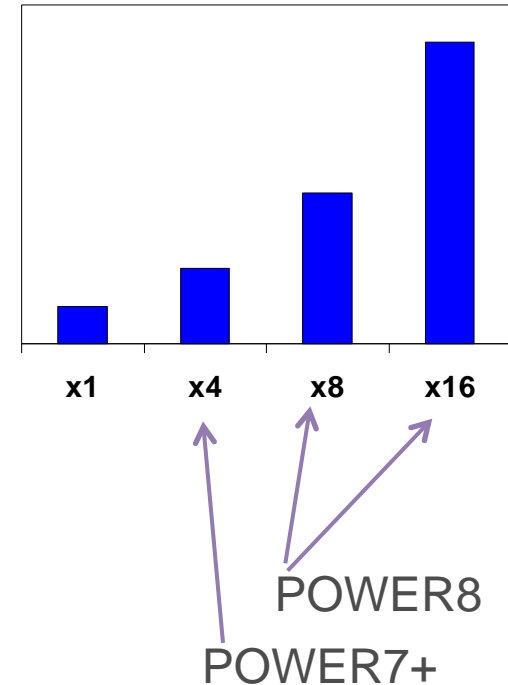

A Gen1 x8 PCIe adapter has a theoretical max (peak) bandwidth of 4 GB/sec.  
 A Gen2 x8 adapter has a peak bandwidth of 8 GB/sec.  
 A Gen3 x8 adapter has a peak bandwidth of 16 GB/sec.



## PCIe x8 and x16



# *POWER8 servers have x8 AND x16 PCIe slots*

Compared to POWER7+ PCIe Gen2 x8 slot, a POWER8 PCIe Gen3 x16 slot has a peak bandwidth of 4X (2X going Gen2 to Gen3 plus 2X going x8 to x16)





- x16 slot/card has more connections than a x8 slot/card
- “x16” or “x8” refers to the number of lanes. 16Gbps/lane (Full Duplex)
- A x8 card can be placed in a x16 slot, but only uses half the connections

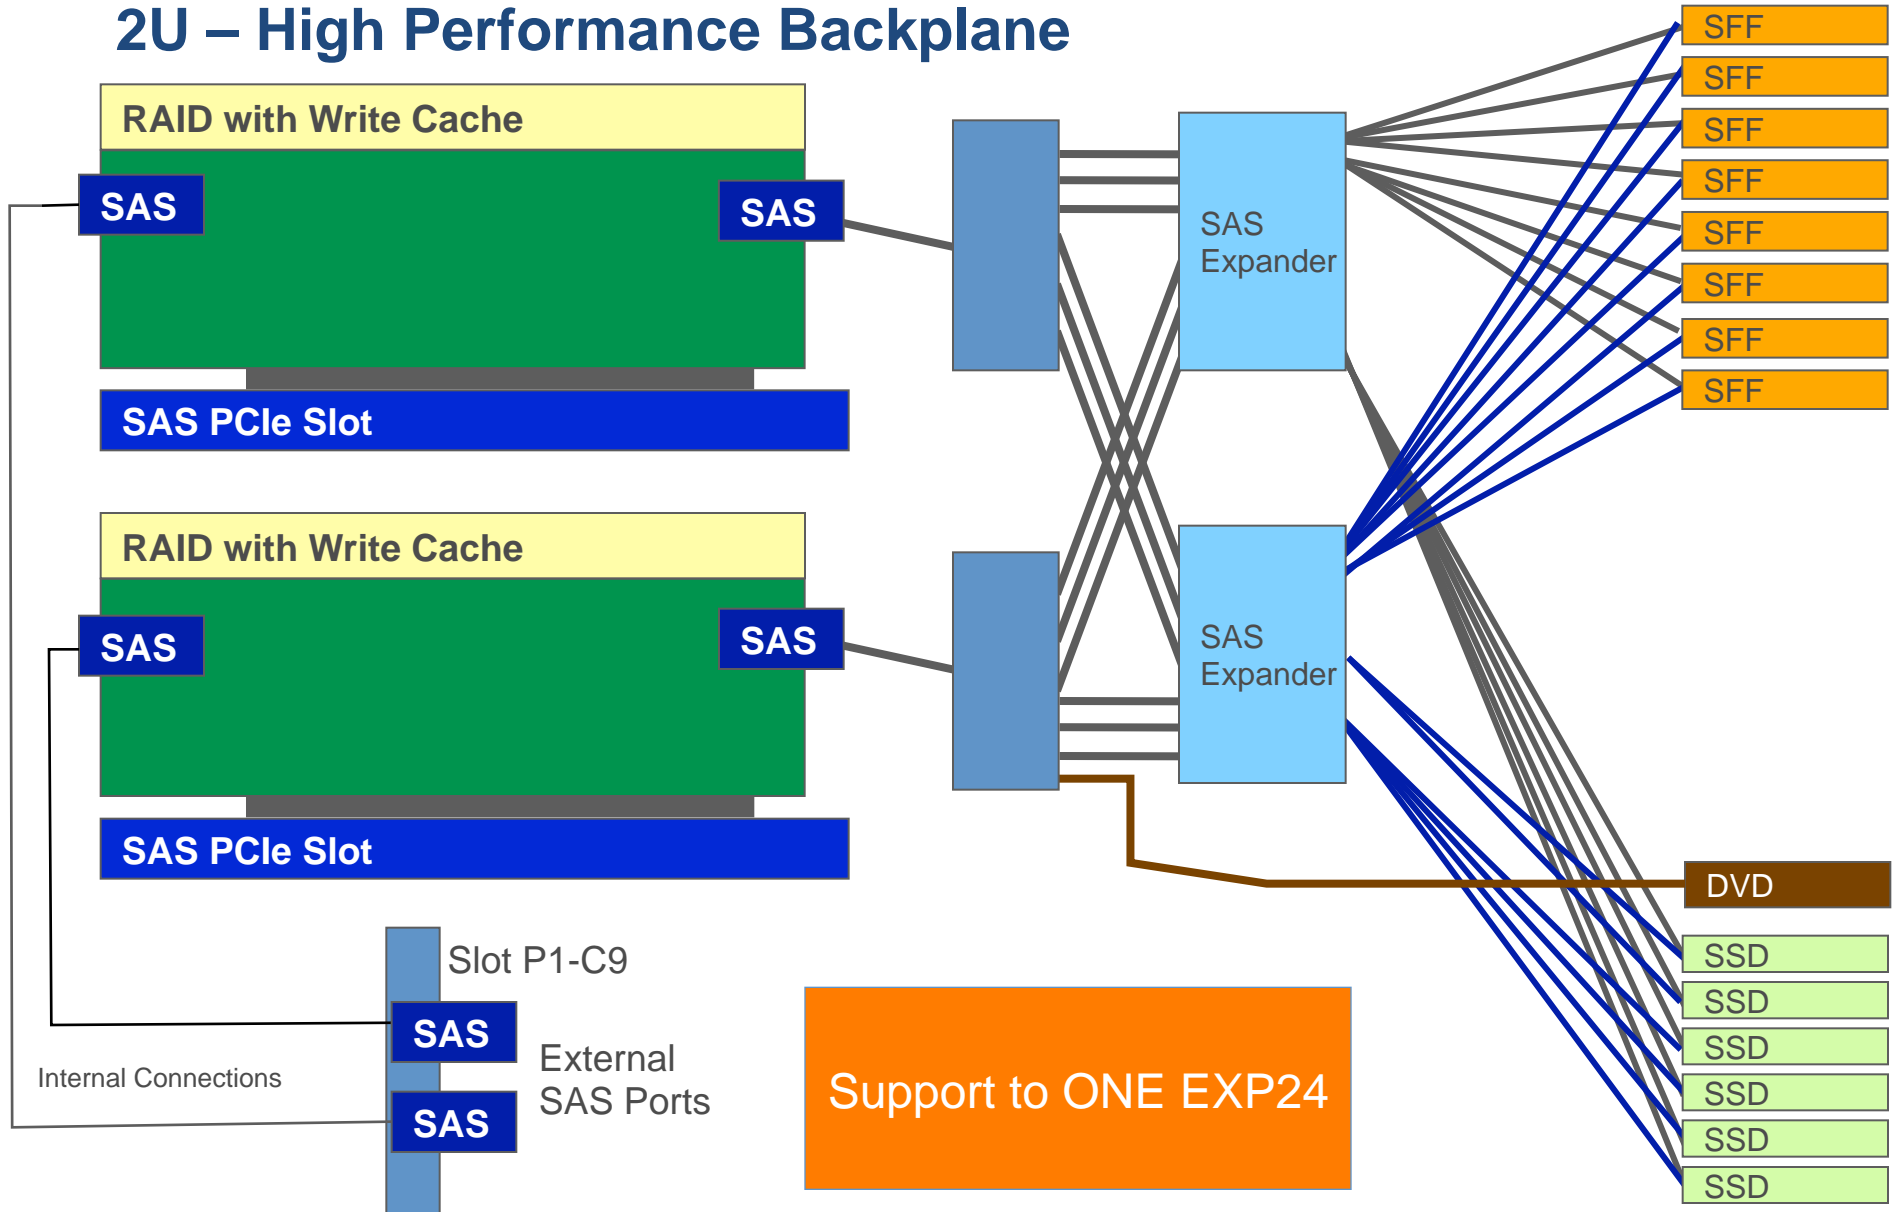
# Standard Backplane



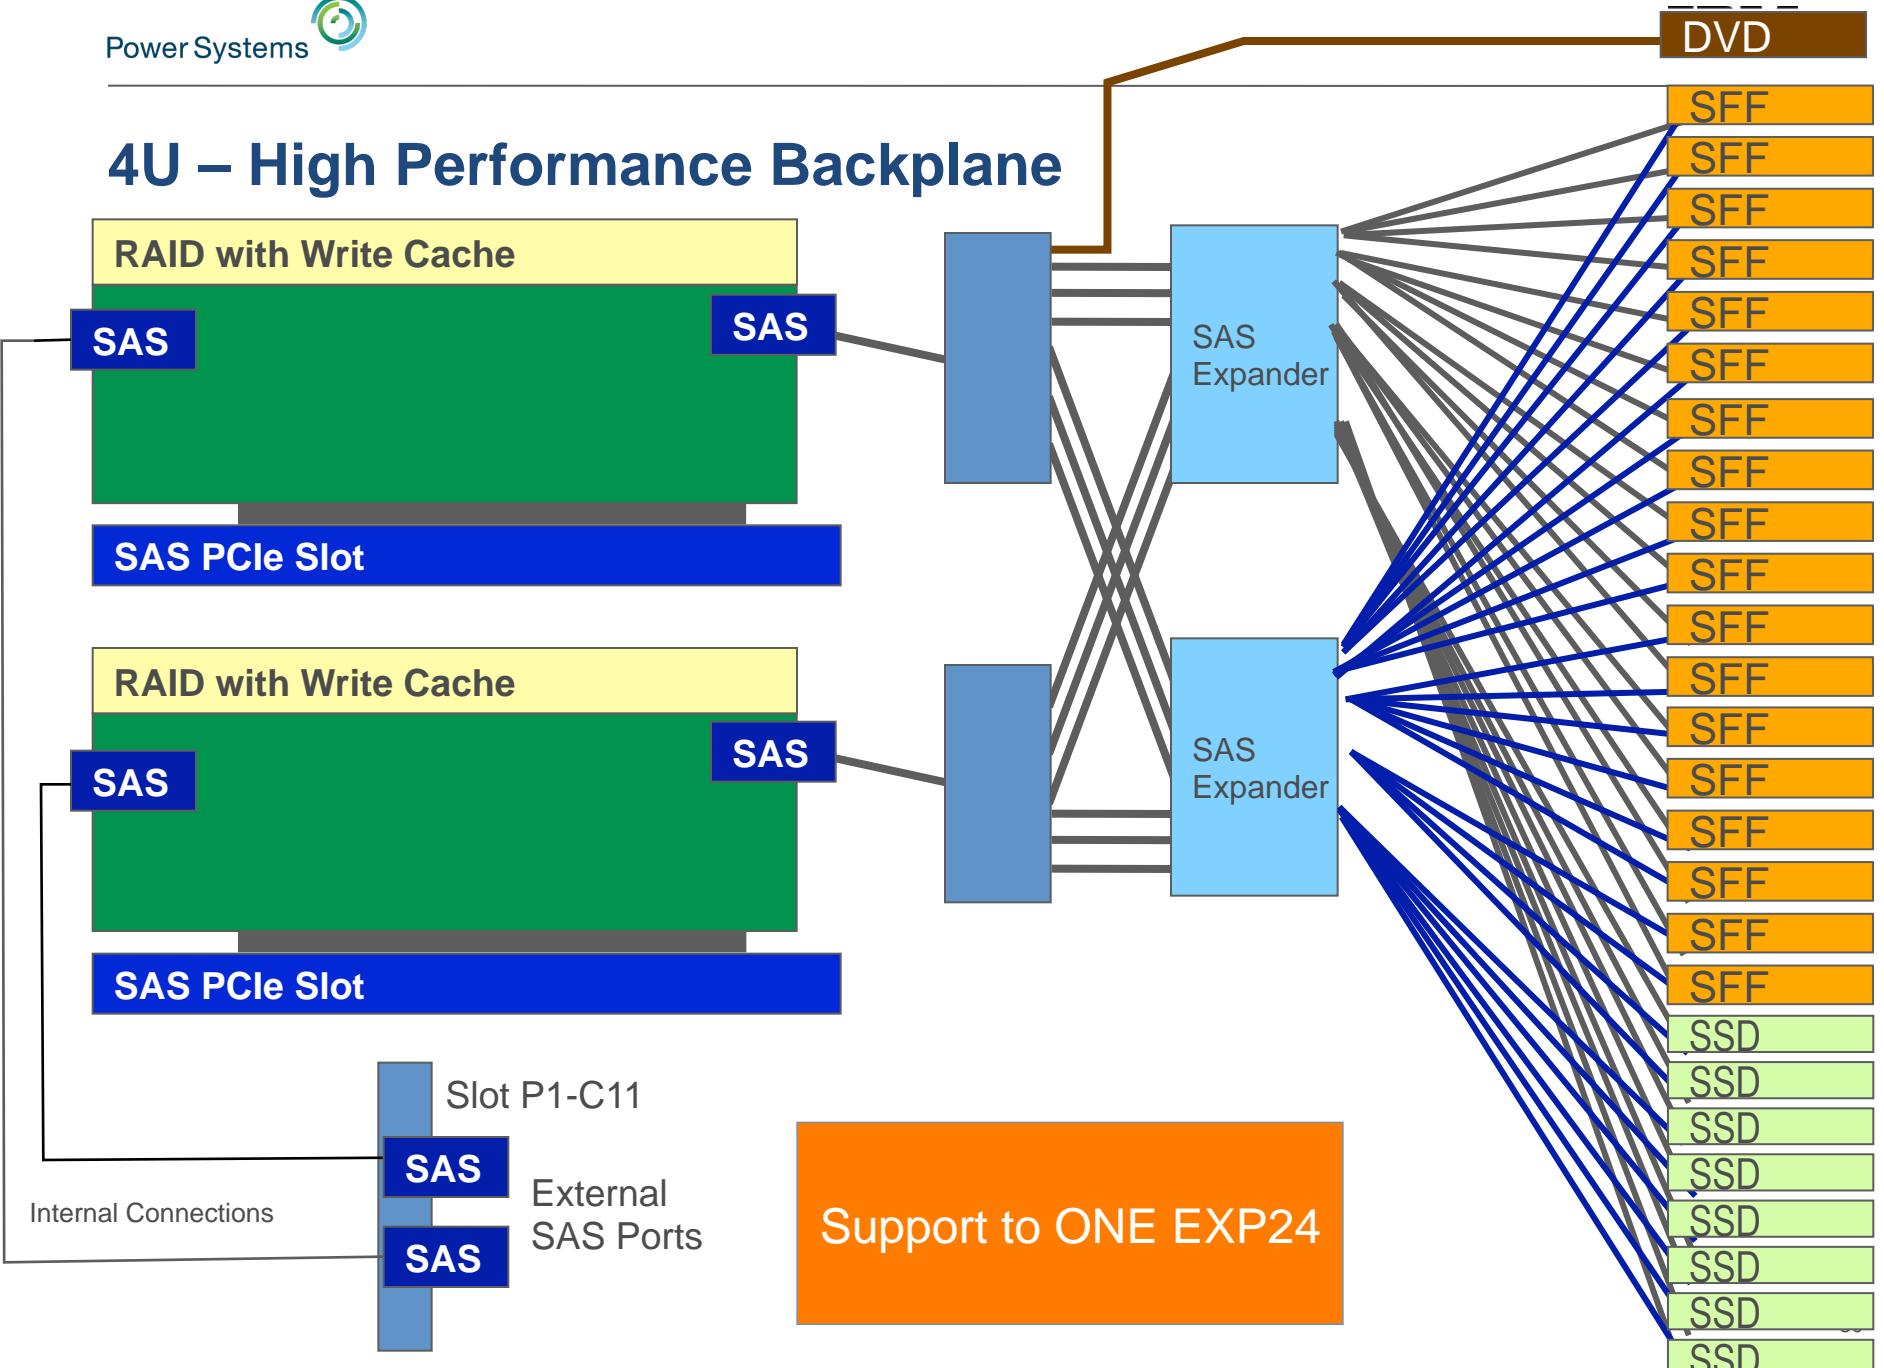
# Standard Backplane + SPLIT



## 2U – High Performance Backplane



# 4U – High Performance Backplane





# PERFORMANCE



## Benchmarks POWER8

- For the 1 and 2 socket POWER8 Systems, IBM has run several benchmarks
  - SAP SD 2-Tier
    - Database + App
  - SPECjEnterprise
    - Database + Java + Storage
  - Oracle Payroll
  - Siebel
  - SPEC
    - Integer
    - Float-point
  - rPerf
    - SMT-1
    - SMT-2
    - SMT-4
    - SMT-8
- All benchmarks show leadership performance for the POWER8 systems.



# SAP SD 2-Tier



## CERTIFICATION

### SAP® Standard Application Benchmarks

The SAP Sales and Distribution (SD) Standard Application Benchmark performed on March 3, 2014, by IBM in Beaverton, OR, USA, was certified by SAP on behalf of the SAP Benchmark Council on April 28, 2014, with the following data:

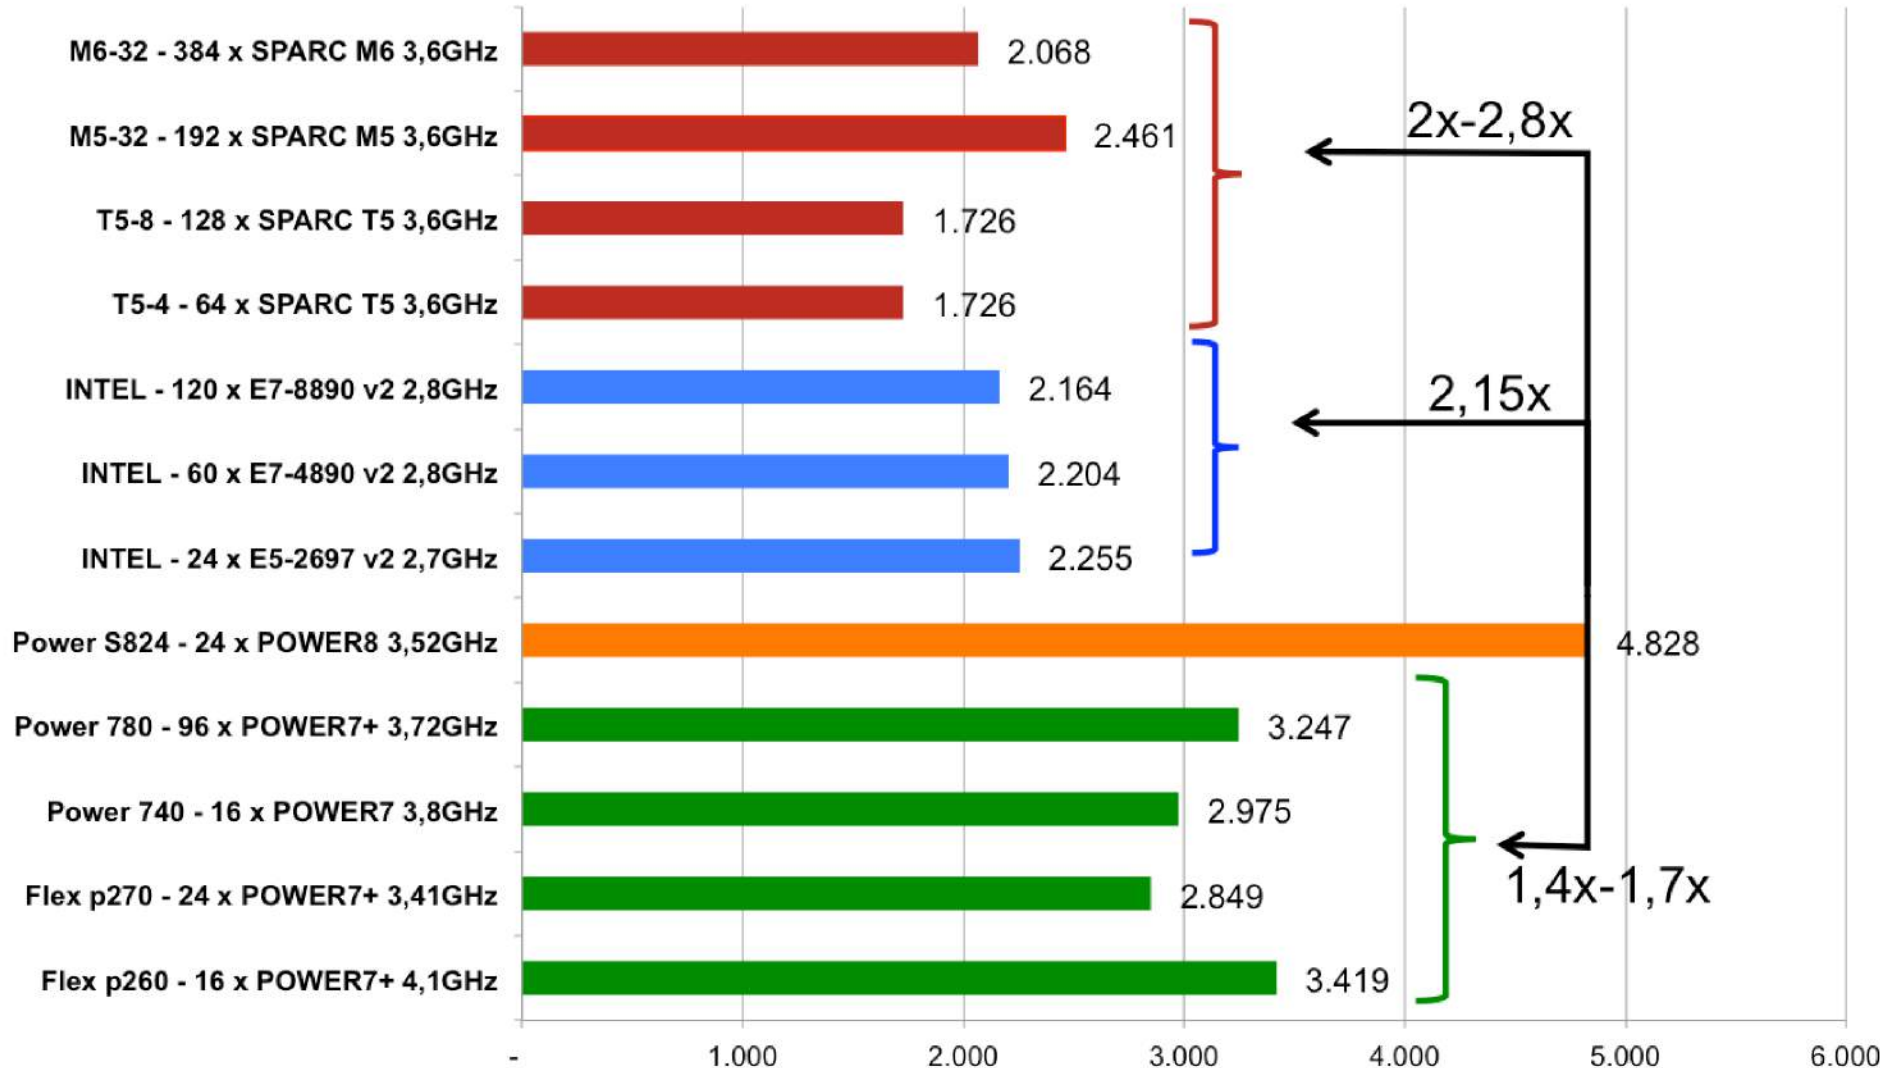
Number of SAP SD benchmark users:	21,212
Average dialog response time:	0.98 seconds
Throughput:	
Fully processed order line items per hour:	2,317,330
Dialog steps per hour:	6,952,000
SAPS:	115,870
Average database request time (dialog/update):	0.011 sec / 0.019 sec
CPU utilization of central server:	99%
Operating system, central server:	AIX 7.1
RDBMS:	DB2 10.5
SAP Business Suite software:	SAP enhancement package 5 for SAP ERP 6.0

**21.212 Users**  
**4.828 SAPs/CORE**  
**99% server usage**  
**0,011 sec dialog**

Configuration:  
 Central server: IBM Power System S824, 4 processors / 24 cores / 192 threads, POWER8, 3.52 GHz, 32 KB (I) and 64 KB (D) L1 cache and 512 KB L2 cache per core, 8 MB L3 cache per core, 512 GB main memory


# SAP SD 2-Tier per Core

## SAP SD 2-Tier - Performance / Core



# SAP SD 2-Tier – POWER8 vs x86

## Total Server Performance - Best x86 vs POWER8






## rPerf – Multiple SMT Levels


	SMT1	SMT2	SMT4	SMT8
<b>Power S814</b>				
<b>6-core 3.0 GHz</b>	<b>48.3</b>	<b>70.1</b>	<b>91.1</b>	<b>97.5</b>
<b>8-core 3.7 GHz</b>	<b>71.4</b>	<b>103.5</b>	<b>134.5</b>	<b>143.9</b>
<b>Power S824</b>				
<b>6-core 3.8 GHz</b>	<b>59.9</b>	<b>86.9</b>	<b>112.9</b>	<b>120.8</b>
<b>12-core 3.8 GHz</b>	<b>116.8</b>	<b>169.4</b>	<b>220.2</b>	<b>235.6</b>
<b>8-core 4.1 GHz</b>	<b>82.3</b>	<b>119.3</b>	<b>155.1</b>	<b>166.0</b>
<b>16-core 4.1 GHz</b>	<b>160.4</b>	<b>232.7</b>	<b>302.4</b>	<b>323.6</b>
<b>24-core 3.5 GHz</b>	<b>209.1</b>	<b>303.2</b>	<b>394.2</b>	<b>421.8</b>
<b>Power S822</b>				
<b>6-core 3.8 GHz</b>	<b>59.9</b>	<b>86.9</b>	<b>112.9</b>	<b>120.8</b>
<b>12-core 3.8 GHz</b>	<b>116.8</b>	<b>169.4</b>	<b>220.2</b>	<b>235.6</b>
<b>10-core 3.4 GHz</b>	<b>88.2</b>	<b>127.8</b>	<b>166.2</b>	<b>177.8</b>
<b>20-core 3.4 GHz</b>	<b>171.9</b>	<b>249.3</b>	<b>324.0</b>	<b>346.7</b>

# rPerf / Core – POWER7 vs POWER8

## rPerf / Core - P7 vs P7+ vs P8 @ 3.5GHz




## rPerf / Core - P7+ vs P8 @ 4.1GHz



# Power 750+ vs POWER8 4U2S


Performance



**P 750+ P8 4U2S**

POWER8 ≥  
Performance

Total # Cores




**P 750+ P8 4U2S**

Fewer cores means Lower  
Software Costs

Does not reflect lower PVU

savings


TCA



**P 750+ P8 4U2S**

Lower investment costs

Performance  
per \$




**P 750+ P8 4U2S**


Better performance  
per dollar spent

# Oracle EBS Payroll Batch XL

## Oracle eBS 12.1.3 Payroll Bacth XL Overall Performance



## Oracle eBS 12.1.3 Payroll Bacth XL Performance/core



# RAS





# RAS Feature Overview

- Standard
- ◻ Optional
- Not Available

RAS Item	POWER7+ 710 / 730	POWER7+ 720 / 740	POWER8 models
<i>Redundant / Hot Swap Fans &amp; Blowers</i>	●	●	●
<i>Hot Swap DASD &amp; Media</i>	●	●	●
<i>Hot Swap PCI Adapters</i>	—	—	●
<i>Concurrent Firmware Update</i>	●	●	●
<i>Redundant / Hot Swap Power Supplies</i>	●	●	●
<i>Dual disk controllers (split backplane)</i>	—	◻	◻
<i>Processor Instruction Retry</i>	●	●	●
<i>Alternate Processor Recovery</i>	●	●	●
<i>Storage Keys</i>	●	●	●
<i>PowerVM™/Live Part. Mobility/Live App Mobility</i>	◻	◻	◻
<i>Dynamic Processor Sparing</i>	—	—	—
<i>Redundant Service Processors</i>	—	—	—
<i>Redundant System Clocks</i>	—	—	—
<i>Hot GX Adapter Add and Cold Repair</i>	—	—	N/A
<i>Dynamic Service Processor &amp; System Clock Failover</i>	—	—	—
<i>Enterprise Memory ( Memory Sparing )</i>	—	—	●
<i>Integrated TPMD Function</i>	—	—	●
<i>Hot GX Adapter Repair</i>	—	—	N/A
<i>Active Memory Mirroring for Hypervisor</i>	—	—	—
<i>Power Pools</i>	—	—	—
<i>Dynamic Processor Sparing</i>	—	—	—

## RAS Features – Always On

- L1/L2/L3 Cache ECC
  - Correct Single Bit Error
  - Detect Double Bit Error
  - Can delete a cache line with a persistent issue
  - Can repair a column of an L3 cache dynamically by using spare capability
- Cores
  - Processor Instruction Retry
  - Alternate Processor Recovery
  - Partition Availability Priority
- Memory
  - Retry Buffer on Controllers
  - Spare data lanes on memory buses
  - L4 on DIMM – Correct single error, Detect double error
  - Memory Scrubbing
  - ChipKill
  - Spare DRAM modules - can handle up to 4 DRAM module ChipKills

# Intel Memory – Performance or RAS not Both

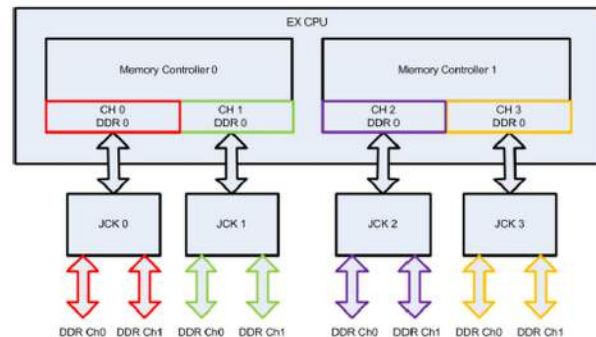
- 2 memory modes
- Lockstep
  - RAS
  - Redundant memory channels
  - 1600MHz
- Performance
  - Performance
  - Non redundant memory channels
  - 1333MHz
- Setting via BIOS
- Reboot

## Intel® C10x<sup>1</sup> Scalable Memory Buffer Performance vs. Lockstep Mode BIOS Options

Two operation modes supported, default setting is Performance Mode

### Lockstep Mode (1:1)

➤ Up to 1600 : 1600

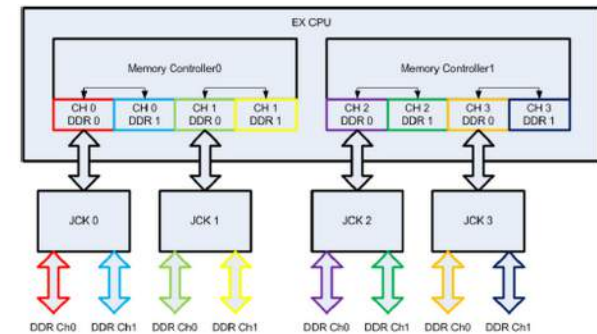


1:1 bus ratio

DDR channels @ DDR3-1600

### Performance Mode (2:1)

➤ Up to 2667:1333



2:1 bus ratio

DDR channels @ DDR3-1333

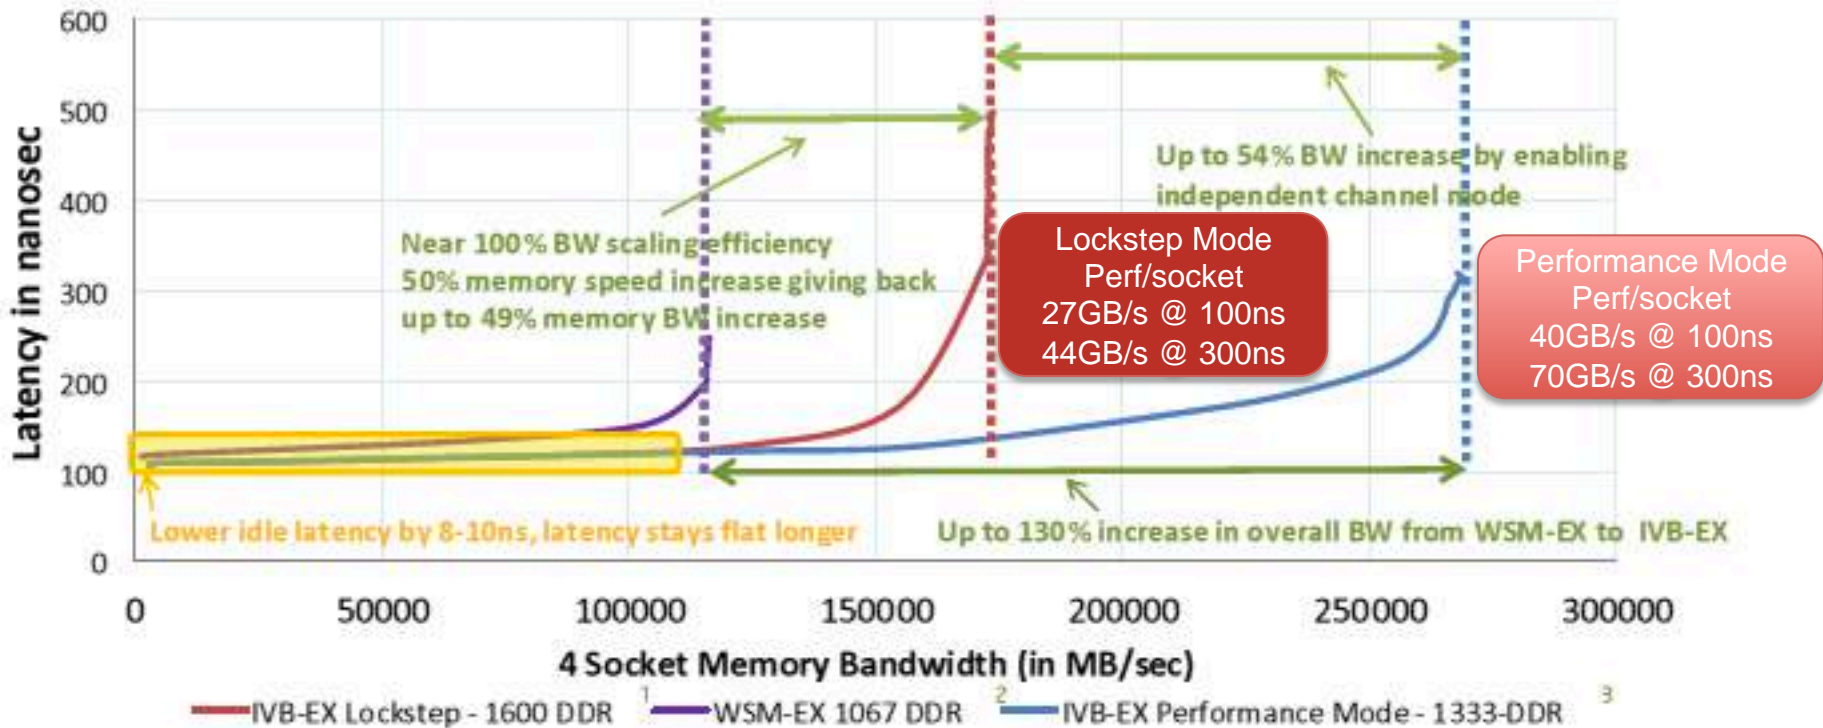
<sup>1</sup> Intel® C10x is the external product name for the JordanCreek Memory Extension Buffer Chip. Two versions exist, one that supports up to 8 DIMMs and a second one that supports up to 12 DIMMs, hence the suffix "x" is used.

# Memory Bandwidth & Latency

## Massive increase in throughput<sup>1</sup>

Results have been measured based on internal Intel tools and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

### 3Read 1Write traffic mix



Performance mode use recommended for highest system memory bandwidth (default setting with released BIOS)



# Solutions






# New Power Systems Solutions for Big Data & Analytics

*Bringing insight to the point of impact with systems and solutions optimized for Big Data & Analytics performance*

**IBM Solution for Hadoop – Power Systems Edition** 

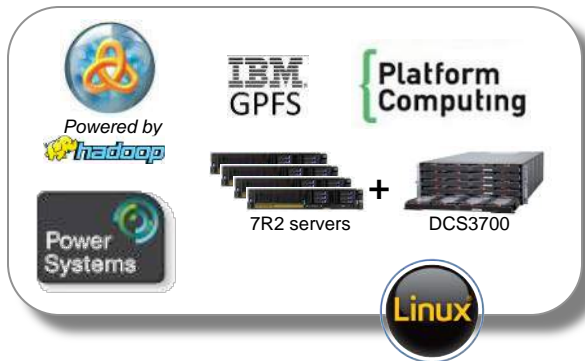
**BLU Acceleration – Power Systems Edition** 


**IBM Solution for Analytics – Power Systems Edition** 


Integrated platform solution for Hadoop analytics software

Integrated platform solution for analytics software


Integrated solution pre-configured with Cognos BI and/or SPSS and/or BLU Acceleration



Powered by  **hadoop**

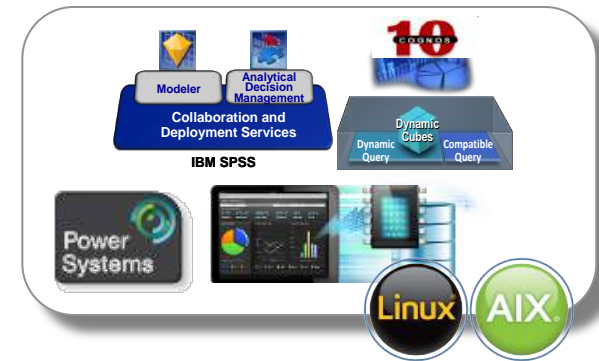
IBM **GPFS**  **Platform Computing**

7R2 servers + DCS3700

Power Systems 



Power Systems 





10 YEARS

Modeler Analytical Decision Management

Collaboration and Deployment Services

IBM SPSS

Dynamic Query Dynamic Cubes Compatible Query

Power Systems  

# Power Systems can deliver insight to the point of impact with big data & analytics **accelerators**

**75% less storage<sup>1</sup>**

## IBM FlashSystem



**82X faster insights<sup>2</sup>**

## Next Generation In-Memory



**24:1 consolidation<sup>3</sup>**

## POWER8 with CAPI Flash Accelerators



1- Source: COCC Cast Study <http://bit.ly/1iQemuu>

2- 82X is based on IBM internal tests as of April 17, 2014 comparing IBM DB2 with BLU Acceleration on Power with a comparably tuned competitor row store database server on x86 executing a materially identical 2.6TB BI workload in a controlled laboratory environment. Test measured 60 concurrent user report throughput executing identical Cognos report workloads. Competitor configuration: HP DL380p, 24 cores, 256GB RAM, Competitor row-store database, SuSE Linux 11SP3 (Database) and HP DL380p, 16 cores, 384GB RAM, Cognos 10.2.1.1, SuSE Linux 11SP3 (Cognos). IBM configuration: IBM S824, 24 cores, 256GB RAM, DB2 10.5, AIX 7.1 TL2 (Database) and IBM S822L, 16 of 20 cores activated, 384GB RAM, Cognos 10.2.1.1, SuSE Linux 11SP3 (Cognos). Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment.

3- 24:1 system consolidation ratio (12:1 rack density improvement) based on a single IBM S824, (24 cores, POWER8 3.5 GHz), 256GB RAM, AIX 7.1 with 40 TB memory based Flash replacing 24 HP DL380p, 24 cores, E5-2697 v2 2.7 GHz), 256GB

# IBM DB2 with BLU Acceleration

## Unmatched Innovation from IBM Research & Development

### Next Generation In-Memory

In-memory columnar processing with dynamic movement of data from storage





### Actionable Compression

Patented compression that preserves order so data can be used without decompressing



### Parallel Vector Processing

Multi-core and SIMD parallelism (Single Instruction Multiple Data)

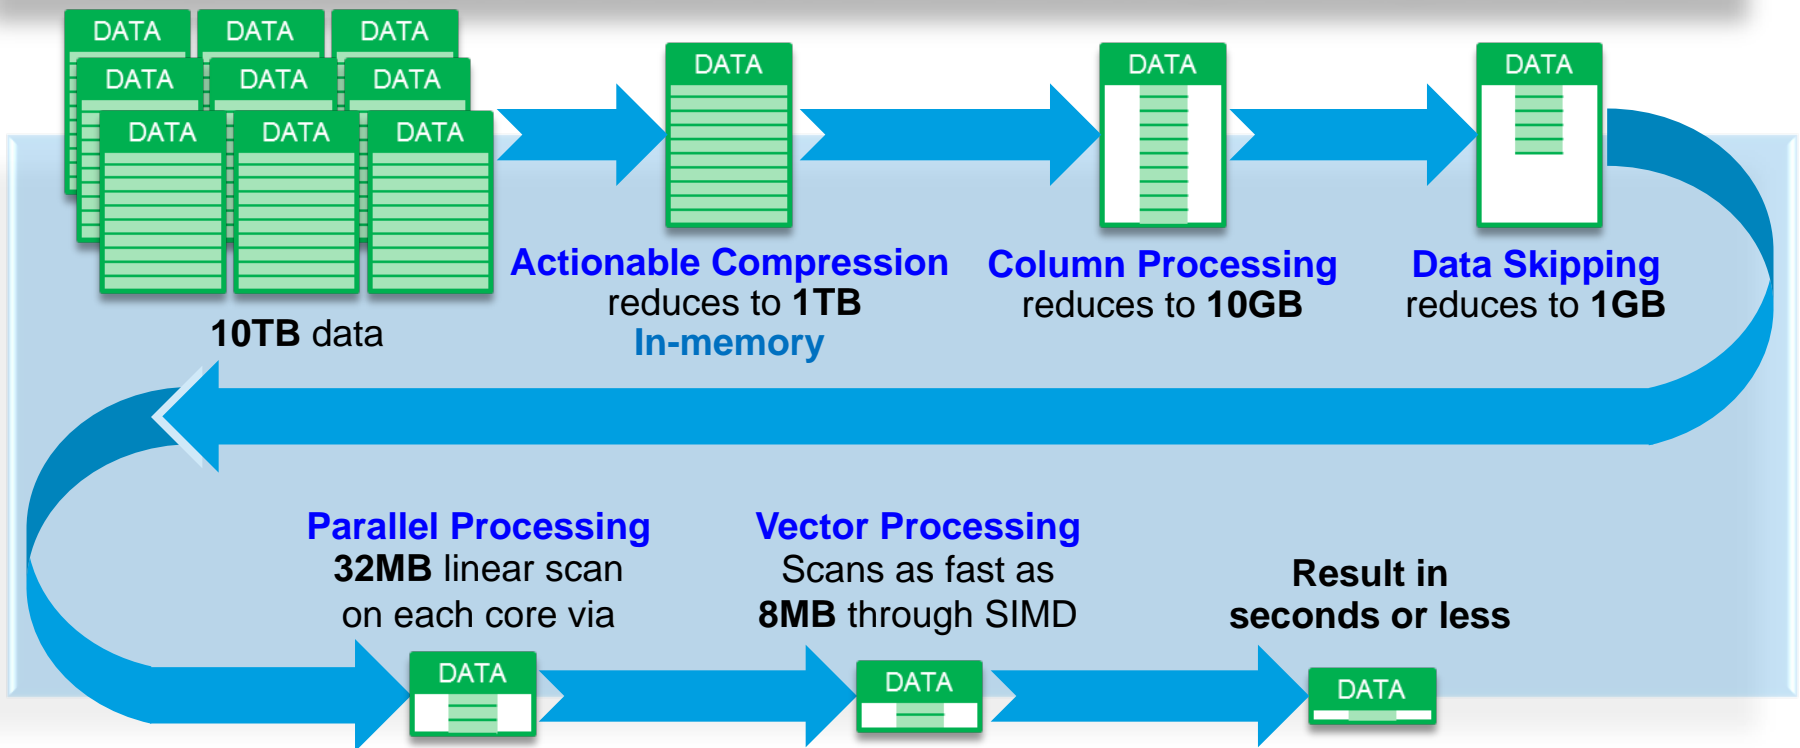


### Data Skipping

Skips unnecessary processing of irrelevant data

# BLU Acceleration Illustration: 10TB query in seconds or less

- The System: **32 cores, 1TB memory, 10TB table with 100 columns and 10 years of data**
- The Query: **How many “sales” did we have in 2010?**
  - `SELECT COUNT(*) from MYTABLE where YEAR = '2010'`
- The Result: **In seconds or less as each CPU core examines the equivalent of just 8MB of data**



## Power Systems clients are seeing Remarkable Results



*“We tested some representative queries taken from our existing SAP application and tested them on DB2 with BLU Acceleration and observed performance improvements in many of our query response times. For example, **one of our most time consuming queries experienced a 50X performance improvement.**”*

*- Richard Simms, Director of Infrastructure, Fossil*



*“I was asked to help one of our analysts who gave up on something he was executing across a billion records, because **after 3 hours, the answer didn't come back.** We took that query (which involved a big join) & ran it on DB2 with BLU Acceleration and it **finished in 10 seconds – 1080X faster.**”*

*- Randy Wilson, Lead DB2 for LUW DBA*



*“The BLU Acceleration technology makes our analytical queries run **4-15x faster and decreases the size of our tables by 10X.** But it's all the things I don't have to do with BLU that make me appreciate the technology even more: **no tuning, no partitioning, no indexes, no aggregates.**”*

*- Andrew Juarez, Lead SAP Basis and DBA*



# DB2 BLU Sizing Validation – System Recommendations

The following table shows various uncompressed database sizes and systems to consider using. A larger system can always contain an optimally configured LPAR.

Total Uncompressed DB Size	Active Uncompressed Data Size	# of Cores	RAM
<1TB	< 330GB	8 -12	64 - 128
1 – 2.5	330 – 800GB	8 - 16	128 - 256
2.5 – 5	800 – 1700GB	16 - 24	256 - 512
5 – 7.5	1.7 – 2.5TB	16 - 32	512 - 1024
7.5 – 10	2.5 – 3.3TB	32 - 48	1024 - 2048
10 – 15	3.3 – 5TB	48 - 64	1024 - 2048
15 – 20	5 – 6.7TB	64 - 128	2048 - 4096
20 +	6.7TB >	64 +	2048 +

These values may be adjusted once the Best Practices and sizing guidelines are published.

# IBM Solution for Analytics - Power Systems Edition



- **POWER8 server options:** Power S814, Power S822, Power S824
- **System Software:** AIX Standard or Enterprise Edition, PowerVM Standard or Enterprise Edition
- **BLU Acceleration options:** IBM DB2 Advanced Workgroup or Advanced Enterprise Edition
- **Analytics options:** Cognos BI, SPSS Modeler, SPSS C&DS, SPSS ADM, InfoSphere DataStage

## Simple to Acquire

*Order server, storage, software and support from a single vendor*

## Simple to Deploy

*Pre-installed and pre-optimized server, storage & software*

## Simple to Implement


*Highly scalable to grow as your analytics need change*

<http://www.ibm.com/systems/power/solutions/analytics/editions/preload-analytics.html>

Includes options for pre-installed DB2 AWSE or DB2 EWSE, Cognos BI, SPSS Modeler, SPSS C&DS, SPSS ADM, or InfoSphere DataStage licenses ready for activation. Registration of valid IBM SWG licenses required prior to server shipment with selected options. Pre-load of Power hardware and software is not available in Greater China

# IBM Solution for Hadoop - Power Systems Edition

**NEW.** A storage-dense integrated platform optimized to simplify and accelerate unstructured big data & analytics



## Speed Matters

Higher ingest rates delivers 37% faster insights than competitive Hadoop solutions with 31% fewer data nodes.<sup>1</sup>

## Availability Matters

Better reliability and resiliency with 73% fewer outages and 92% fewer performance problems over x86.<sup>2</sup>

*Integrated platform solution  
for Hadoop ready for  
analytics software*

1) Based on STG Performance testing comparing to Cloudera/HP published benchmark


2) CLAIMS: Solitaire Interglobal Paper - Power Boost Your Big Data Analytics Strategy – <http://www.ibm.com/systems/power/solutions/assets/bigdata-analytics.html?LNK=wf>

# IBM Solution for Hadoop – Power Systems Edition

Key requirements & design parameters – focused on customer value

- Best-in-class hardware
- Dense storage subsystem
- Advanced software capabilities
- Better reliability & management
- Best in class file system
- Automated cluster provisioning
- IBM Platform Symphony

*IBM InfoSphere BigInsights*  
or Open-source Hadoop



*IBM Platform Symphony*  
IBM Platform Cluster Manager



*Distributed File System*  
IBM Elastic Storage, HDFS



*Linux Operating Environment*  
RHEL SUSE




*IBM Power Systems*  
IBM Power 7+, Power8



# Standard Configurations – POD – based design

Power Big Data clusters are built using a simple building block approach to tailor the mix of CPU and storage to application requirements.



## Compute Dense Data POD



**Four servers per storage subsystem**

## Balanced Data POD



**Two servers per storage subsystem**

## Storage Dense Data POD



**One server per storage subsystem**







## SQL on Hadoop – BigInsights (BigSQL) on Power8



**ETA 3Q14**  
**BigInsights 3.0.0.1**


- BigSQL allows you to run SQL on your Hadoop data
  - All data stays in a Hadoop filesystem (HDFS or GPFS-FPO)
- Uses IBM's mature relational MPP database technology
- Provides for enterprise-level security
  - Authentication, authorization (fine-grain access), auditing, etc.
- Other enterprise features:
  - Federation
  - Use of standard SQL, and all the SQL / Business Analytics tools
  - Mixed workload cluster management
- Performance results
  - Ran TPC-DS queries only on a 10TB data set
  - All 99 queries run without any changes
  - Ran a multi-user throughput test
  - Comparison against Hive running on Intel based cluster

# POWER8 delivers **2.5x performance** on Big Data/Hadoop

POWER8 requires 60% fewer servers based on the best x86 published result



- POWER8 S822L delivers over 2.5x the performance of the best published x86 system
  - ... and continues to offer far superior RAS
  
- POWER8 exploits additional cores, more threads, larger caches, memory bandwidth
  
- Terasort is a popular benchmark to measure the performance of a Hadoop solution
  - Sorts a large dataset (10 TB) in parallel
  - Exercises the Map-reduced framework and Hadoop Distributed File System (HDFS)

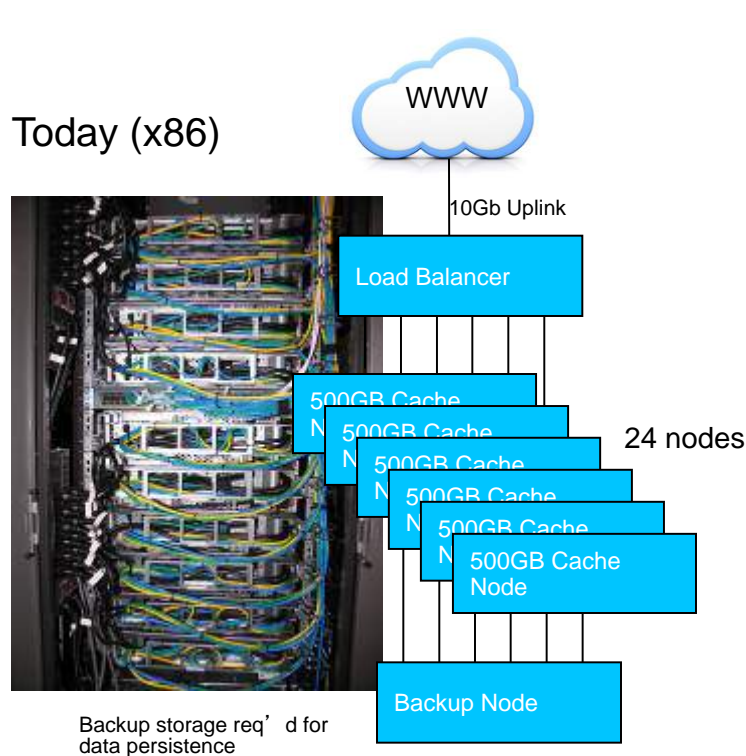


IBM Analytics Stack: IBM Power System S822L; 24 cores / 192 threads, POWER8; 3.0GHz, 512 GB memory, RHEL 6.5, InfoSphere BigInsights 3.0.0.0


[http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/te\\_tera.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/te_tera.pdf)

# “In-Memory” NoSQL/KVS Infrastructure via CAPI-attached Flash

12TB, 6K users, 2GB/User Example – Power Systems delivers 24:1 consolidation over x86<sup>1</sup>



26 servers, switches, etc. = >1 rack  
 In-memory database w/ 12TB DRAM  
 ~18kW



## Sizings

### Power + CAPI Flash Advantage

- 24:1 physical server consolidation
- >1 rack to 4U Density (>12x)
- 18kW  $\diamond$  1.5kW Energy (12x)
- \$24/user  $\diamond$  \$7.5/user Cost (3.2x)
- 4.5k  $\diamond$  60k Users per rack (13x)<sup>2</sup>



1) 24:1 system consolidation ratio (12:1 rack density improvement) based on a single IBM S824, (24 cores, POWER8 3.5 GHz), 256GB RAM, AIX 7.1 with 40 TB memory based Flash replacing 24 HP DL380p, 24 cores, E5-2697 v2 2.7 GHz), 256GB RAM, SuSE Linux 11SP3. Inbound network limits performance to 1M IOPs in both scenarios, equal capacity (#user, data) in both cases. x86 cost includes 10k\$ for 2x 1U switches

2) Normalized to a single rack

# NC State analyzes unstructured big data with Linux on Power

## Largest technical college in the Southeast deploys Cloud based big data & analytics solution on IBM PowerLinux

### Need

- NC State University tries to teach students how to look at a business problem, analyze them, and come to a decision-based understanding of where they' re going to go with a new business.

### Solution

- IBM Power Linux servers
- IBM InfoSphere BigInsights
- IBM Content Analytic Studio
- Deployed in an IBM Cloud

NCSU ([Case Study](#), [YouTube](#))

### Results

- Establish a Cloud computing environment that allows students to use products such as IBM Content Analytic Studio for doing natural language processing.

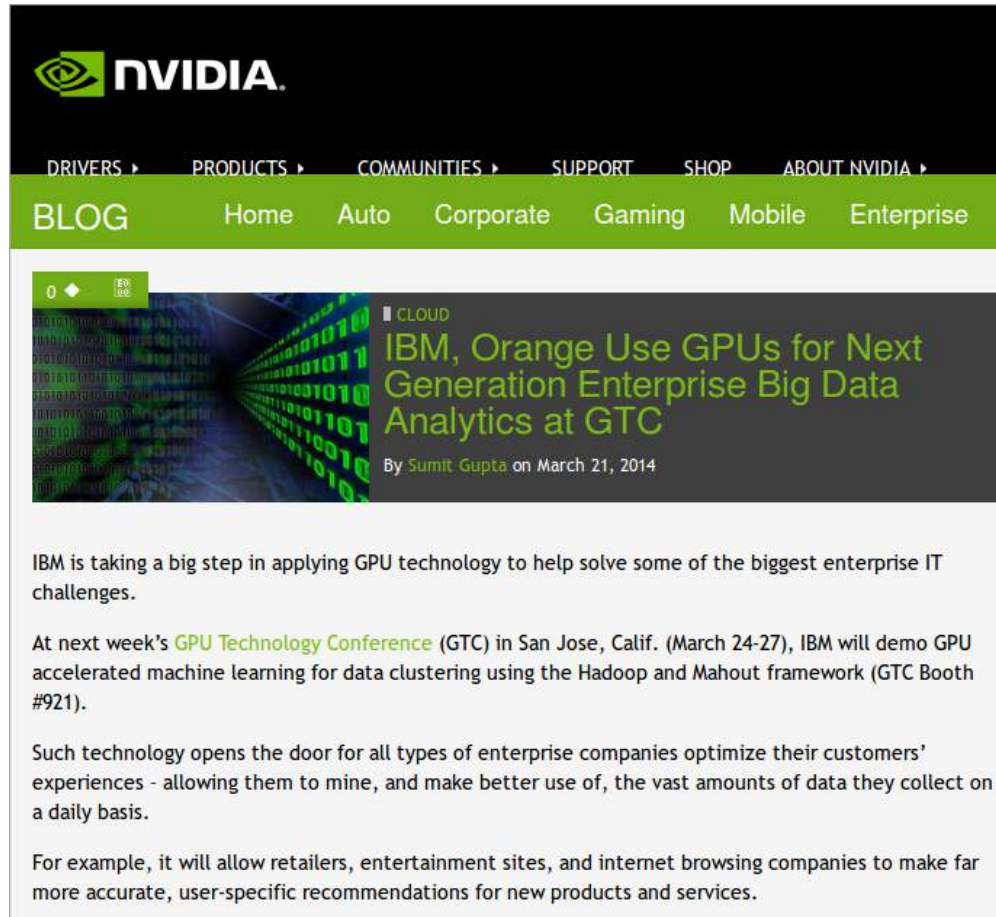


*"By working smarter and doing smarter computing, you'll be able to optimize your systems and, again, leverage your talent for looking for important knowledge."*

- Dr. Michael Kowolenko, Senior Research Scholar College of Management, NCSU



# GPU processing and Big Data



The screenshot shows the NVIDIA website's blog section. At the top, the NVIDIA logo is displayed on a black background. Below it is a navigation bar with links for DRIVERS, PRODUCTS, COMMUNITIES, SUPPORT, SHOP, and ABOUT NVIDIA. A green banner below the navigation bar contains the word "BLOG" and a menu with links for Home, Auto, Corporate, Gaming, Mobile, and Enterprise. The main content area features a large article preview with a background image of binary code. The article title is "IBM, Orange Use GPUs for Next Generation Enterprise Big Data Analytics at GTC" and it is attributed to Sumit Gupta on March 21, 2014. The article text discusses IBM's use of GPU technology for data clustering and machine learning at the GPU Technology Conference (GTC).

IBM is taking a big step in applying GPU technology to help solve some of the biggest enterprise IT challenges.

At next week's [GPU Technology Conference \(GTC\)](#) in San Jose, Calif. (March 24-27), IBM will demo GPU accelerated machine learning for data clustering using the Hadoop and Mahout framework (GTC Booth #921).

Such technology opens the door for all types of enterprise companies optimize their customers' experiences - allowing them to mine, and make better use of, the vast amounts of data they collect on a daily basis.

For example, it will allow retailers, entertainment sites, and internet browsing companies to make far more accurate, user-specific recommendations for new products and services.

**IBM is demonstrating the use of GPU accelerators on a distributed computing system (required for such an enormous data set) for clustering using Hadoop. With GPU accelerators working alongside IBM Power CPUs, the demo runs eight-times faster than with a Power system without GPUs.**

<http://blogs.nvidia.com/blog/2014/03/21/ibm-orange-gpus-big-data/>





# HMC Version 8

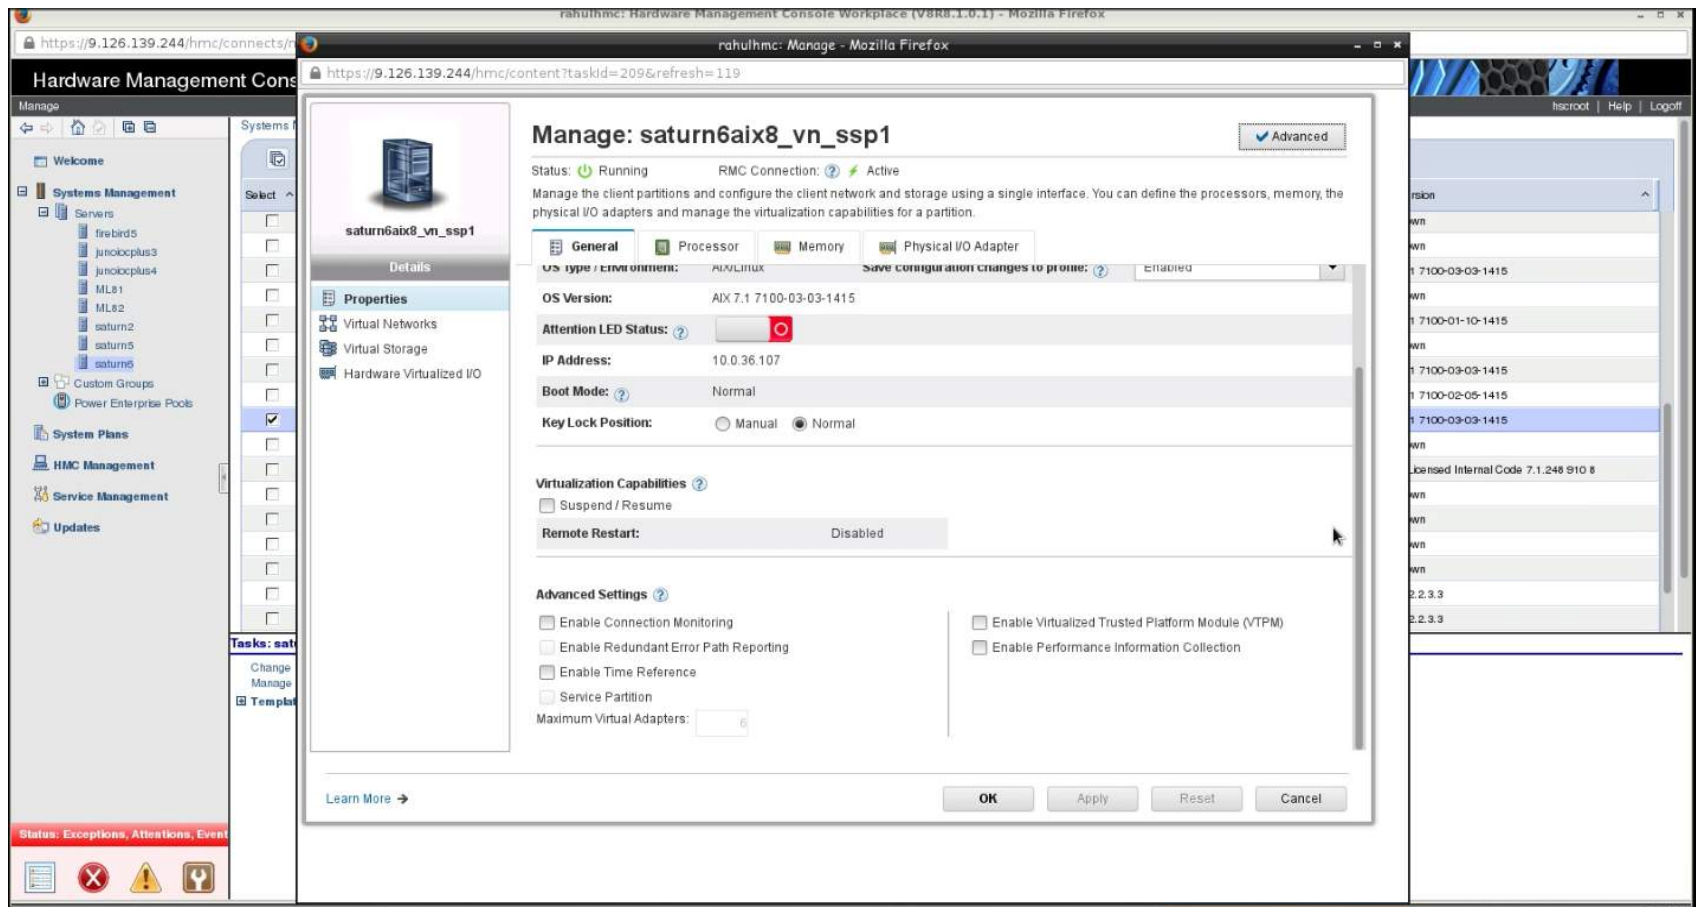


- **Hardware Support / Serviceability**
  - Support for new POWER8 processor-based systems
  - Support for new I/O adapters
  - New rack-mount HMC – 7042-CR8
  - Call Home Serviceable Event Manager (SEM)
  - Server Firmware Entitlement
  - Additional tracing during Serviceable Event creation
- **PowerVM Support**
  - Integrated Facility for Linux
  - HMC CLI for Remote Restart
  - Dynamic reset of Partition Remote Restart Attribute
  - IBM i 5250 virtual Terminal changes
  - Schedule DPO operation from GUI
- **HMC Console Management**
  - Security: Support of NIST 800-131a
  - HMC Installation Changes
- **HMC Usability Enhancements**
  - PowerVM Management
  - Performance Management
- Redbook: *IBM Power Systems Hardware Management Console Version 8 Release 8.1.0 Enhancements* - <http://www.redbooks.ibm.com/redpieces/pdfs/sg248232.pdf>

## HMC GUI Highlights

- **Seamless integration with existing HMC dashboard**
  - New GUI tasks are available right inside current HMC dashboard
  - New GUI tasks can launch existing HMC tasks from itself.
  - Look and feel in accordance with IBM OneUI guidelines.
  - GUI design after regular feedback from real customers and client council meetings.
  
- **Change in Concepts**
  - The VIOS is no more a partition, its considered part of Virtualization layer.
  - Storage assignments made from a Logical Partition perspective
  - Introduction of new concepts such Virtual Network and Network Bridge.
  - Abstracting low level details such as adapters in Virtual storage assignment.
  
- **New Features**
  - Introduction of Shared Storage Pool management.

# HMC New Management GUI - Managing partitions

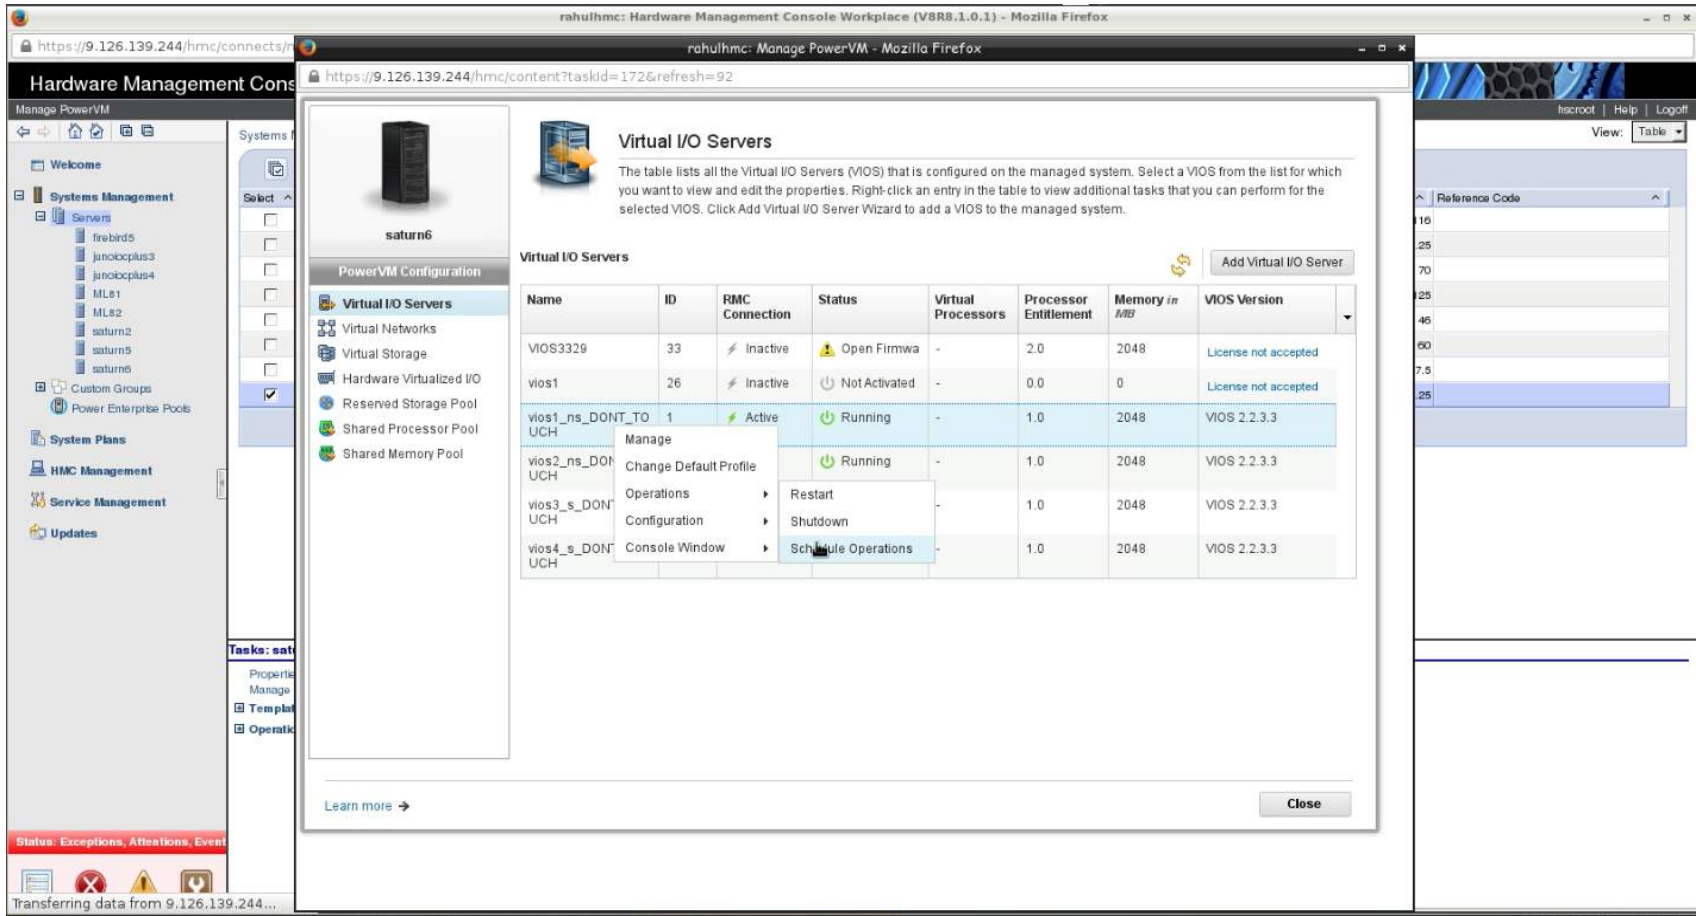


The screenshot displays the Hardware Management Console (HMC) Management GUI. The main window is titled "Manage: saturn6aix8\_vn\_ssp1" and shows the configuration for a running partition. The interface includes a left-hand navigation pane with categories like "Systems Management", "System Plans", and "HMC Management". The main content area is divided into several sections:

- Details:** Shows the partition name "saturn6aix8\_vn\_ssp1" and a "Details" button.
- Properties:** Lists various properties such as "Virtual Networks", "Virtual Storage", and "Hardware Virtualized I/O".
- General:** Contains the following information:
  - Status: Running
  - RMC Connection: Active
  - OS type / Environment: AIX/Linux
  - OS Version: AIX 7.1 7100-03-03-1415
  - Attention LED Status: On (indicated by a red light icon)
  - IP Address: 10.0.36.107
  - Boot Mode: Normal
  - Key Lock Position: Normal (selected)
- Virtualization Capabilities:**
  - Suspend / Resume: Disabled
  - Remote Restart: Disabled
- Advanced Settings:**
  - Enable Connection Monitoring:
  - Enable Redundant Error Path Reporting:
  - Enable Time Reference:
  - Service Partition:
  - Maximum Virtual Adapters: 6
  - Enable Virtualized Trusted Platform Module (VTPM):
  - Enable Performance Information Collection:

At the bottom of the configuration window, there are buttons for "OK", "Apply", "Reset", and "Cancel". A "Learn More" link is also present. The right-hand side of the screenshot shows a list of other partitions, with "1 7100-03-03-1415" selected.

# HMC New Management GUI - Managing VIOS



The screenshot displays the 'Virtual I/O Servers' management page in the HMC GUI. The system being managed is 'saturn6'. The page contains a table of configured VIOS servers and a context menu for the selected 'vios4\_s\_DON\_UCH' server.

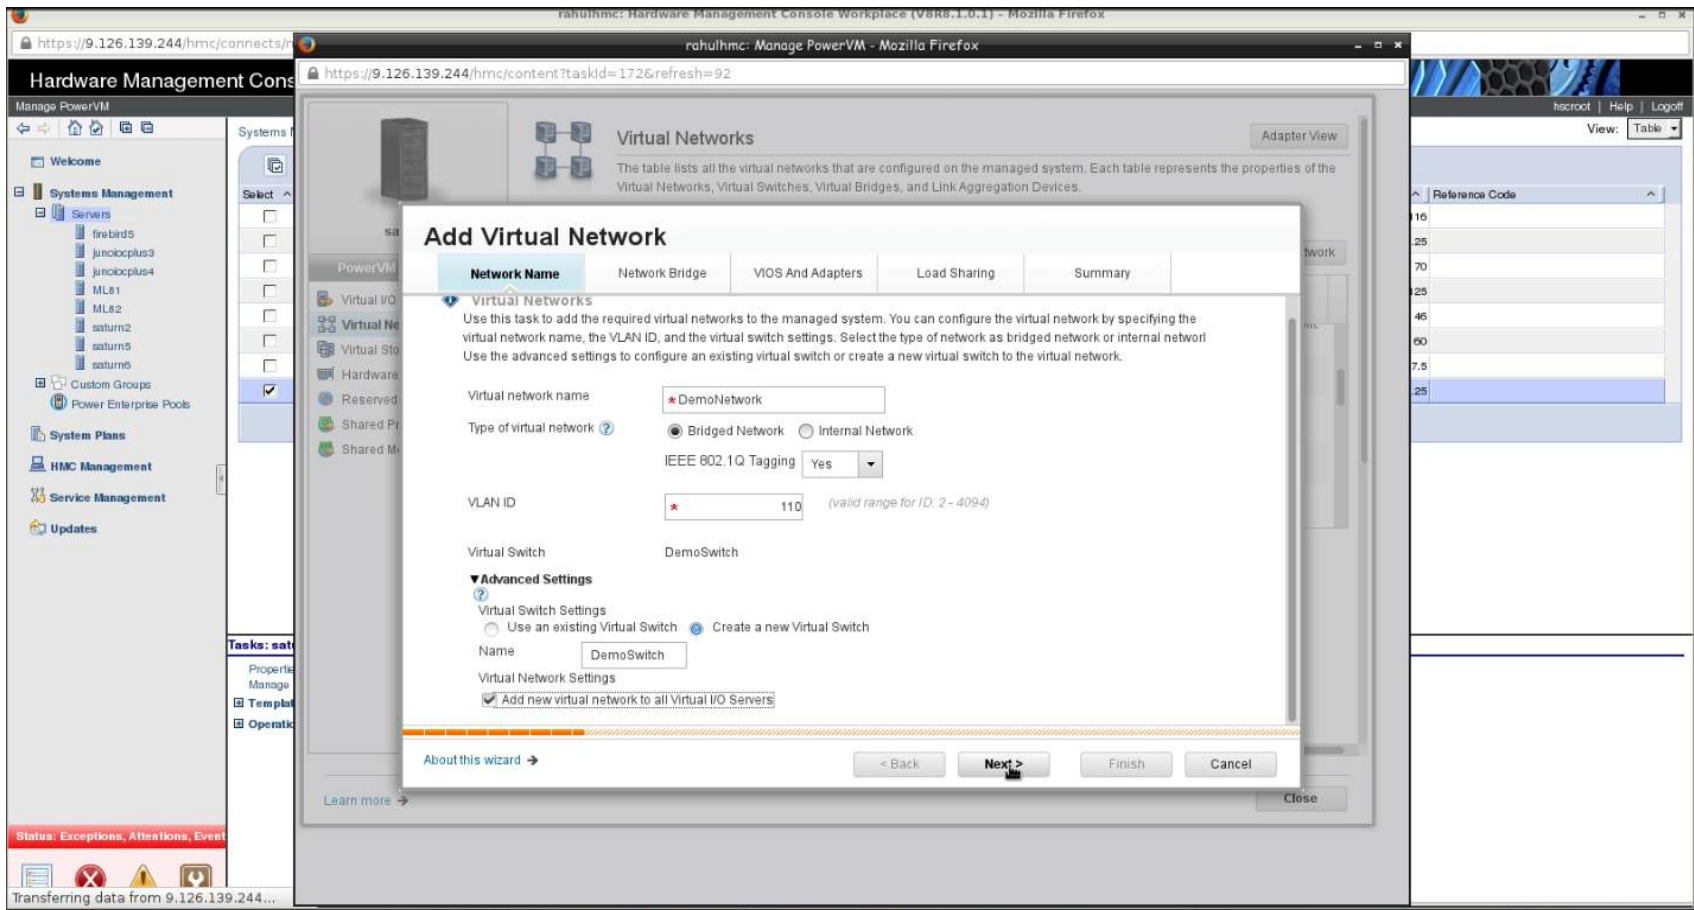
Name	ID	RMC Connection	Status	Virtual Processors	Processor Entitlement	Memory in MB	VIOS Version
VIOS3329	33	Inactive	Open Firmware	-	2.0	2048	License not accepted
vios1	26	Inactive	Not Activated	-	0.0	0	License not accepted
vios1_ns_DONT_TO_1 UCH	1	Active	Running	-	1.0	2048	VIOS 2.2.3.3
vios2_ns_DON_UCH			Running	-	1.0	2048	VIOS 2.2.3.3
vios3_s_DON_UCH				-	1.0	2048	VIOS 2.2.3.3
vios4_s_DON_UCH				-	1.0	2048	VIOS 2.2.3.3

The context menu for 'vios4\_s\_DON\_UCH' includes the following options:

- Manage
  - Change Default Profile
- Operations
  - Restart
  - Shutdown
  - Schedule Operations
- Configuration
- Console Window



# HMC New Management GUI - Managing Virtual Networks

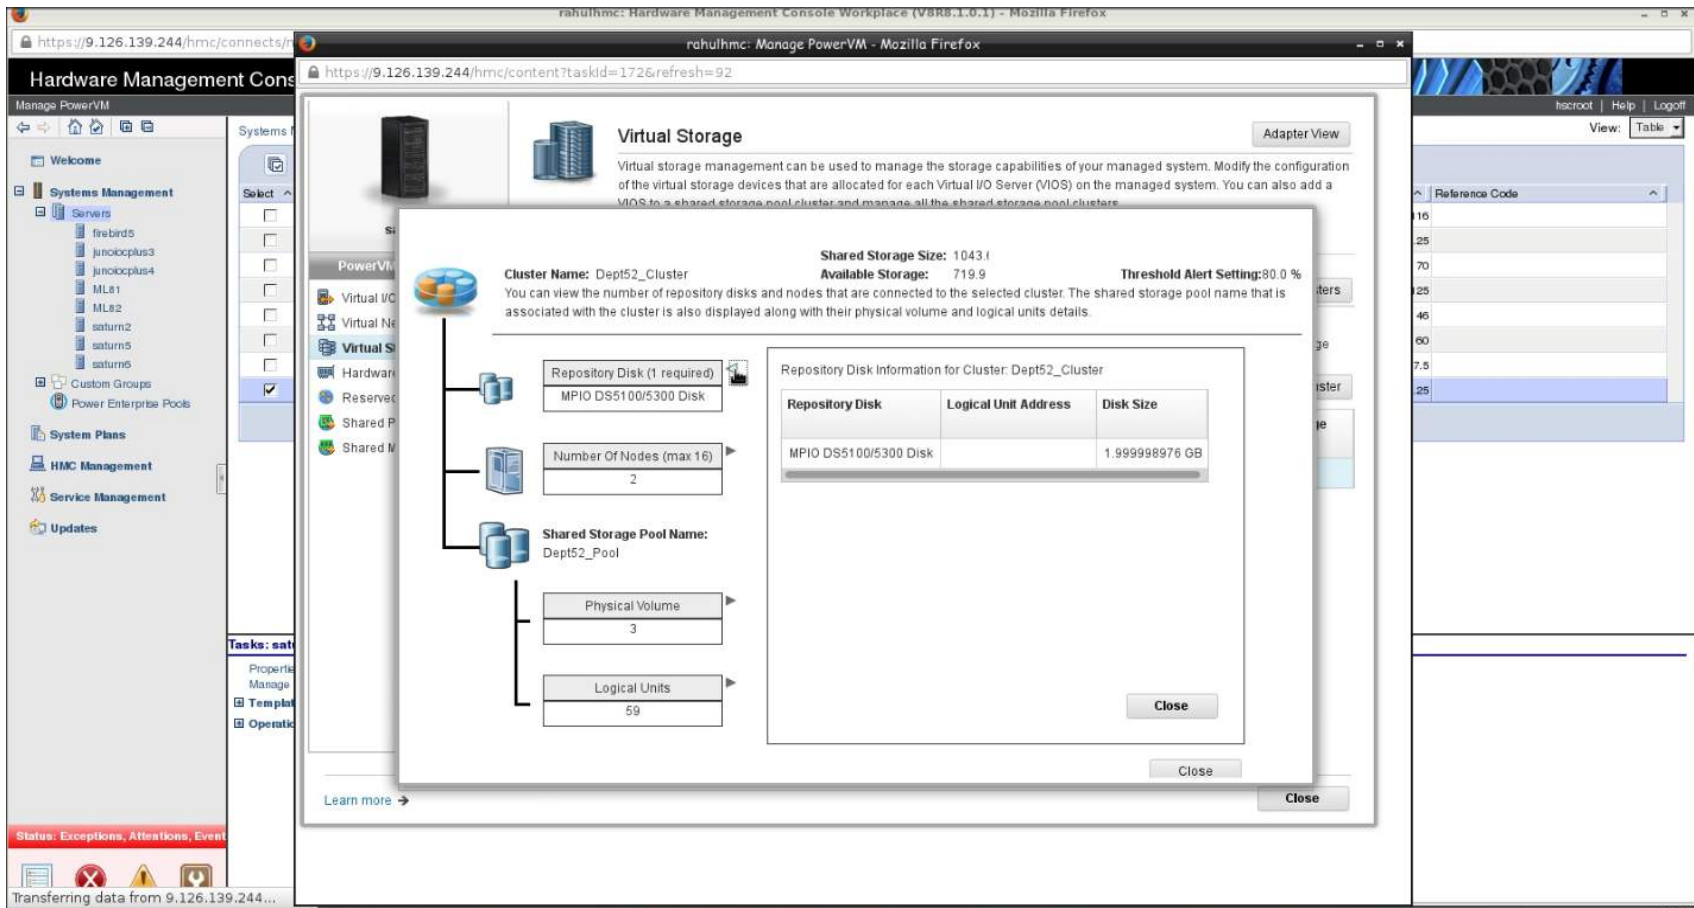


The screenshot displays the HMC New Management GUI interface. The main window is titled "Virtual Networks" and contains a table listing virtual networks. A modal window titled "Add Virtual Network" is open, showing the following configuration details:

- Network Name:** DemoNetwork
- Type of virtual network:** Bridged Network (selected)
- IEEE 802.1Q Tagging:** Yes
- VLAN ID:** 110 (valid range for ID: 2 - 4094)
- Virtual Switch:** DemoSwitch
- Advanced Settings:**
  - Virtual Switch Settings:** Create a new Virtual Switch (selected)
  - Name:** DemoSwitch
  - Virtual Network Settings:** Add new virtual network to all Virtual I/O Servers (checked)

The wizard includes navigation buttons: < Back, Next >, Finish, and Cancel. A "Close" button is also visible at the bottom right of the modal window.

# HMC New Management GUI - Managing Virtual Storage



The screenshot displays the HMC New Management GUI interface. The main window is titled "Virtual Storage" and provides information for a cluster named "Dept52\_Cluster".

**Virtual Storage Summary:**

- Shared Storage Size: 1043.1
- Available Storage: 719.9
- Threshold Alert Setting: 80.0 %

**Cluster Name:** Dept52\_Cluster

You can view the number of repository disks and nodes that are connected to the selected cluster. The shared storage pool name that is associated with the cluster is also displayed along with their physical volume and logical units details.

**Repository Disk Information for Cluster: Dept52\_Cluster**

Repository Disk	Logical Unit Address	Disk Size
MPIO DS5100/5300 Disk		1.999998976 GB

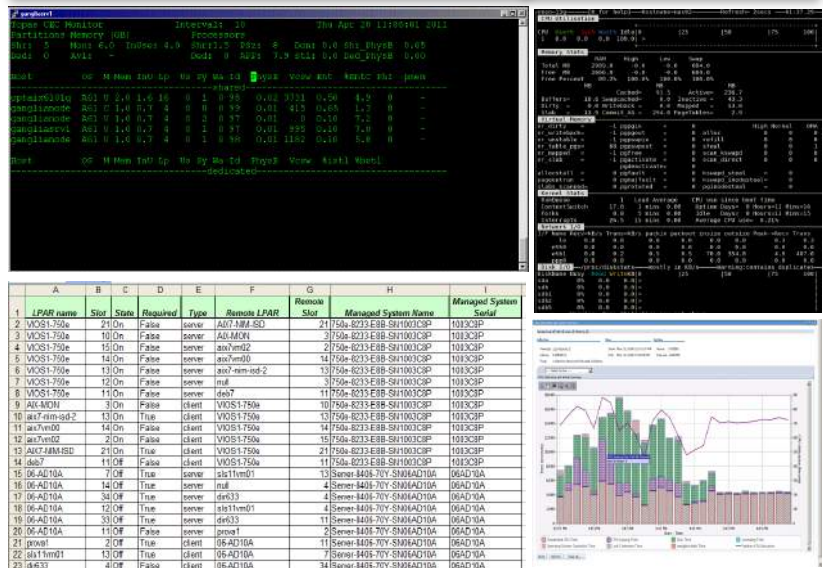
**Other Configuration Details:**

- Repository Disk (1 required): MPIO DS5100/5300 Disk
- Number Of Nodes (max 16): 2
- Shared Storage Pool Name: Dept52\_Pool
- Physical Volume: 3
- Logical Units: 59

The interface includes a sidebar with navigation options like "Systems Management", "System Plans", and "HMC Management". A status bar at the bottom indicates "Status: Exceptions, Attentions, Event" and "Transferring data from 9.126.139.244...".

# HMC Performance Monitoring

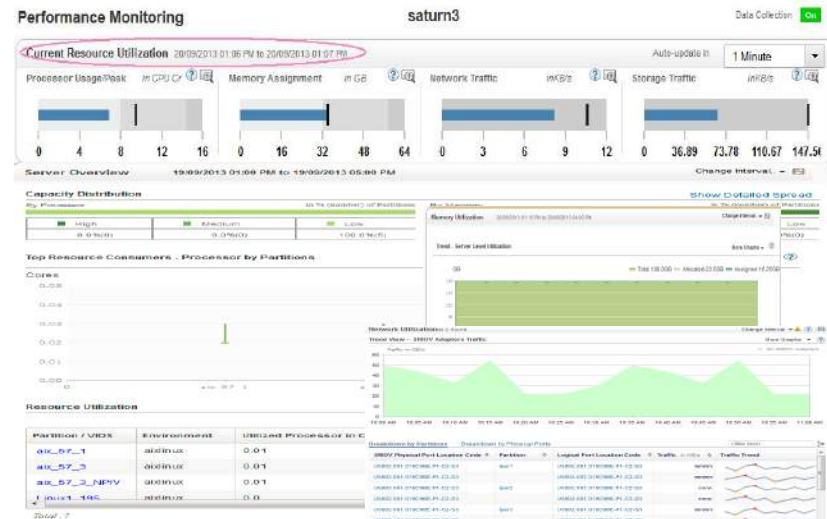
## HMC Today



LPAR name	Slot	State	Required	Type	Remote LPAR	Remote Slot	Managed System Name	Managed System Serial
1 VIOS1-750a	21	On	False	server	AD7-IBM-SD	21	750a-R213-E88-SM1001CP	1013C3P
2 VIOS1-750a	10	On	False	server	AD1-IBM-SD	10	750a-R213-E88-SM1001CP	1013C3P
3 VIOS1-750a	15	On	False	server	an7-m02	15	750a-R213-E88-SM1001CP	1013C3P
4 VIOS1-750a	14	On	False	server	an7-m00	14	750a-R213-E88-SM1001CP	1013C3P
5 VIOS1-750a	13	On	False	server	an7-mm-nsf-2	13	750a-R213-E88-SM1001CP	1013C3P
6 VIOS1-750a	12	On	False	server	an7	12	750a-R213-E88-SM1001CP	1013C3P
7 VIOS1-750a	11	On	False	server	an7	11	750a-R213-E88-SM1001CP	1013C3P
8 VIOS1-750a	3	On	False	client	VIOS1-750a	10	750a-R213-E88-SM1001CP	1013C3P
9 ALL-ADMN	13	On	True	client	VIOS1-750a	11	750a-R213-E88-SM1001CP	1013C3P
10 an7-mm-nsf-2	13	On	True	client	VIOS1-750a	13	750a-R213-E88-SM1001CP	1013C3P
11 an7-m00	14	On	False	client	VIOS1-750a	14	750a-R213-E88-SM1001CP	1013C3P
12 an7-m02	15	On	False	client	VIOS1-750a	15	750a-R213-E88-SM1001CP	1013C3P
13 AD7-IBM-SD	21	On	True	client	VIOS1-750a	21	750a-R213-E88-SM1001CP	1013C3P
14 an7	11	On	False	client	VIOS1-750a	11	750a-R213-E88-SM1001CP	1013C3P
15 06-AD10A	7	Off	True	server	sis11m01	13	Server-848-76Y-SM06AD10A	06AD10A
16 06-AD10A	14	Off	True	server	an7	4	Server-848-76Y-SM06AD10A	06AD10A
17 06-AD10A	34	Off	True	server	an7	4	Server-848-76Y-SM06AD10A	06AD10A
18 06-AD10A	12	Off	True	server	sis11m01	4	Server-848-76Y-SM06AD10A	06AD10A
19 06-AD10A	33	Off	True	server	an7	11	Server-848-76Y-SM06AD10A	06AD10A
20 06-AD10A	11	Off	False	server	an7	2	Server-848-76Y-SM06AD10A	06AD10A
21 an7	2	Off	True	client	06-AD10A	11	Server-848-76Y-SM06AD10A	06AD10A
22 an11m01	13	Off	True	client	06-AD10A	7	Server-848-76Y-SM06AD10A	06AD10A
23 an63	4	Off	False	client	06-AD10A	34	Server-848-76Y-SM06AD10A	06AD10A



## HMC w/ new enhancements



**Performance Monitoring** saturn3

Current Resource Utilization: 2013/03/13 01:06 PM to 2013/03/13 01:07 PM

Auto-update: 1 Minute

Processor Usage/Peak: 0 4 8 12 16  
Memory Assignment: 0 16 32 48 64  
Network Traffic: 0 3 6 9 12  
Storage Traffic: 0 36.89 73.78 110.67 147.54

Server Overview: 19/03/2013 01:00 PM to 19/03/2013 05:00 PM

Capacity Distribution: High, Medium, Low

Top Resource Consumers - Processor by Partitions

Partition / LPAR	Usage	Usage Percentage	Used/Total Processor
06_07_1	0.01	0.01%	0.01
an7_3	0.01	0.01%	0.01
an7_3_NPV	0.01	0.01%	0.01
an7_100	0.01	0.01%	0.01



- Disjoint set of tools
- Multiple agents need to be installed in OS
- Minimal or Lack of Visualization

- Integrated Visual Monitor in HMC
- Standard set of Interfaces (REST APIs) for external APIs to consume data

## Integrated Performance & Capacity Metrics

Provide full PowerVM performance and capacity metrics via a single touch-point (HMC).

- Data collection and aggregation of performance metrics from Hypervisor and VIOS
- Identify performance bottlenecks and facilitate early problem detection and notification
- REST based API to access all platform (Hypervisor & VIOS) metrics for IBM Internal and Third Party tools
- Comprehensive Visualization to view **aggregated historical performance data**
- It is not a real time performance monitoring tool. It provides snapshots of resource utilization every minute
- Main objective is performance management and capacity planning - not monitoring of specific process/application in detail
-

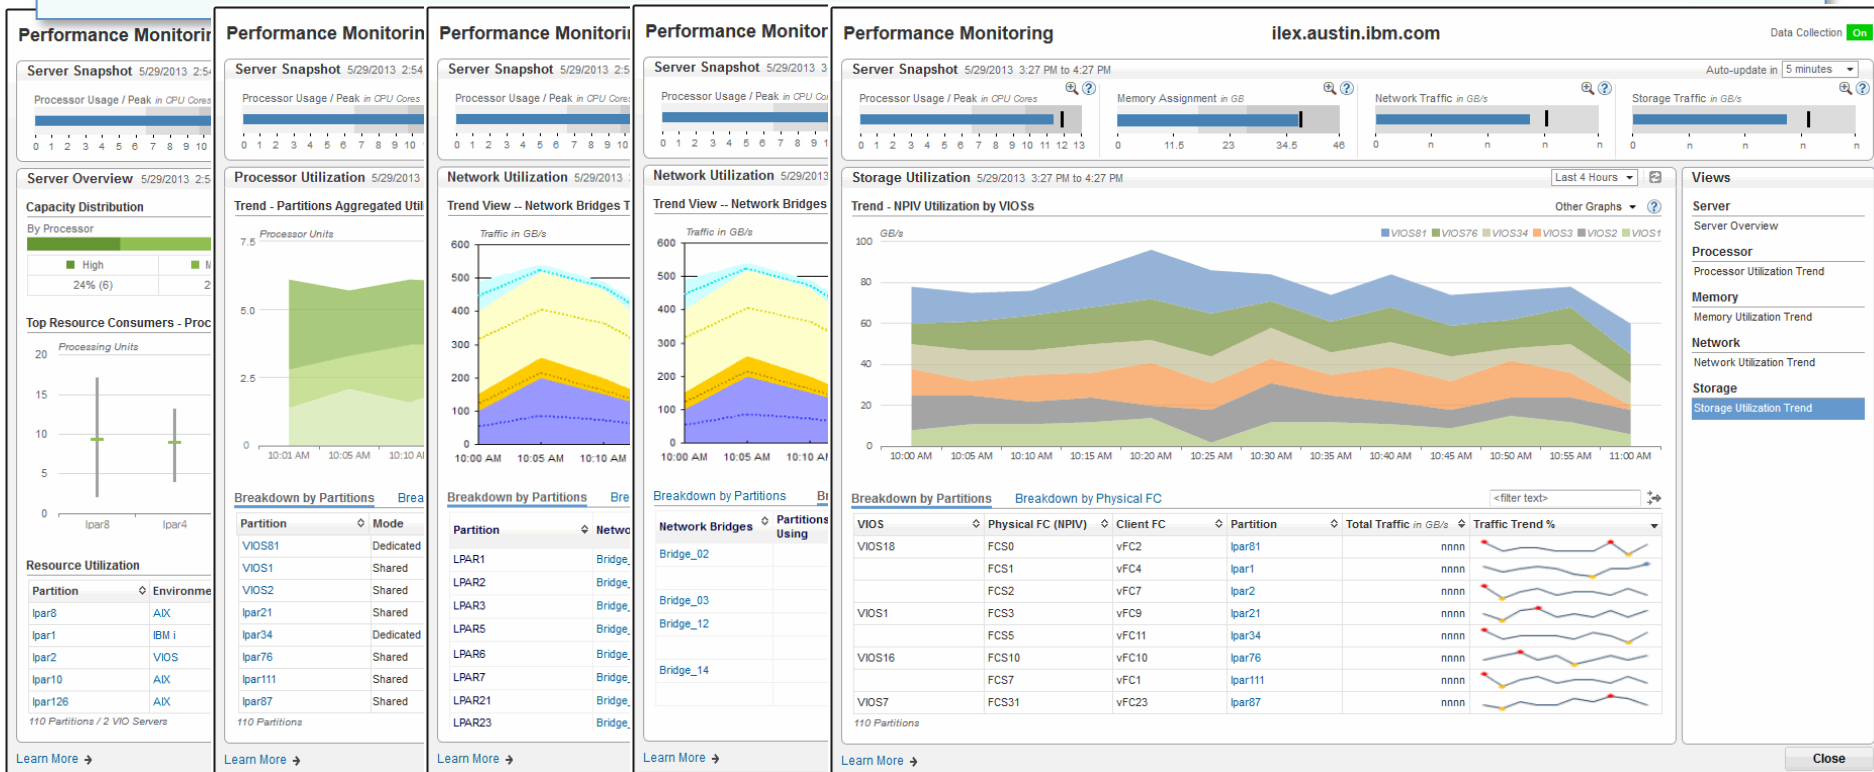
## Performance & Capacity Metrics Details

Provides both Compute & IO Resources Utilization

- Following new metrics available and exposed through new HMC APIs
  - Physical System Level Processor & Memory Resource Usage Statistics
    - System Processor Usage Statistics (w/ LPAR, VIOS & Power Hypervisor usage breakdown)
    - System Dedicated Memory Allocation and Shared Memory Usage Statistics (w/ LPAR, VIOS & Power Hypervisor usage breakdown)
  - Advanced Virtualization Statistics
    - Per LPAR Dispatch Wait Time Statistics
    - Per LPAR Placement Indicator (for understanding whether the LPAR placement is good / bad based on score)
  - Virtual IO Statistics
    - Virtual IO Server' s CPU / Memory Usage (Aggregated, Breakdown)
    - SEA Traffic & Bandwidth usage Statistics (Aggregated & Per Client, Intra/Inter LPAR breakdown)
    - NPIV Traffic & Bandwidth usage Statistics (HBA & Per Client breakdown)
    - vSCSI Statistics (Aggregated & Per Client Usage)
    - VLAN Traffic & Bandwidth usage Statistics (Adapter & LPAR breakdown)
    - 1 • SRIOV Traffic & Bandwidth usage Statistics (Physical & Virtual Function Statistics w/ LPAR
    - 2 breakdown)
    - 0



# Integrated Performance Monitoring in HMC




- Resource utilization dashboard
  - Last few minutes Processor, memory & I/O utilization snapshot (1, 5, 10, 15 minutes refresh)
  - Server & LPAR level Capacity Distribution, Top Resource consumption information
  - Basic Utilization Trends for all resources w/ LPAR & Resource level breakdowns

## HMC Management Demo Videos


- Managing partitions
  - [ManageLpar-1-4.avi](#)
- Managing VIOS
  - ManagePowerVM-1.avi

# PowerVM Live Partition Mobility Performance Enhancements

- LPM performance has been continuously improving
  - 5 times faster in 2014 than in 2011



- VIOS 2.2.3 brought significant performance enhancements for LPM (2013)
  - Ability to achieve full bandwidth on a 10G adapter when transferring memory



## PowerVM Live Partition Mobility Performance Enhancements

- HMC V8R8.1.0 brings further enhancements by optimizing the device selection for LPM
  - During a Live Partition Mobility activity the HMC makes *find\_device* calls to the Virtual I/O Server involved in the migration when virtual SCSI (vSCSI) device mapping is specified.
  - In previous versions of HMC the number of *find\_device* calls was equal to the number of vSCSI adapters in each Virtual I/O Server.
  - In HMC V8R8.1.0 the logic has been improved so that only one *find\_device* call is made to each Virtual I/O Server. Depending on the version of Virtual I/O Server and the configuration of the backing devices each *find\_device* call could range between one and 15 seconds to complete, so a reduction in the number of *find\_device* calls improves Live Partition Mobility performance.

## Fix Level Recommendation Tool - LPM Validation

- The FLRT Live Partition Mobility (LPM) report provides recommendations for LPM operations based on source and target input values. These recommendations may include recommended fixes, including interim fixes, for known LPM issues. These recommendations are based on testing and field experience. The levels provided in the input screen are based on the supported levels for the selected machine type models.
- <https://www14.software.ibm.com/webapp/set2/flrt/home>

# Fix Level Recommendation Tool - LPM Validation

## Fix Level Recommendation Tool

The FLRT Live Partition Mobility (LPM) report provides recommendations for LPM operations based on source and target input values. These recommendations may include recommended fixes, including interim fixes, for known LPM issues. These recommendations are based on testing and field experience. The levels provided in the input screen are based on the supported levels for the selected machine type models. To view minimum requirements for Power systems for AIX, Red Hat Linux, and SUSE Linux visit the [System Software Maps](#) page. [Clear all](#)

Power systems
  LPM report
  Other systems
  Load saved inventory

### Source System

<input checked="" type="checkbox"/> Server Machine-Type-Model	9117-MMA	
<input checked="" type="checkbox"/> System firmware	EM350_159	
<input checked="" type="checkbox"/> Hardware Management Console	V7 R770	
<input checked="" type="checkbox"/> PowerVM Virtual I/O Server	2.2.2.2	
<input checked="" type="checkbox"/> OS Family	AIX	
<input checked="" type="checkbox"/> AIX	6100-09-01	

### Target System

<input checked="" type="checkbox"/> Server Machine-Type-Model	8284-22A	
<input checked="" type="checkbox"/> System firmware	SV810_061	
<input checked="" type="checkbox"/> Hardware Management Console	V8 R810	
<input checked="" type="checkbox"/> PowerVM Virtual I/O Server	2.2.3.3	

Name of report (optional):

LPM

Submit



Date: 2014.06.24 Database last updated 2014.06.22  
 Report name: LPM  
 Source MTM: 9117-MMA  
 Target MTM: 6264-22A

[+ Expand all](#) [- Collapse all](#) [PDF URL](#) [PDF page](#)

**Source system** [-] [D]

	Input level	Recommended level and updates
 System firmware Release date EoS/PS	<a href="#">EM350_159</a> 2014.06.25 Last	None
 HMC Software Release date EoS/PS	<a href="#">V7 R770</a> 2013.02.20 2015.02.26	<a href="#">V7 R770 SP4</a> 2014.07.21 2015.02.26
 AIX Release date EoS/PS	<a href="#">6100-09-01</a> 2013.11.19 2016.11.19	<a href="#">mcpk15.140306.epkg.Z</a> 2014.04.22

▼ Notes

[mcpk15.140306.epkg.Z](#) : [mcpk15.140306.epkg.Z](#) sets the default RMC packet size to 1500 bytes to avoid MTU imbalance problems. This fix does not require a reboot.

 VIOS Release date EoS/PS	<a href="#">2.2.2.2</a> 2013.02.20 2015.10.01	<a href="#">2.2.3.4</a> 2014.02.21 2015.10.01  <a href="#">IV55956m4a.140226.epkg.Z</a> 2014.04.22 <a href="#">mcpk14.131211.epkg.Z</a> 2014.04.22
--	---	---

▼ Notes

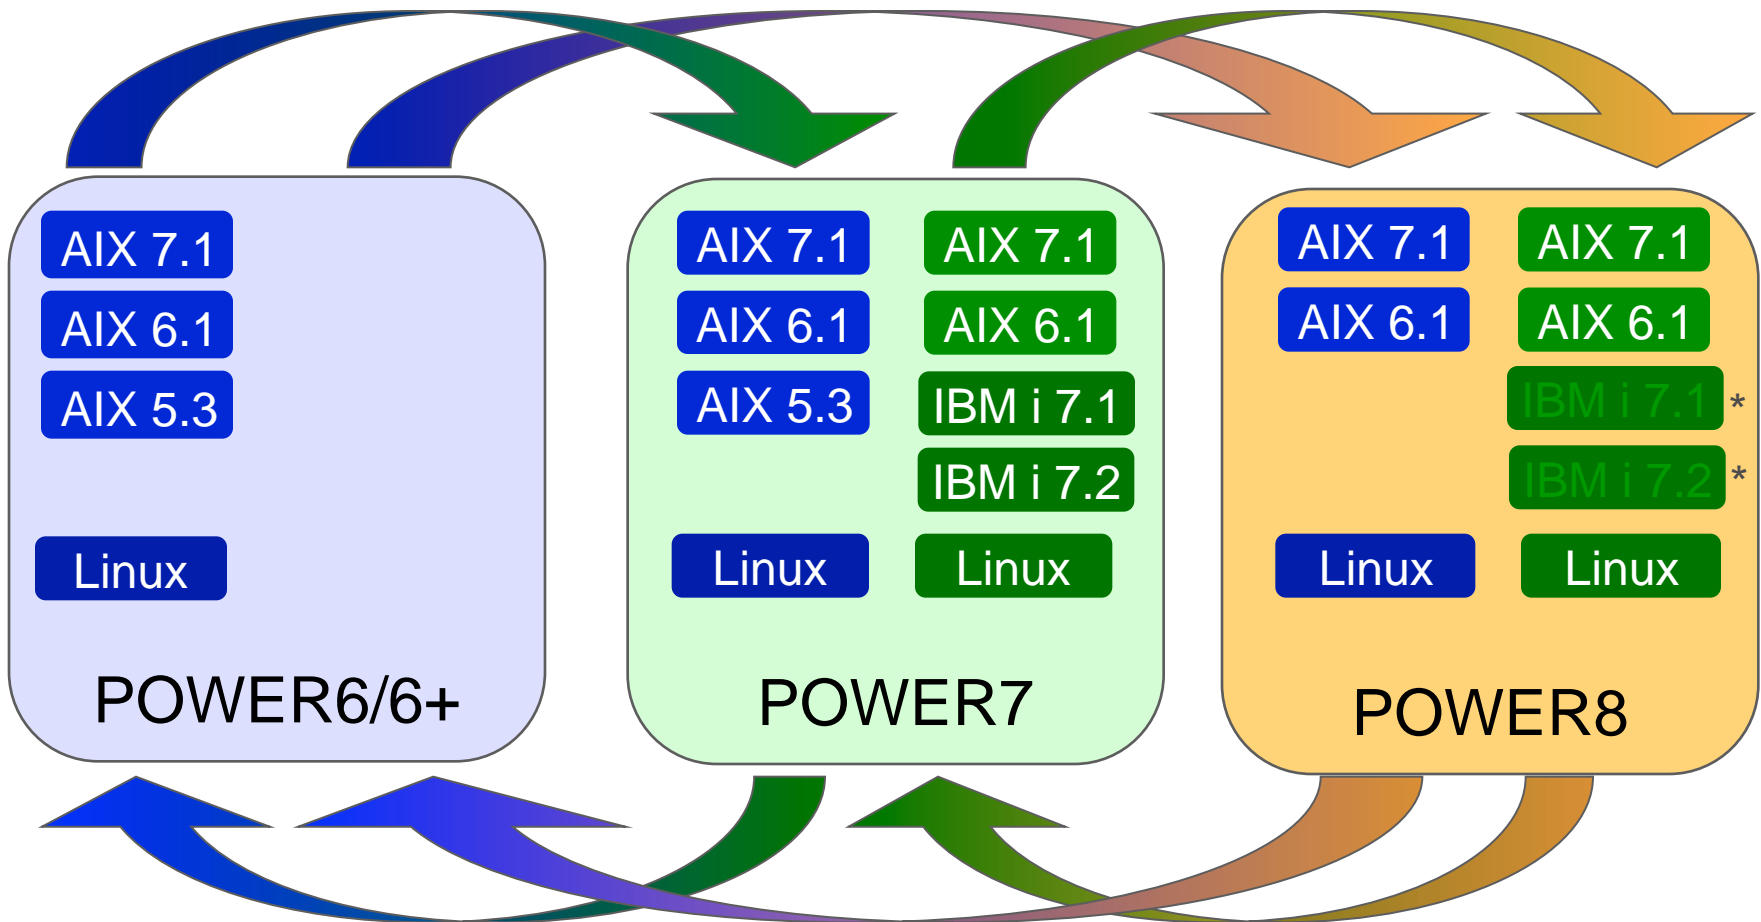
[IV55956m4a.140226.epkg.Z](#) : [IV55956m4a.140226.epkg.Z](#) is only required if using NPIV.

[mcpk14.131211.epkg.Z](#) : [mcpk14.131211.epkg.Z](#) sets the default RMC packet size to 1500 bytes to avoid MTU imbalance problems. This fix does not require a reboot.

**Target system** [-] [D]

	Input level	Recommended level and updates
 System firmware Release date EoS/PS	<a href="#">SV510_061</a> 2014.07.29 Not Announced	None
 HMC Software Release date EoS/PS	<a href="#">V8 R610</a> 2014.06.10 Not Announced	<a href="#">MH01421</a> 2014.06.10
 VIOS Release date EoS/PS	<a href="#">2.2.3.3</a> 2014.06.12 2016.10.01	None

# POWER6 / POWER7 / POWER8 Partition Mobility



Leverage POWER6 / POWER7 Compatibility Modes

LPAR Migrate between POWER6 / POWER7 / POWER8 Servers

Can not move POWER8 Mode partitions to POWER6 or POWER7 systems.



# OS Support



<b>POWER6 and POWER6+ MODES</b>	<b>POWER7 MODE *</b>	<b>POWER8 MODE</b>
<b>2-Thread SMT</b>	<b>4-Thread SMT, IntelliThreads</b>	<b>8-Thread SMT</b>
<b>8 Protection Keys *(16 in P6+ Mode)</b>	<b>32 Protection Keys User Writeable AMR</b>	<b>32 Protection Keys User Writeable AMR</b>
<b>VMX (Vector Multimedia Extension / AltiVec)</b>	<b>VSX (Vector Scalar Extension)</b>	<b>VSX2, In-Core Encryption Acceleration</b>
<b>Affinity OFF by Default</b>	<b>CPU/Memory Affinity Enhancements ON by Default, HomeNode, 3-tier Memory, MicroPartition Affinity</b>	<b>HW Memory Affinity Tracking Assists, Micropartition Prefetch, Concurrent LPARs per Core</b>
<b>64-core/128-thread Scaling</b>	<b>64-core / 256-thread Scaling 256-core / 1024-thread Scaling</b>	<b>&gt; 1024-thread Scaling Hybrid Threads Transactional Memory Active System Optimization HW Assists</b>
<b>N/A</b>	<b>Active Memory Expansion</b>	<b>HW Accelerated/Assisted Active Memory Expansion</b>
<b>N/A</b>	<b>P7+ : AME compression acceleration and Encryption acceleration</b>	<b>Coherent Accelerator / FPGA Attach</b>

\* POWER7+ mode does not exist - POWER7 and POWER7+ processors use the same processor mode



# AIX Support of Scale-out Servers

	11 / 2012	2 / 2012	3 / 2013	5 / 2013	8 / 2013	9 / 2013	10 / 2013	12 / 2013	2Q / 2014	3Q / 2014
<b>AIX 6 TL7</b>	SP6		SP7			SP8		SP9		SP10
<b>AIX 6 TL8</b>	SP1	SP2				SP3		SP4		SP5
<b>AIX 6 TL9</b>							SP1		SP3	
<b>AIX 7 TL1</b>	SP6			SP7	SP8			SP9		SP10
<b>AIX 7 TL2</b>	SP1	SP2								SP5
<b>AIX 7 TL3</b>							SP1		SP3	

**P7 or P6 Modes with Virtual I/O**

**P7 or P6 Modes with Full I/O Support**

**P8, P7 or P6 Modes with Full I/O Support**


# AIX 7.1 TL 3 SP 3 Exploitation of POWER8 Features (June 2014)

- Security Features
  - Random number generator
  - On-chip crypto
  - In-core crypto
- Performance Optimization
  - Transactional Memory enablement for middleware and applications
  - New POWER8 instructions
  - Dynamic Code Optimization (for Java and IBM JIT technology)
  - SMT 8 (SMT 4 as default mode)
  - 24x7 Monitor
- Virtualization
  - Compression accelerator (used by Active Memory Expansion)



# AIX exploitation of POWER8 accelerators

## AIX OpenSSL Acceleration with POWER8 in-core Crypto Support



Results from AIX Development internal testing

- In-core Crypto is enabled on AIX 7.1.3.3 with OpenSSL version 1.0.1.510
- IBM GSKit 8.0.50.X with HW crypto Support Planned for 3Q 2014
- IBM Software products which embed GSKit will pick up HW acceleration as they qualify 8.0.50.X

# IBM i 7.1 TR8      Cores and Threads for P7/P8

	IBM i 7.1
<b>POWER7</b>	Max Scale = 32 cores (SMT4) Max Partition = <b>64 cores (SMT4)</b> Threads = ST, SMT2, SMT4 up to 256 threads in single partition
<b>POWER8</b>	Max Scale = 32 cores (SMT8) Max Partition = <b>64 cores (SMT4)</b> Threads = ST, SMT2, SMT4, SMT8 up to 256 threads in single partition

### Thread Model:

IBM i 7.1 will support ST, SMT2, SMT4, and SMT8 and will remain limited to 256 threads in a partition. The ability to switch SMT mode without IPL will be PTFed back to IBM i 7.1.

# IBM i v7.2      Cores and Threads for P7/P8

	IBM i 7.2
<b>POWER 7</b>	<p>Max Scale = 32 cores (SMT4)</p> <p>Max Partition = <b>96 cores (SMT4)</b></p> <p>Threads = ST, SMT2, SMT4 up to 384 threads in single partition</p>
<b>POWER 8</b>	<p>Max Scale = 48 cores (SMT8)</p> <p>Max Partition = <b>96 cores (SMT8)</b></p> <p>Threads = ST, SMT2, SMT4, SMT8 up to 768 threads in single partition</p>

### Thread Model:

IBM i 7.2 will support ST, SMT2, SMT4 and SMT8, up to a limit of 768 threads in a single partition. It will support switchable SMT (not requiring a partition IPL to change modes)



# IBM i 7.2

## What's new:

- DB2 for i significant enhancements in the areas of Security, Performance and Usability
- IBM Navigator for i provides system and message monitoring as well as batch process capture and modeling
- Rational Developer for i supports Free Format RPG
- IBM Integrated Application Server built on WebSphere Liberty profile
- POWER8 and POWER7+ refreshed firmware - 770 & 780 supported by IBM i 7.2
- Support added for USB 3.0, SRIOV, WAN over LAN
- PowerHA Express Edition with HyperSwap
- BRMS Hub Function
- SSO application suite extended to include FTP and Telnet authenticating with Kerberos

## Features / Business Value:

- Client enjoy new features and functions of DB2, WebSphere, Security, Systems Management, Business resiliency and development tools/compiler
  - ❖ Fully integrated and tested
- Secures critical business data

## Client Benefits:

- Provides support for new features of POWER8 and POWER7 systems
- Protects critical business data
  - ❖ Ensures security of critical business data which is critical in today's environment of mobile, web and social access
- Improves system management – ease-of-use and easier problem resolution
- Improved business resiliency – BRMS and PowerHA
- Many additional benefits

# IBM i 7.2 Highlights

Powerful new features of DB2® for i ensure security of the data in a modern environment of mobile, social and network access

- Role based security
- Temporal support now included in DB2 for i

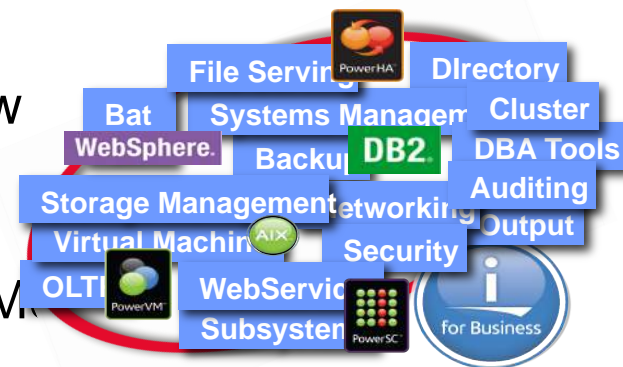
IBM Navigator for i extends system management capabilities to manage and monitor performance services

Announcing a new resiliency option; PowerHA Express for i with HyperSwap

Additional flexibility in configuration for I/O with new virtualization: Ethernet through SRIOV and WAN' over LAN"

Backup and Recovery Management Systems (BRM) introduces a Hub function for easier management

Support for new POWER8 systems, and new features for POWER8 and POWER7



# Linux on Power Systems: A powerful combination....

## IBM Power Systems

Optimized for big data and analytics  
 Workload-based resource allocation  
 Automated, policy-based security



## Linux Operating System

Rich opportunities for innovation  
 Significant cost savings  
 Industry-leading flexibility & performance



## Delivering leadership solutions across IBM

- Optimized for mobile, big data, analytics and social
- Industry specific capabilities and reference configurations

## Building a robust ISV ecosystem

- Power Systems Linux Centers worldwide
- Strong development support for community

## Delivering industry leading virtualization

- Commitment to OpenStack & wide range of server deployment options
- Platform, virtualization & cloud management solutions



IBM Software





# Linux on Power Systems Solutions

**IFL: Power Cores & Memory dedicated for Linux workloads**

**PowerLinux 7R1**



- 1-socket, 2U
- POWER7+ processor
- Up to 8 cores
- 256 GB memory
- Linux only

**PowerLinux 7R2**



- 2-socket, 2U
- POWER7+ processor
- Up to 16 cores
- 512 GB memory
- Linux only

**PowerLinux 7R4**



- 4-socket, 5U
- Up to 32 POWER7+ cores
- 1 TB memory
- Hot-swap PCI adapters
- Linux only

**1, 2, and 4 socket systems**



**New**

**Power 770 IFL**



**Power 780 IFL**



**Power 795 IFL**



**Power Enterprise IFLs**

**Available for POWER7 & POWER7+ Models  
Power 770 / 780 / 795**



# Power Enterprise Linux Engines

## Power IFL: Power Integrated Facility for Linux

- Power Cores dedicated for Linux workloads
- Discounted pricing for Cores, Memory, and PowerVM



## Structure of the offering

- Single Feature Code: 4 Cores & 32 GB memory
- Available for Power 770, Power 780 and Power 795
- PowerVM EE licensed for Power IFLs
- HWMA and SWMA for PowerVM for Linux priced separately
- Can have multiple FC's on a single server

## T's & C's

- Clients must sign a contract addendum agreeing to segregation of cores and memory for Linux activities
- No converting to AIX & IBM i environments

## Linux license acquired separately



Power  
770



Power  
780



Power  
795

# Power IFL Structure

## Power IFL Cost Savings

**Lower Processor core activations**  
**Lower Memory activations**  
**Lower PowerVM Costs**



**Lower Software Cost PVU = 70**



**Lower Software Maint**  
**For PowerVM for PowerLinux**




**Lower cost core Hardware Maint after warranty**



# Power Systems' Linux momentum is growing

## Ecosystem Leadership




**OpenPOWER Foundation**  
Open development alliance based on POWER hardware and software

## IBM Software Optimization




**IBM WATSON™**



**400+ Global Linux ISVs**  
200+ new ISVs across GCG with 400 applications


**150+ SWG Optimized Offerings**

WebSphere	Rational	Industry Solutions
Tivoli	Cognos	Security
Information Management	SPSS	



**Power Systems Linux Centers**

- Beijing
- Austin
- New York
- Montpellier



**Industry-Leading Benchmarks**

- Over twice the throughput of x86/VMware with TPOX on DB2
- 36% faster than x86 with SPEC Java on WebSphere

A new generation of **IBM Power Systems** optimized for scale-out data and cloud infrastructures



## Linux-only systems

offer better economics for Linux deployments

## Ubuntu Server for Power

brings the fastest growing cloud OS to Power Systems

## Power KVM

simplifies management of heterogeneous data centers

# Target workloads for Linux on Power Systems

## Virtualized Application Infrastructure

Java, PHP, web, networking

## Analytics & Research

Business intelligence, structured & unstructured data, predictive analytics, HPC

## Data Services

Database, enterprise content management

## Business Applications

Mobile data access, ERP, cloud delivery, application development

## Workload Requirements

Highly threaded  
Throughput oriented  
Scale out capable  
Economic efficiency

Compute intensive  
High memory bandwidth  
Floating point  
High I/O rates

Handle peak workloads  
Scalability  
High quality of service  
Resiliency and security

High quality of service  
Scalability  
Flexible infrastructure  
Large memory footprint



## What Linux Distributions in various Power Environments?

<i>Linux</i>	<i>Release</i>	<i>Endian</i>	<i>Dedicated LPAR</i>	<i>PowerVM Guest</i>	<i>PowerKVM Guest</i>
<b>Redhat</b>	<b>5.10</b>	<b>Big</b>	<b>Y</b>	<b>Y</b>	<b>N</b>
<b>Redhat</b>	<b>6.4</b>	<b>Big</b>	<b>Y</b>	<b>Y</b>	<b>N</b>
<b>Redhat</b>	<b>6.5</b>	<b>Big</b>	<b>Y</b>	<b>Y</b>	<b>Y</b>
<b>SUSE</b>	<b>11 SP3</b>	<b>Big</b>	<b>Y</b>	<b>Y</b>	<b>Y</b>
<b>Ubuntu*</b>	<b>14.04</b>	<b>Little</b>	<b>N</b>	<b>N</b>	<b>Y</b>

\*Exploits P8

1. Select the applications you want to run on Linux on Power
2. Then look at the Linux distributions that are available for those apps
3. Pick your Linux distribution of choice



# Linux for Scale-out Servers

	1H / 2014	2H / 2014
<b>RHEL6</b>	<b>RHEL 6.5 POWER7 Mode on POWER8</b>	<b>RHEL 6.6 POWER7 mode on POWER8</b>
<b>RHEL 7</b>	<b>RHEL 7.0</b>	
<b>SLES 11</b>	<b>SLES 11 + SP3 POWER7 Mode on POWER8</b>	
<b>SLES 12</b>		<b>SLES 12 (LE) POWER8 Support</b>
<b>Ubuntu (LE)</b>	<b>14.04.00/01 POWER8 Support</b>	<b>14.04.00/02</b>
<b>Debian</b>	<b>LE Introduction POWER8 Support</b>	<b>LE Update</b>
<b>VIOS</b>	<b>2.2.3.3</b>	<b>2.2.3.X</b>

# Learn more about Linux on Power Systems

## Power Systems Linux Portal

(Product Information)

[www.ibm.com/systems/power/software/linux/](http://www.ibm.com/systems/power/software/linux/)



[@ibmpowerlinux](https://twitter.com/ibmpowerlinux)



## The PowerLinux Community (developerWorks)

[www.ibm.com/developerworks/group/tpl/](http://www.ibm.com/developerworks/group/tpl/)



[plus.google.com/communities/100156952249293416679](https://plus.google.com/communities/100156952249293416679)

# OpenPower















ENTERPRISE

google

IBM

Intel

microchips

## This Google Motherboard Means Trouble for Intel

BY ROBERT MCMILLAN 04.29.14 | 3:15 PM | PERMALINK

Facebook Share 1.4k | Twitter Tweet 604 | Google+ 283 | LinkedIn Share 565 | Pinterest Pin it 1




Image: Google

# OpenPOWER™ Foundation Members



## How will the OpenPOWER Consortium benefit clients?

- New Innovators on the Power platform will create more value
- OpenPOWER technology creates greater choice for our customers
- Innovation will increase adoption of the Power Systems platform



## What does this mean to the industry?

- Game changer on the competitive landscape of the server industry
- Will enable and drive innovation in the industry
- Provide more choice in the industry



# PowerKVM



## Power Virtualization Options

### PowerKVM



Initial Offering: Q2 2014

- **PowerKVM: Open Source option for virtualization on Power Systems for Linux workloads.**
- **For clients that have Linux centric admins. (RHEL 6.5 & SLES 11.3)**

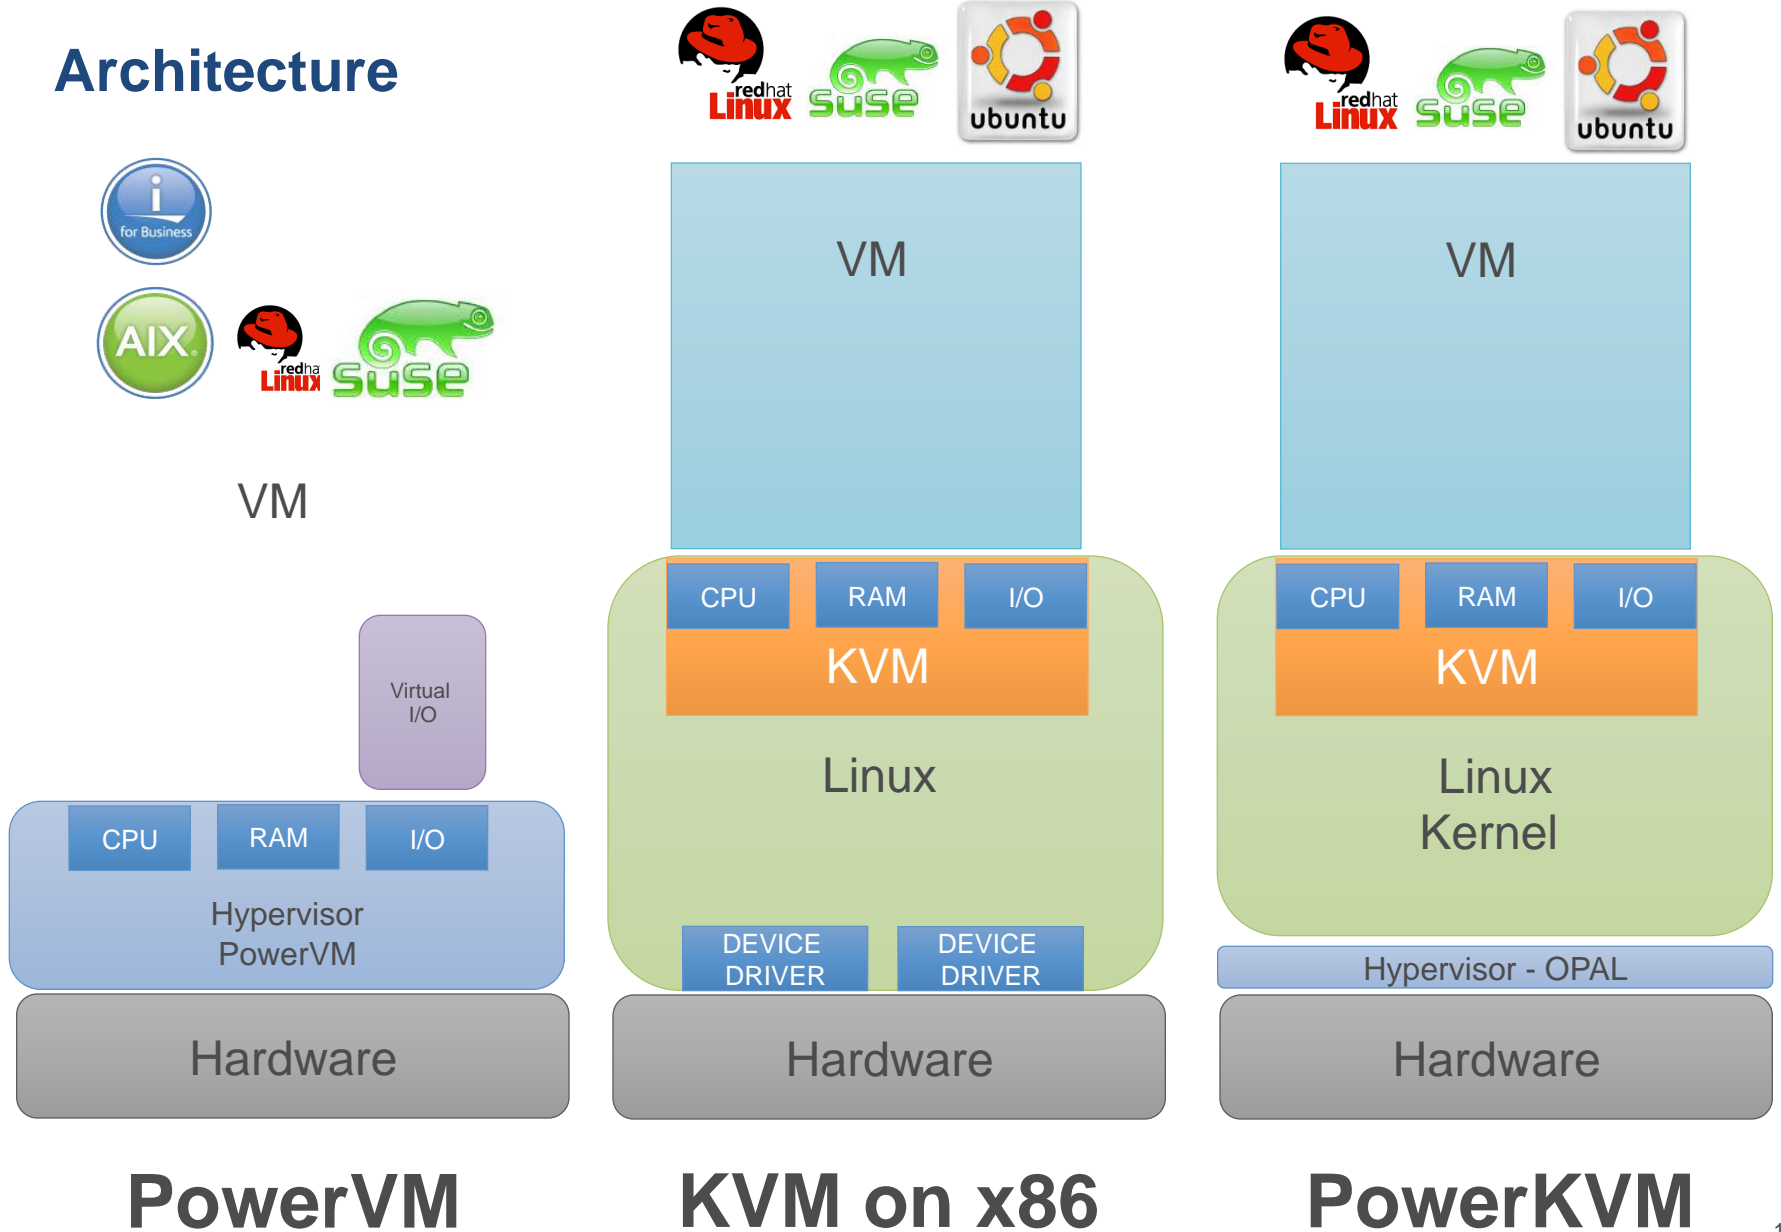
### PowerVM



Initial Offering: 2004

**PowerVM: Provides virtualization of Processors, Memory, Storage, & Networking for AIX, IBM i, and Linux environments on Power Systems.**

# Architecture



**PowerVM**

**KVM on x86**

**PowerKVM**

# PowerKVM v2.1

## Open Virtualization Choice for Linux-only Scale-out Servers



- ✓ Optimize **Linux Workload Consolidation and scale out** of workloads at a lower cost of ownership
- ✓ Maintain flexibility and agility by exploiting **Open Source Community**
- ✓ **Leverage traditional Linux admin skills** on Power Systems to administer virtualization
- ✓ Use **open source tools like OpenStack** to manage virtualization



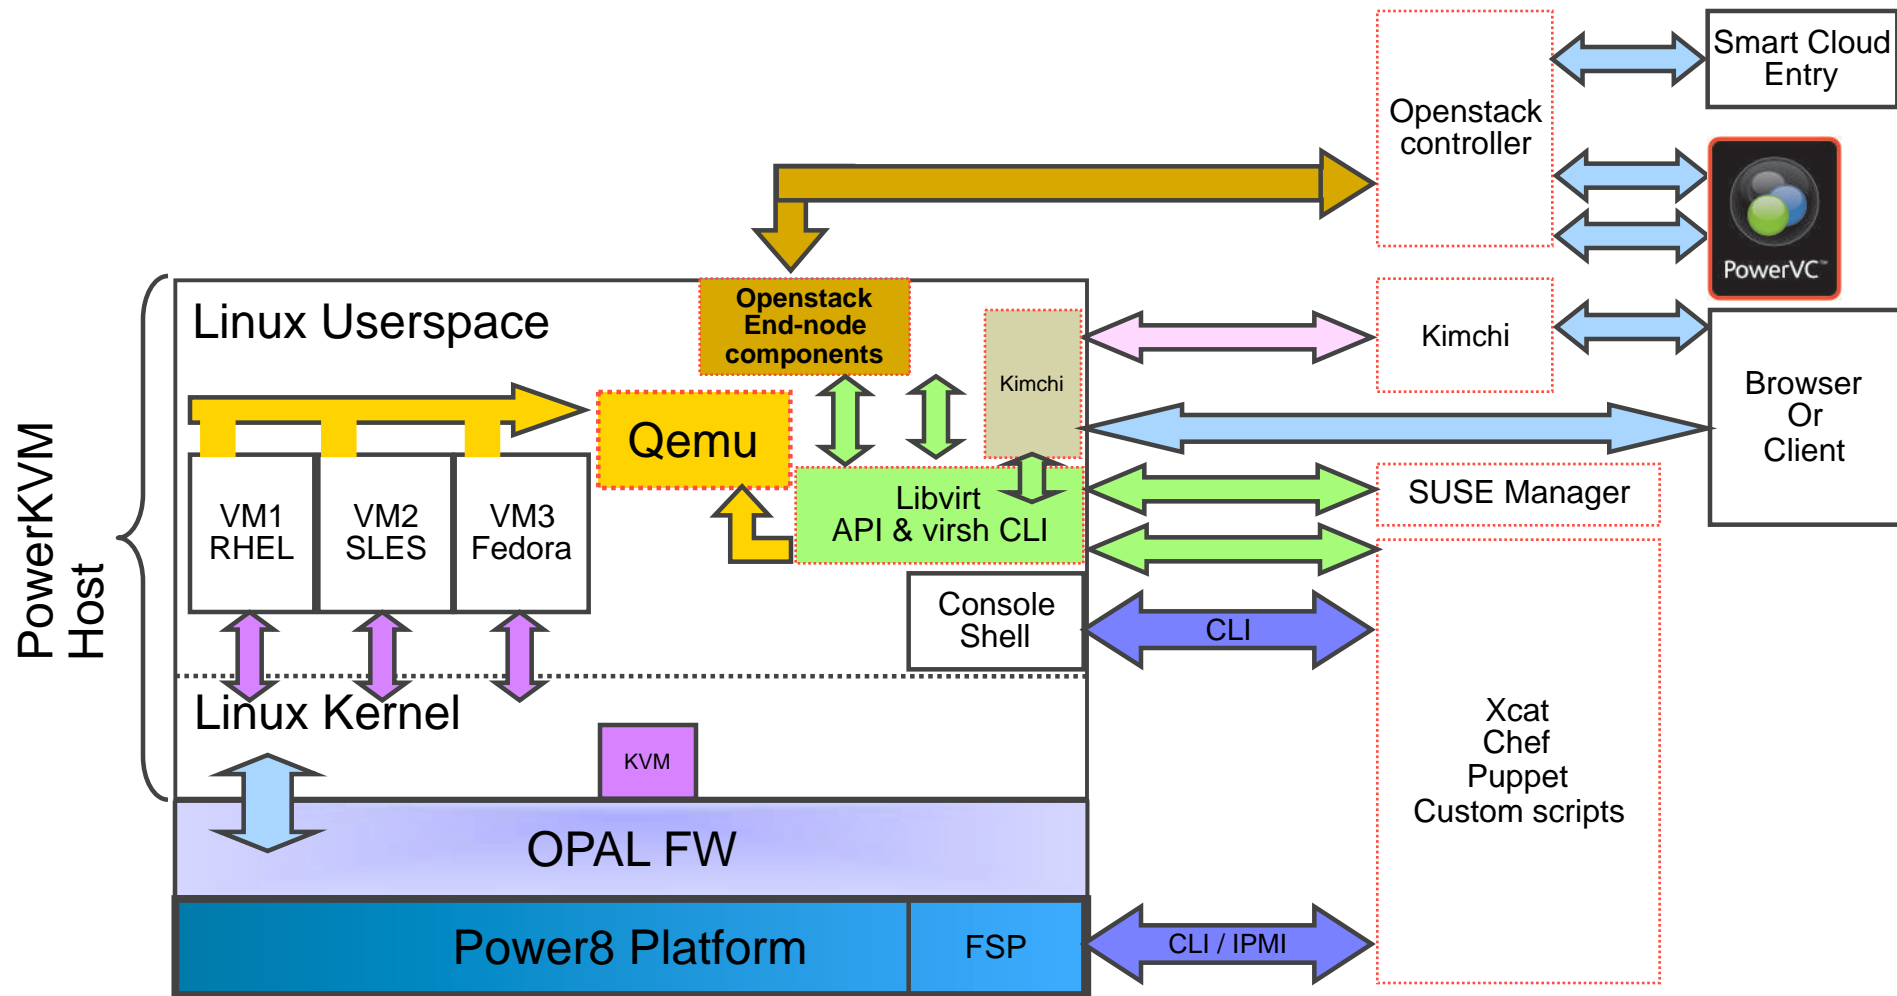
**Announce – 4/28 GA – 6/10**

- ✓ *Reduces IT Infrastructure costs*
- ✓ *Optimize Linux workload consolidation at a lower cost*
- ✓ *Simplify your virtualization management using open source tools*

- Kernel-Based Virtual Machine(KVM) Linux based virtualization For Scale Out POWER8 Linux Servers
- Processor and memory sharing and over commitment enables higher VM and workload consolidation
- Dynamic addition & removal of virtual devices
- Live VM Migration enables higher availability and allows workload balancing
- Exploits P8 Features like Micro-Threading providing greater scheduling granularity vs x86 virtualization
- Exploits performance, scalability and security built into Linux
- Managed by PowerVC and open source tools which provides flexible familiar Linux admin tools
- Supports Redhat, SUSE, Ubuntu Linux Guests




# PowerKVM




# Core Handling


## Full Core



SMT1-8



## SPLIT Core

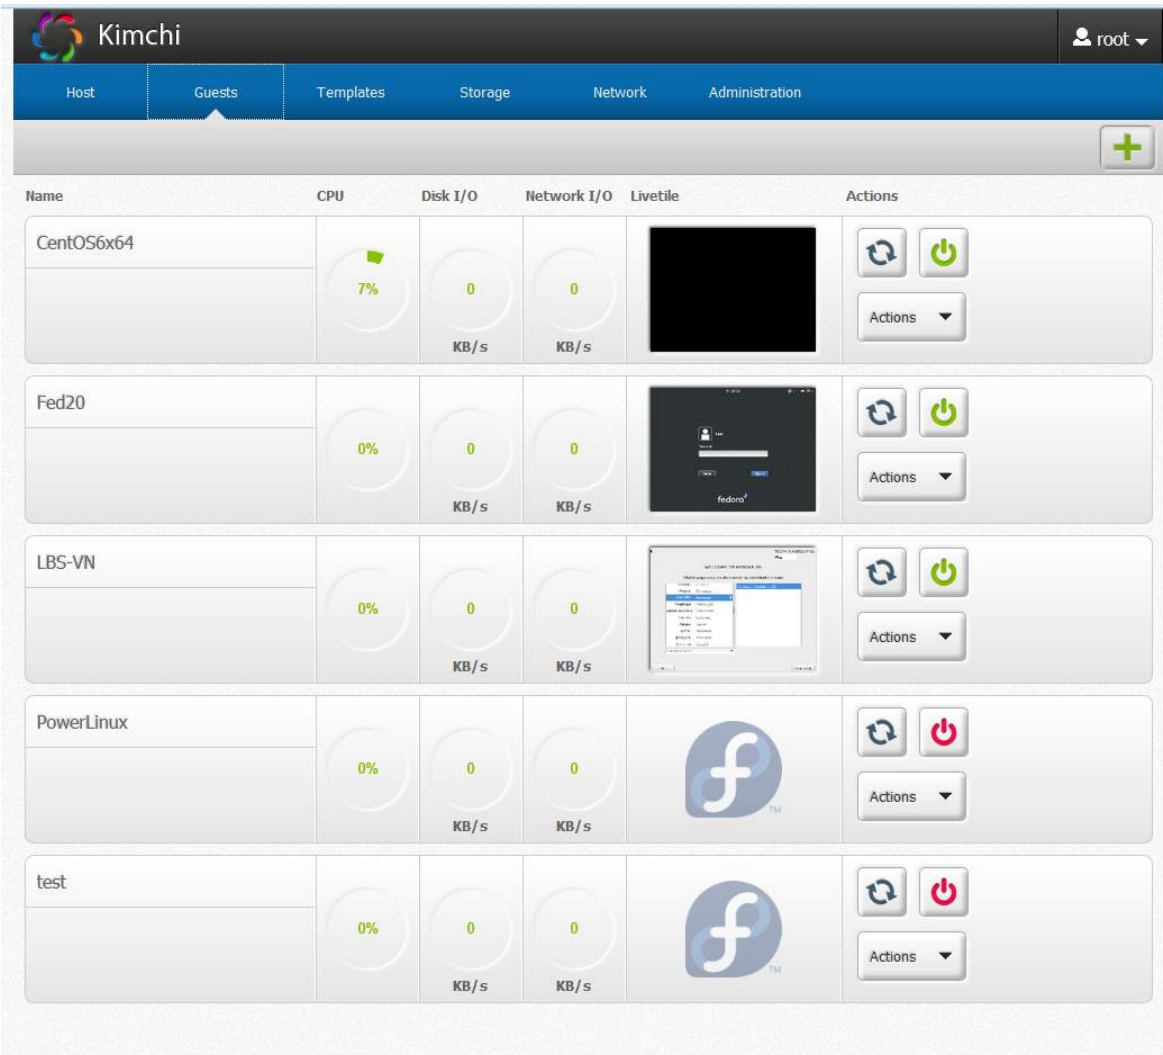

















SMT1-2

## PowerVM vs PowerKVM Comparison

	<i>PowerVM</i>	<i>PowerKVM</i>
<b>GA Availability</b>	<b>2004</b>	<b>Q2 2014</b>
<b>Supported Hardware</b>	<b>All P6, P7, P7+, P8 Systems</b>	<b>PowerLinux P8 Systems S812L &amp; S822L</b>
<b>Supported OS</b>	<b>AIX, IBM i &amp; Linux</b>	<b>Linux</b>
<b>Workload Mobility</b>	<b>Supports AIX, IBM i &amp; Linux</b>	<b>Linux</b>
<b>Basic Virtualization Management</b>	<b>IVM / HMC / FSM</b>	<b>Virtman/libvirt</b>
<b>Advanced Virtualization Management</b>	<b>PowerVC/VMControl</b>	<b>PowerVC, Vanilla OpenStack</b>
<b>Admin Type</b>	<b>Power Centric</b>	<b>Linux/x86 Centric</b>
<b>Established Security Track Record on Power</b>	<b>Yes</b>	<b>No</b>
<b>Open Source Hypervisor</b>	<b>No</b>	<b>Yes</b>

# Kimchi Local Administration



Name	CPU	Disk I/O	Network I/O	Livetile	Actions
CentOS6x64	7%	0 KB/s	0 KB/s		  Actions
Fed20	0%	0 KB/s	0 KB/s		  Actions
LBS-VN	0%	0 KB/s	0 KB/s		  Actions
PowerLinux	0%	0 KB/s	0 KB/s		  Actions
test	0%	0 KB/s	0 KB/s		  Actions

Provides simple graphical web interface to initially configure the PowerKVM Host and to manage basic virtualization for a small configuration.

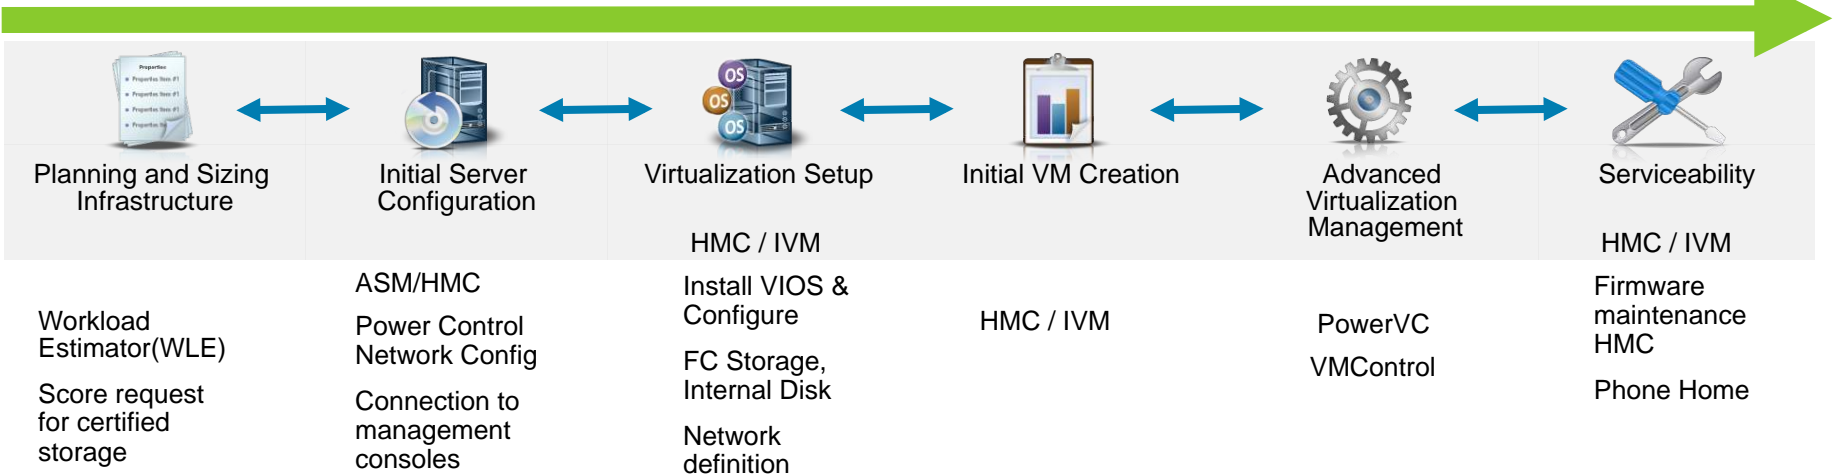
Included in PowerKVM distribution.

Function includes

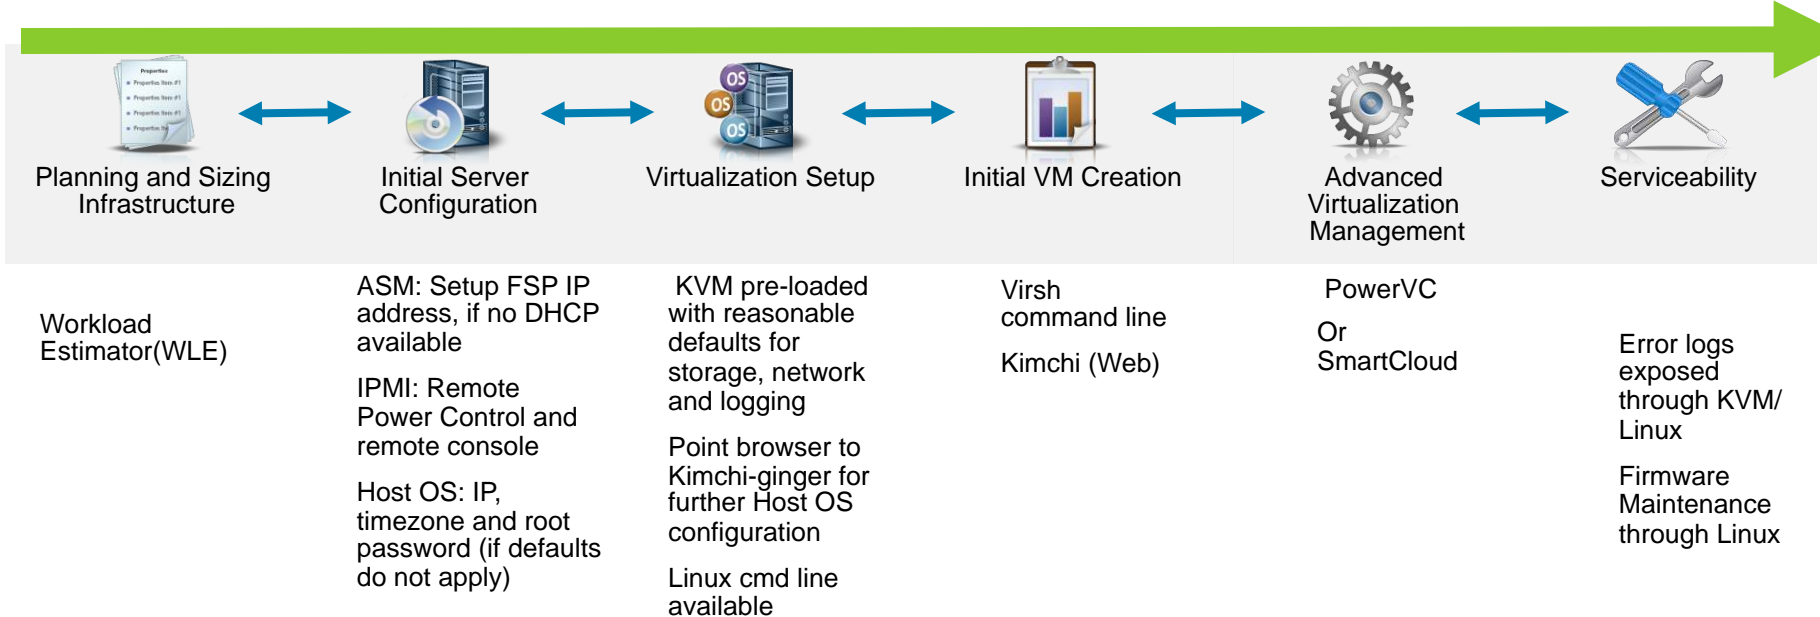
- ✓ Initial host setup
- ✓ Firmware update
- ✓ Backup of configuration
- ✓ Simple VM setup
- ✓ Start and stop of VMs
- ✓ Host monitoring
- ✓ Use of Templates
- ✓ View VM guest console

# PowerVM VS KVM Out of Box Experience

## PowerVM



## PowerKVM



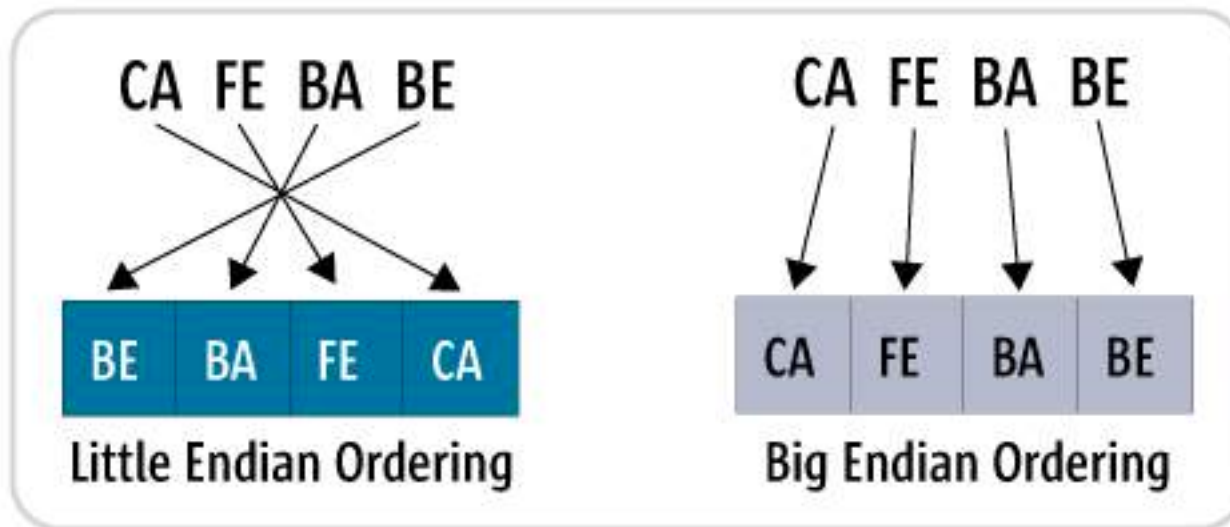
## PowerKVM Offering Structure

- PowerKVM has a single Standard Edition = PID 5765-KVM
- Socket Based Pricing with Subscription model for the PowerKVM host
  - Pricing is socket based and requires buying a socket pair subscription 8x5 or 24x7 offering
- POWER8 Scale-out Linux Only Servers have a choice of either PowerKVM or PowerVM
  - Only one virtualization option can be chosen
  - Either PowerVM or PowerKVM can run on the system but not both
- When PowerKVM is ordered the feature code drives proper firmware defaults
  - PowerKVM is preloaded on internal disk with option to disable preload
  - Guest Linux ISO can be preloaded for convenience to build guest VMs



## Endianness

- Endianness is how a system handles data



- x86 is Little Endian
- PowerVM is Big Endian
  - AIX, IBM i, RHEL, SUSE on top of PowerVM must be Big-Endian
- POWER is Bi-endian !!!

# Endianness

- Endianness is how a system handles data

```
#include <stdio.h>
int main(void) {
    int val;
    unsigned char *ptr;

    ptr = (char*) &val; /* pointer ~ptr~ points to ~val~ */
    val = 0x89ABCDEF; /* four bytes constant */
    printf("%X.%X.%X.%X\n", ptr[0], ptr[1], ptr[2], ptr[3]);
    exit(0);
}
```

On an x86 system, the result is:

```
EF.CD.AB.89
```

On a POWER system, the result is:

```
89.AB.CD.EF
```

- Linux Apps with this kind of programming logic had to have its code changed to run on Power properly



## PowerKVM Positioning

- First release available in 2014
- Focus: **New Linux workloads** for Power Systems
- **Seamless transition for existing Linux admins** to adopt Power Linux Virtualization without any training
- **No HMC or other traditional IBM consoles**
  - Normal Linux management and OpenStack options
- **PowerKVM only supports Linux guest VMs**
- **Cloud potential:** Have many more small VMs than traditional Power Virtualization
- **POWER8 PowerLinux hardware only**
- **Live Workload mobility** support between PowerKVM servers
- **Open Source Hypervisor:** Hardware is abstracted by firmware
- **Managed by OpenStack(PowerVC) or by off the shelf OpenStack or local Linux Tools**

# More Information on PowerKVM

[ibm.com/systems/power/software/linux/powerlinux/powerkvm](http://ibm.com/systems/power/software/linux/powerlinux/powerkvm)

IBM Systems > Power Systems > Software > Linux > PowerLinux solutions >

## IBM PowerKVM – open virtualization

**PowerKVM – the open virtualization choice for Power scale-out Linux Systems**

Overview


Features & benefits

### Highlights

- Open source server virtualization for IBM scale-out Linux servers
- Optimize Linux workload consolidation at a lower cost of ownership
- Exploit the advantage of performance, scalability and security built into Linux and the Kernel-Based Virtual Machine (KVM) hypervisor
- Avoids high cost proprietary x86 virtualization
- Managed just like any other KVM host – OpenStack, libvirt and open Linux tools help you avoid vendor lock-in
- KVM enables single cross platform virtualization which simplifies management
- Maintain flexibility, extensibility and agility by exploiting open source solutions
- Leverage traditional Linux admin skills on Power Systems
- Exploit new advanced POWER8 processor features
- Eliminate scheduled downtime by using live VM migration




IBM® PowerKVM™ provides an open virtualization choice for IBM scale-out Linux systems based on the POWER8™ technology. This solution includes the Linux open source technology of KVM virtualization, and is designed to complement the performance, scalability and security qualities of Linux. This provides an open extendable solution for running VMs on Linux scale-out servers that enables cloud deployments, scale-out processing, and big data solutions reducing complexity and cost.

### Learn more



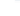


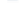
-  [View IBM PowerKVM Open Virtualization data sheet \(763KB\)](#)
- [→ IBM PowerVM for IBM PowerLinux](#)
- [↔ Get Adobe Reader](#)

### Contact IBM



-  [Chat now](#)
-  [Email IBM](#)
- [→ Find a Business Partner](#)
-  [Call IBM: 1-866-883-8901](#)  
Priority code: 101AR13W

### Browse Power Systems

-  [Hardware](#)
-  [Operating systems](#)
-  [System software](#)
-  [Solutions](#)
-  [Migrate to Power](#)
-  [Advantages](#)
- [→ Community](#)
- [→ Success stories](#)
- [→ News](#)
- [→ Support & services](#)
- [→ Resources](#)
- [→ Education](#)

Power Systems



Open Innovation to Put Data to Work

IBM