

Les nouvelles solutions Power Systems

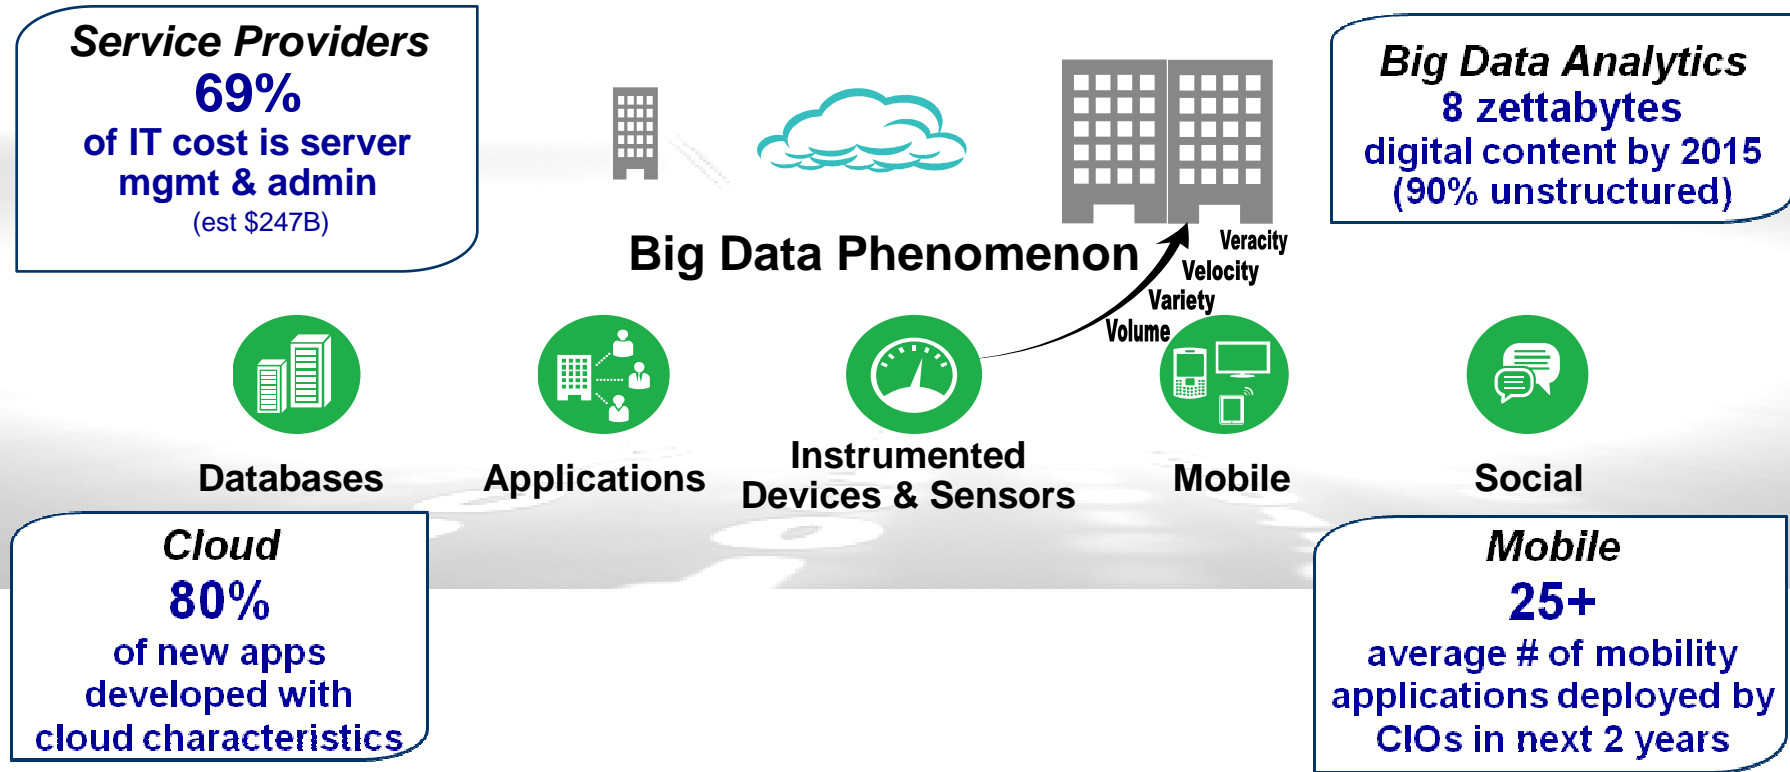
GID SAS

Lyon, 22 Mai 2014

Didier FERAUD
CTS Power STG
didier_feraud@fr.ibm.com



Business needs are driving infrastructure transformation



CIOs are turning to innovative technologies to address these challenges

Source: "Hunting and Harvesting in a Digital World - Insights From the 2013 Gartner CIO Agenda Report". Gartner Executive Program Survey of More Than 2,000 CIOs Shows Digital Technologies Are Top Priorities in 2013. <http://www.gartner.com/technology/cio/cioagenda.jsp>

First generation of systems built with open innovation to put data to work

Optimize Data and Analytics

- **Solutions & operating systems** optimized for new POWER8 big data & analytics innovations
- **Chip designs** for Java apps, big data/analytics
- **I/O** - PCI-3 for more disk and SSD for faster access and greater efficiency



Enhance Cloud Efficiency

- **Solutions & operating systems** optimized for POWER8 Java & virtualization innovations
- **Virtualization and cloud management** built on OpenStack, KVM, PowerVM
- **Chip designs** for cloud and MSP environments, i.e. KVM micro larger number of partitions



Enable Open Innovation

- Accelerate key emerging workloads with open source community and **OpenPOWER Foundation**
- Engaging new, leading **ecosystem partners to deliver new innovations**



ISVs are critical to the IBM Power Ecosystem


Data & Analytics



Cloud



Mobile



Focus Areas in 2014

- Leverage IBM Ecosystem - Differentiate
- Gain Platform share in Big Data, Mobile and Cloud
- Build Regional ISV Ecosystem
- Develop Open Software Linux community for Power
- Create incremental value around existing Power base



Available on All Linux Servers



POWER8 Scale Out Systems

Power S812L / S822L



Scale-out Linux server based on open technology

Power S814



Scale-out technology for mid-sized business solutions

POWER8 architecture

Power S822



Scale-out application server for secure infrastructure built on open technology

Power S824



Scale-out technology server for faster insights from data

Scale Out / Scale In definitions ...

- To *scale horizontally* (or *scale out*) means to add more nodes to a system, such as adding a new computer to a distributed software application. An example might be scaling out from one Web server system to three.
- To *scale vertically* (or *scale up*) means to add resources to a single node in a system, typically involving the addition of CPUs or memory to a single computer. Such vertical scaling of existing systems also enables them to use [virtualization](#) technology more effectively, as it provides more resources for the hosted set of [operating system](#) and [application](#) modules to share.

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**

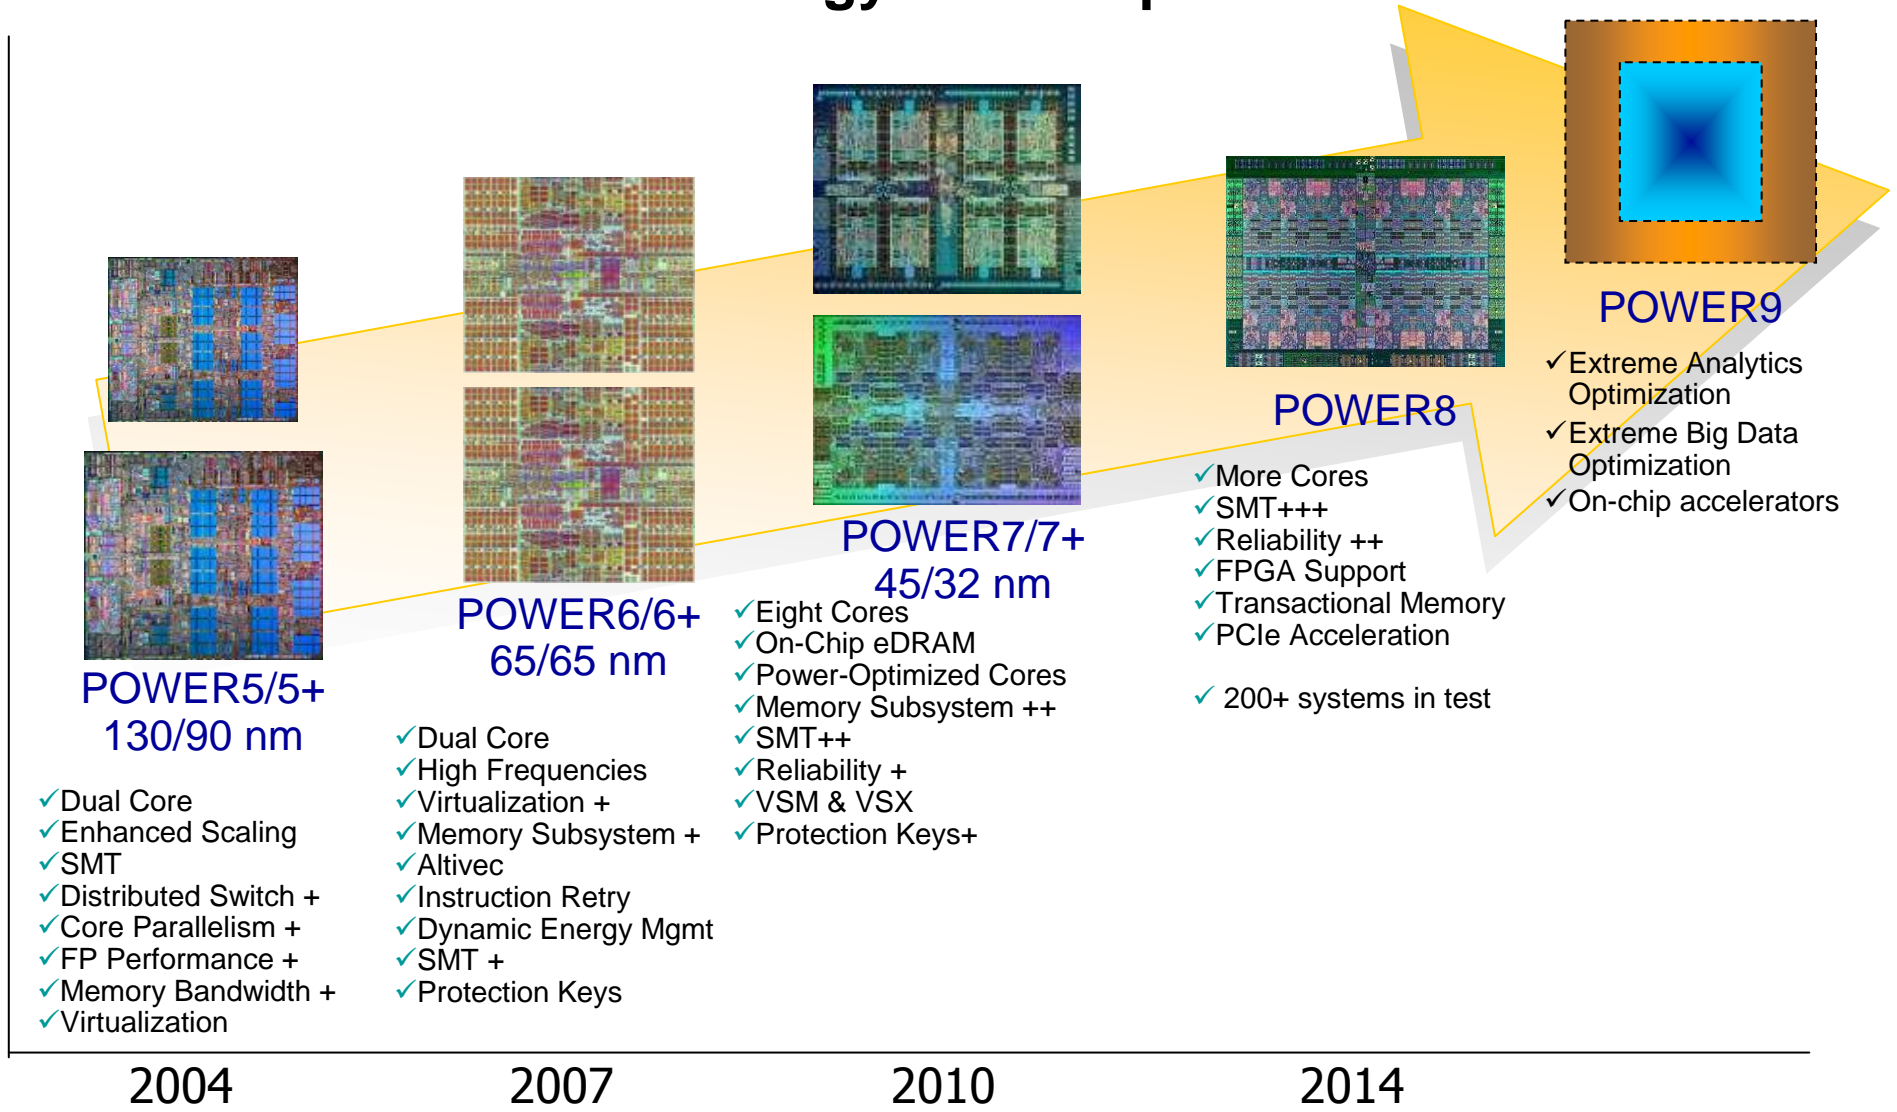
Marketing name = “socket” then “number EIA/U”

For example: **S 8 2 4**





S for Scale Out


Power Processor Technology Roadmap



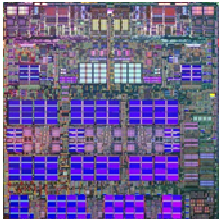
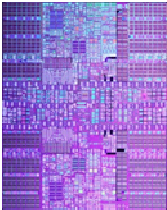
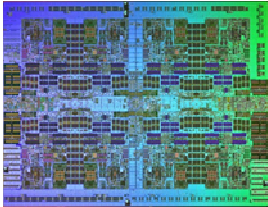
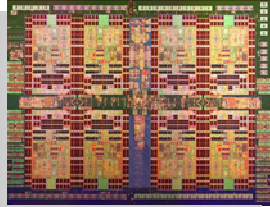
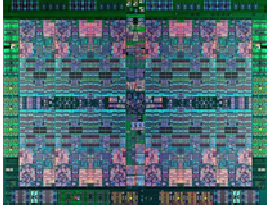
POWER8 Processor



POWER8



- ➔ ***Significant Performance at Thread, Core, and System***
- ➔ ***Optimization for VM Density & Efficiency***
- ➔ ***Strong Enablement of Autonomic System Optimization***
- ➔ ***Excellent Big Data Analytics Capability***


	POWER5 2004	POWER6 2007	POWER7 2010	POWER7+ 2012	POWER8 2014
					
Technology	130nm SOI	65nm SOI	45nm SOI eDRAM	32nm SOI eDRAM	22nm SOI eDRAM
Compute					
Cores	2	2	8	8	12
Threads	SMT2	SMT2	SMT4	SMT4	SMT8
Caching					
On-chip	1.9MB	8MB	2 + 32MB	2 + 80MB	6 + 96MB
Off-chip	36MB	32MB	None	None	128MB
Bandwidth					
Sust. Mem.	15GB/s	30GB/s	100GB/s	100GB/s	230GB/s
Peak I/O	6GB/s	20GB/s	40GB/s	40GB/s	96GB/s

POWER8 Highlights

.... PERFORMANCE

The Register: “it most certainly does belong in a **badass** server, and Power8 is by far one of the most elegant chips that Big Blue has ever created”

- **Huge performance improvement**
 - 1.5X to 1.7X thread-level improvement*
 - 2X core improvement*
 - 3X socket improvement*
- **SMP scaling**
 - 16 sockets, 192 cores
 - Lower latency, high speed
- **Coherent Accelerators (CAPI)** – *memory-space addressable*



POWER8 Highlights

... *MEMORY and BANDWIDTH*


Linley Group microprocessor report: “The Power8 specs are mind boggling... IBM’s newest server processor will smash existing performance records, particularly for memory-intensive applications”

- **Caching Structure**

- L1 to L4 cache with NUCA Non-Uniform Cache Architecture
- 4 TB/s L2 bandwidth* per chip (4GHz 12core)
- 3 TB/s L3 bandwidth* per chip (4GHz 12core)

- **Memory Subsystem**

- 196-230 GB/s sustained* bandwidth per chip
- 410 GB/s bandwidth* at DRAM level per chip
- 1 TB memory per socket (e.g. 4sockets = 4 TB)
- Transactional Memory




POWER8 Highlights

... *I/O, BANDWIDTH, VALUE*

TechInvestor/VentureBeat: “IBM preps its massive 12-headed Power 8 chip”

- **Balance I/O capability**
 - PCIe Gen3 **on-chip** I/O connectivity and protocol (Low latency)
 - I/O bandwidth 96 GB/s per socket
 - Flexible chip interface
- **Energy 3X capacity per watt improvement***
- **Improved RAS**
- **P6/P7/P8 modes and Live Partition Mobility**
- **Workload density / LPAR density**



POWER8 Processor

Cores

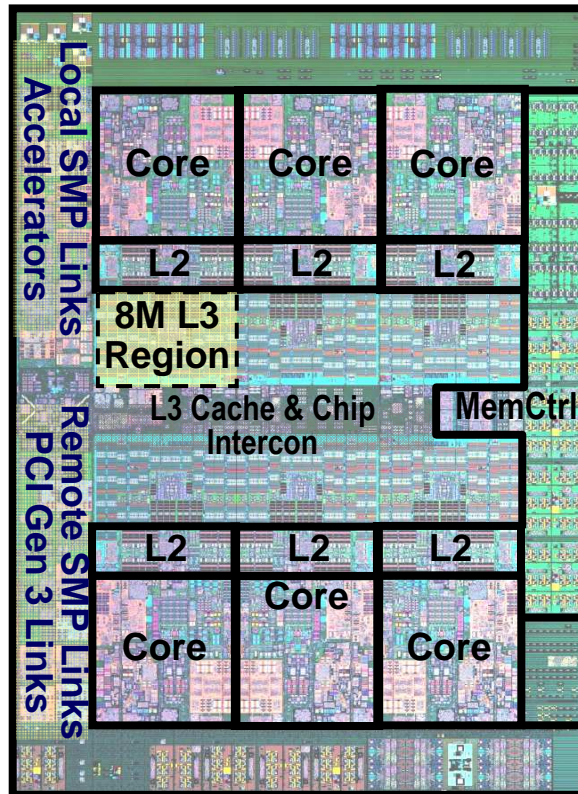
- 6 cores (SMT8)
- 8 dispatch, 10 issue, 16 exec pipe
- 2X internal data flows/queues
- Enhanced prefetching
- 64K data cache, 32K instruction cache

Accelerators

- Crypto & memory expansion
- Transactional Memory
- VMM assist
- Data Move / VM Mobility

Technology


22nm SOI, eDRAM, 15 ML 650mm²



Energy Management


- On-chip Power Management Micro-controller
- Integrated Per-core VRM
- Critical Path Monitors

POWER8 DCM



Caches

- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)



Memory

- Up to 230 GB/s sustained bandwidth
- Up to 1TB per socket


Bus Interfaces

- Durable open memory attach interface
- Integrated PCIe Gen3
- SMP Interconnect
- CAPI (Coherent Accelerator Processor Interface)




POWER8 Multi-threading Options

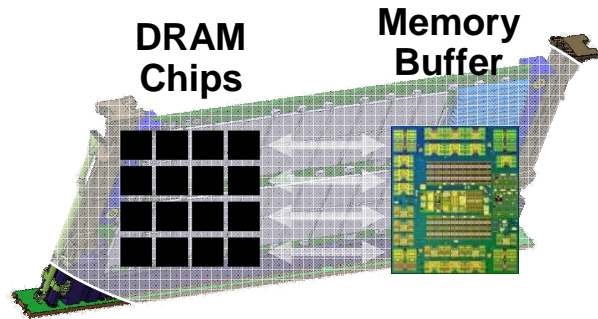
- **SMT1:** Largest unit of execution work
- **SMT2:** Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT4:** Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT8:** Smallest unit of work, but provides the maximum amount of execution work per cycle
- Can **dynamical** shift between modes as required: SMT1 / SMT2 / SMT4 / SMT8
- Mixed SMT modes supported within same LPAR
- ☞ Requires use of “Resource Groups”



POWER8 I/O Bandwidth



Memory Buffer Chip *with 16MB Cache !!!*



Intelligence Moved into Memory

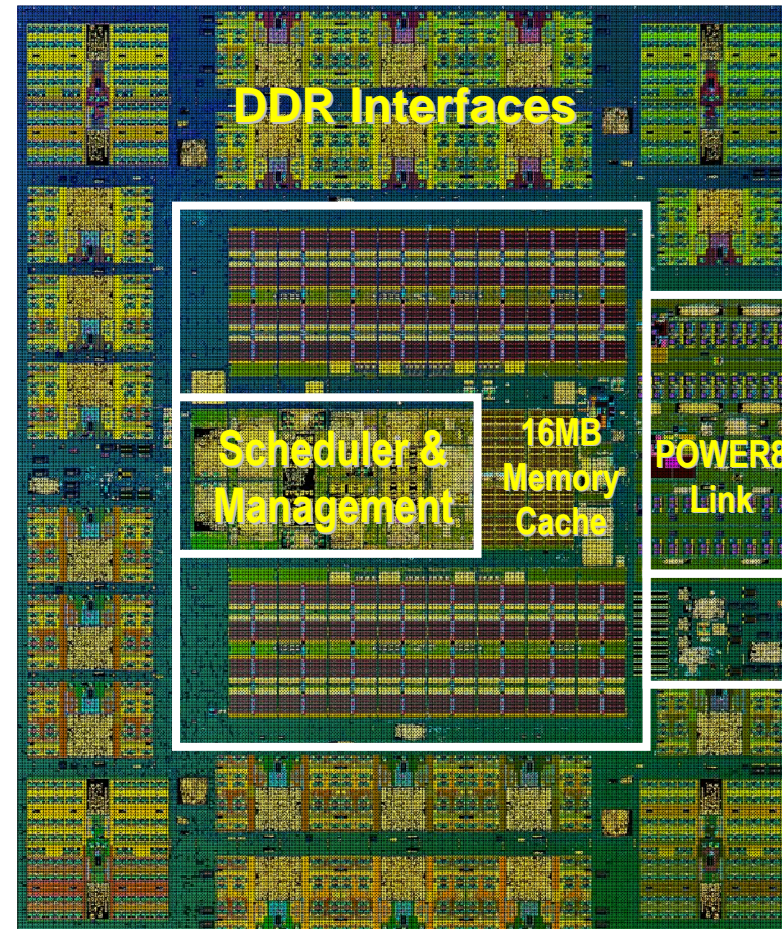
- Scheduling logic, caching structures
- Energy Mgmt, RAS decision point
 - Formerly on Processor
 - Moved to Memory Buffer

Processor Interface

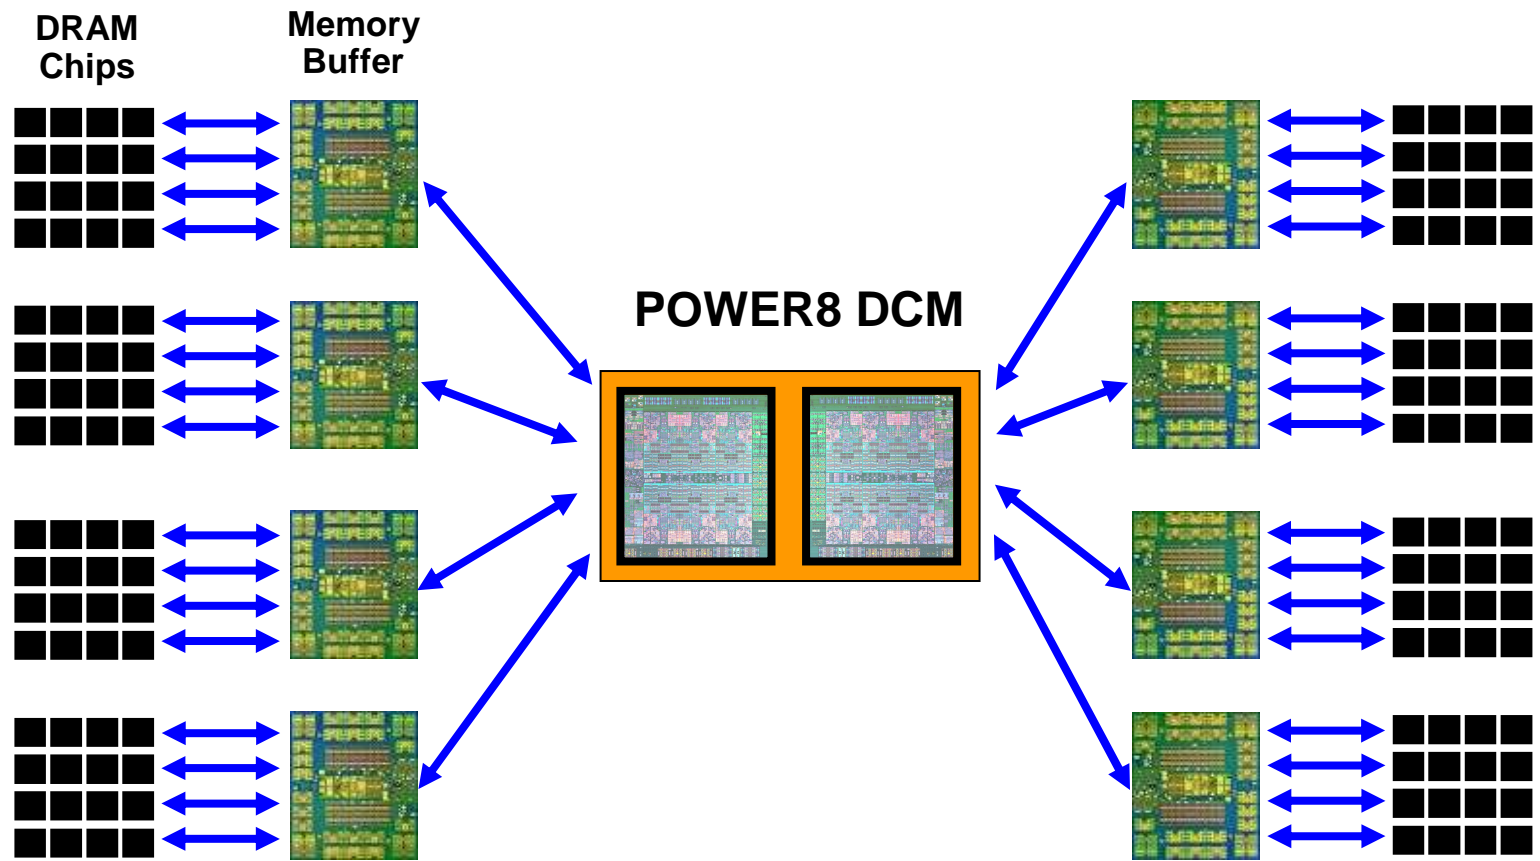
- 9.6 GB/s high speed interface
- More robust RAS
- “On-the-fly” lane isolation/repair
- Extensible for innovation build-out

Performance Value

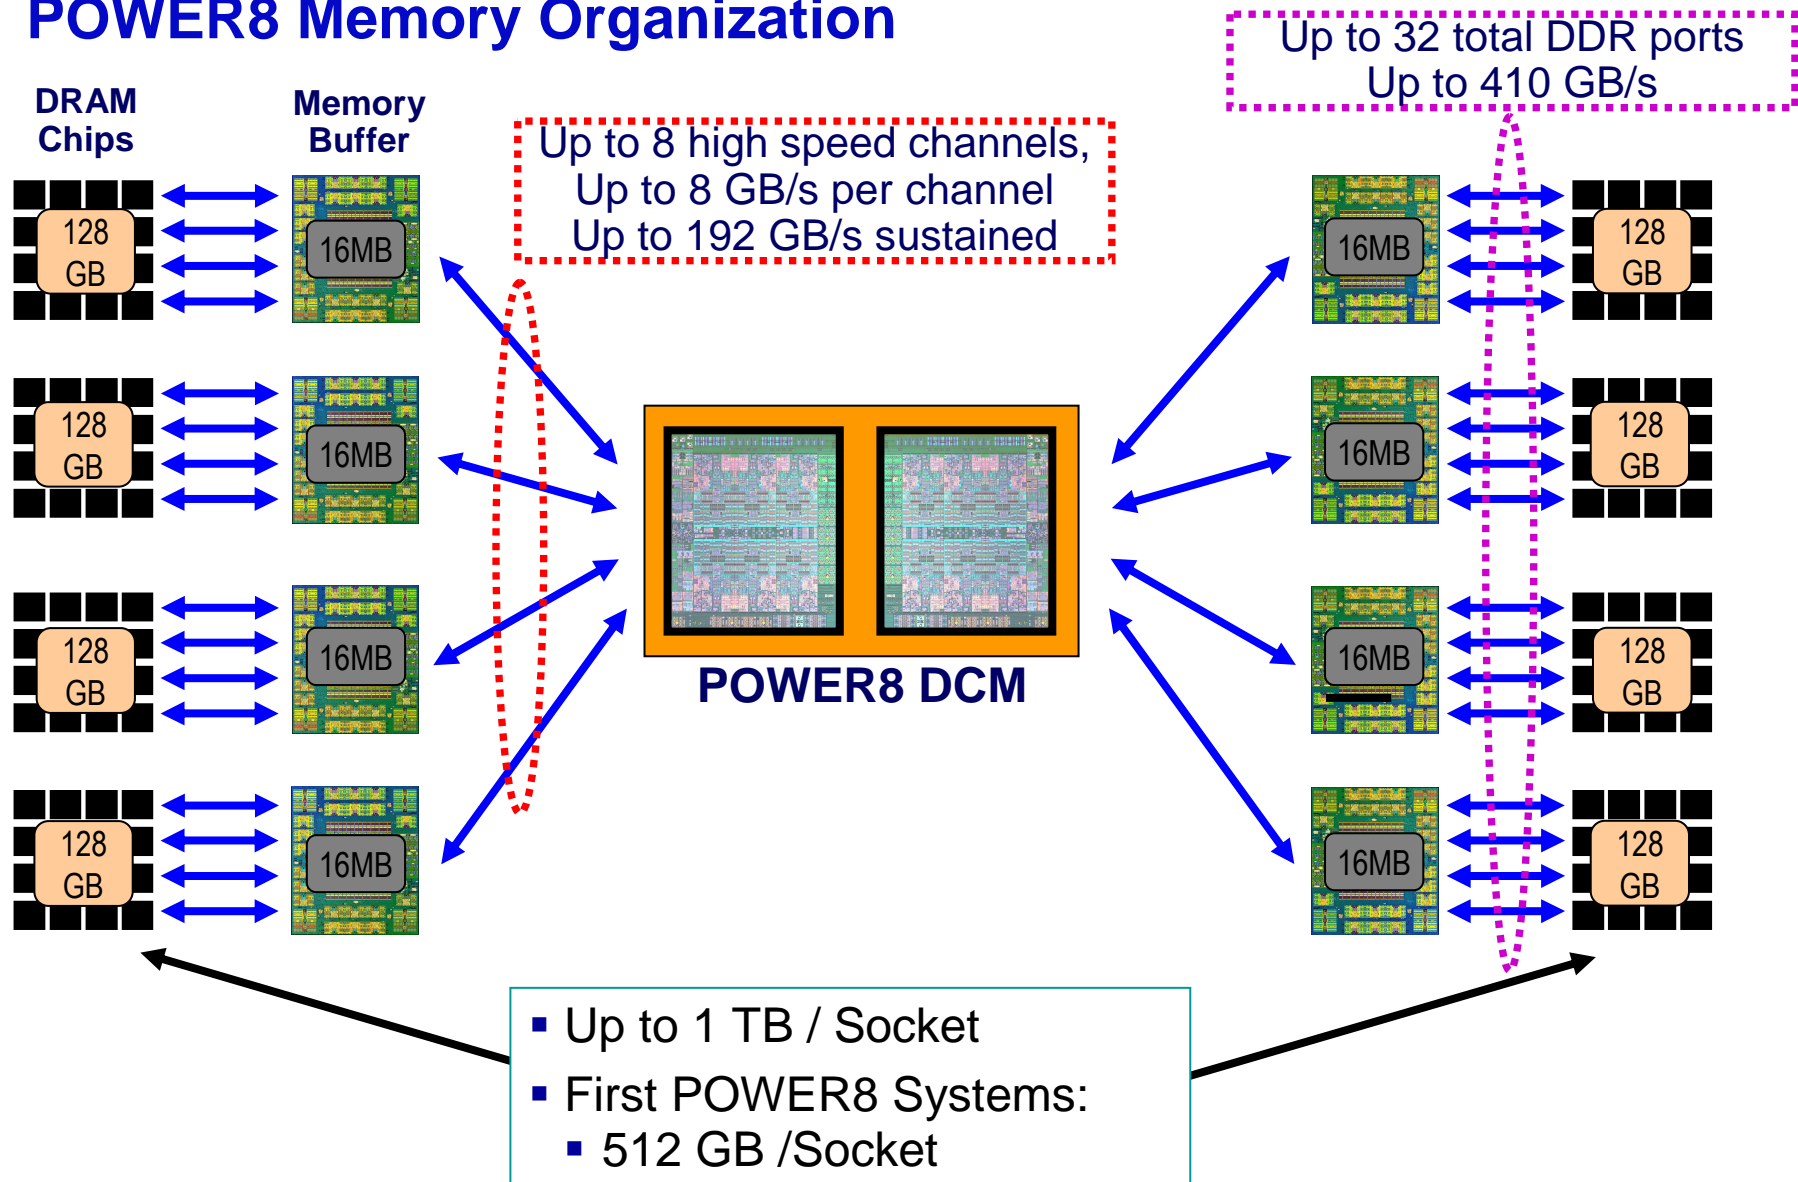
- End-to-end fastpath and data retry (latency)
- Cache → latency/bandwidth, partial updates
- Cache → write scheduling, prefetch, energy
- 22nm SOI for optimal performance / energy
- 15 metal levels (latency, bandwidth)



POWER8 Memory Organization (Max Config shown)



POWER8 Memory Organization



Transactional Memory

Technique that allows a **group** of instructions including **updates to memory image** to execute **speculatively and atomically**. This group of instructions is called a transaction

Value

- Reducing programming development
- Reducing customer cost (higher SLA / fewer images and higher scalability)
- Improving performance of legacy software with large sequential components

Power8 Support


- New instructions mark beginning and end of transaction
- Hardware ensures region is performed atomically using speculation
- Speculation recovery performed in hardware, both registers and memory
- “Flattened” Nesting
- Hardware tracks nesting of transactions / Treats them all as a single large transaction

How to determine potential benefits?


- If the application is highly threaded or manages many resources that might have light weight locking requirements, TM is likely to help
- If read-only transactions typically do not reference the same data as concurrent transactions that write data, TM may help
- If the application is coarse-grained locking and does not scale well, TM may help
- If the application has fine-grained locking, TM may not help
- If the transactions reference a large amount of data, TM is unlikely to help

POWER8 Integrated PCI Gen 3


POWER7



GX Bus



PCIe Gen2



PCI Devices


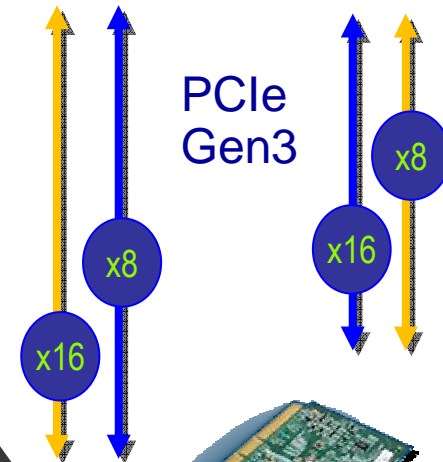
Native PCIe Gen 3 Support

- 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



PCI Devices

CAPI

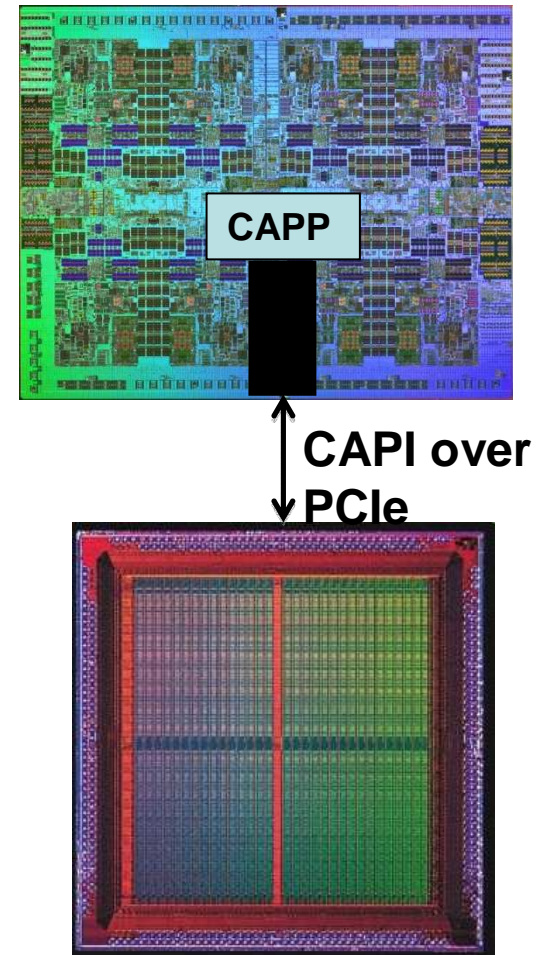
Statement of Direction



CAPI Enables off-chip accelerator to act as an on-chip processor

- **How do we do it ?**
 - Unit on CPU that extends on chip coherence to attached accelerator
- **Efficient model for accelerator**
 - Eliminates the data copy required for other accelerators like GPUs
 - Low latency accesses possible since host CPU not involved
- **Flexible**
 - Can choose from a large mix of FPGA devices
 - Enables correct mix of CPU and FPGA performance
 - Any PCIe device can be enabled

CAPP and PCIe unit not to scale




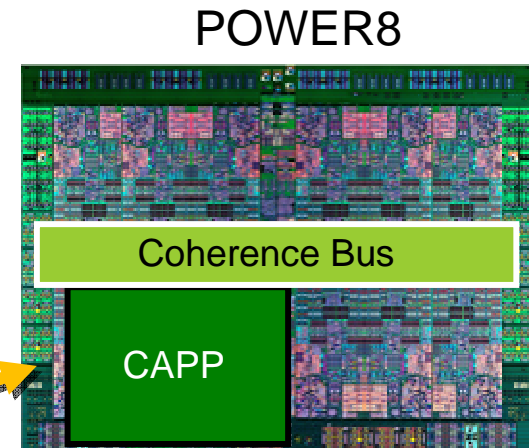
POWER8 CAPI (Coherent Accelerator Processor Interface)

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



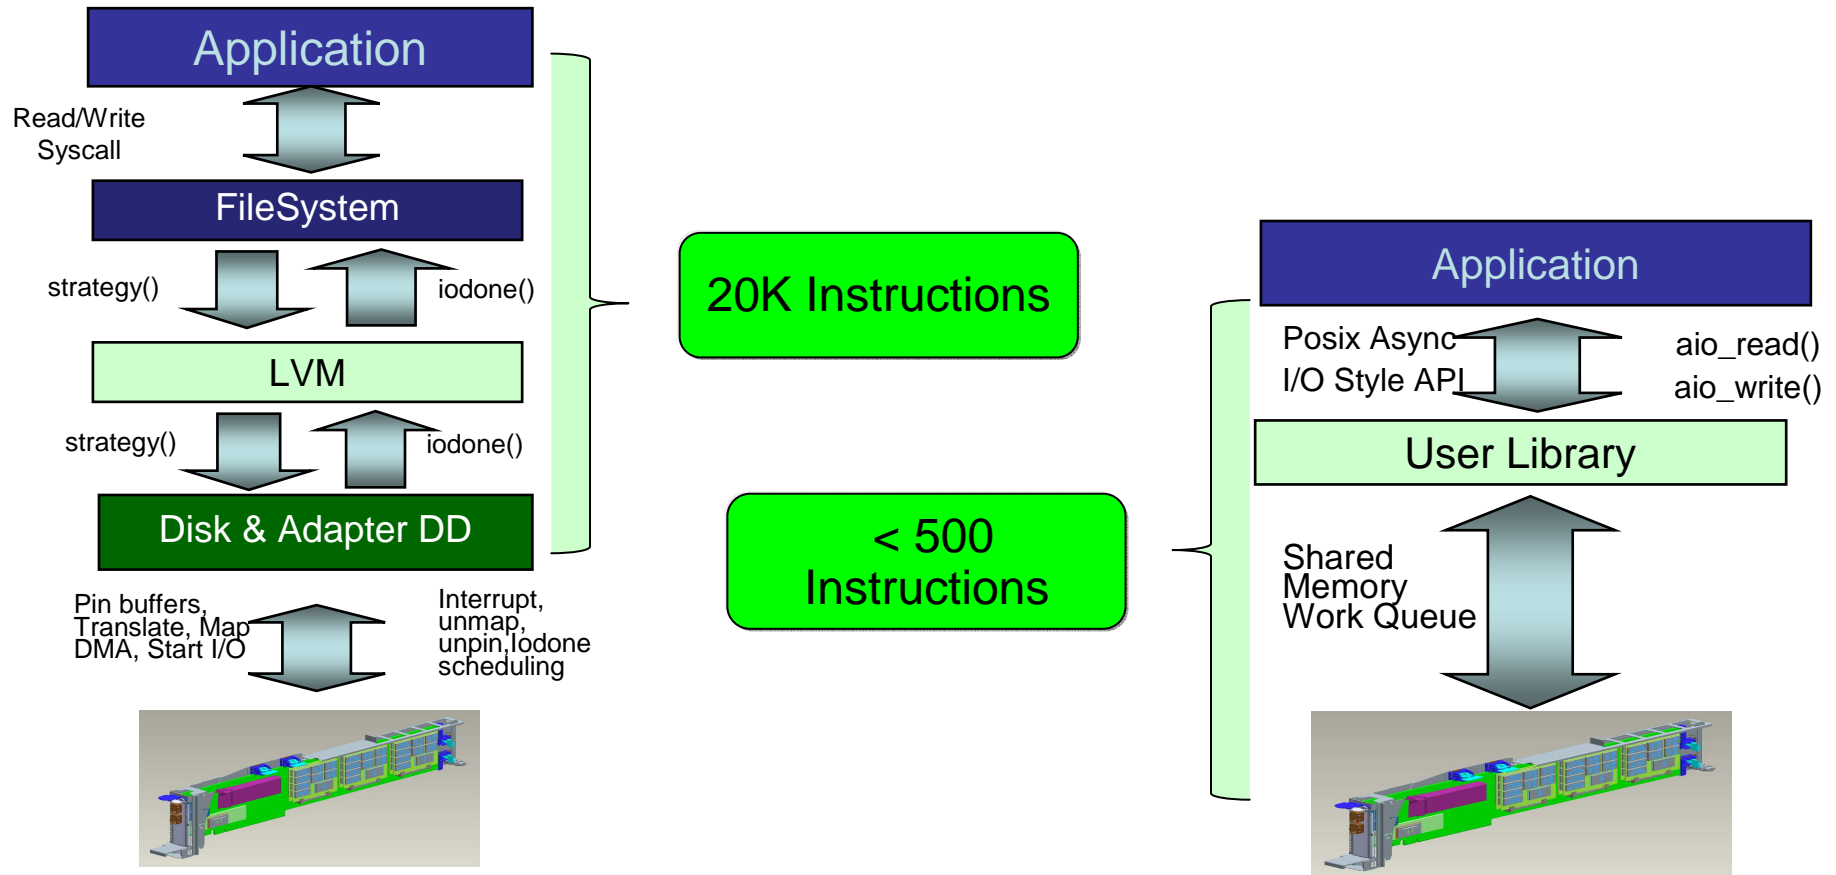
Processor Service Layer (PSL)

- Present robust, durable interfaces to applications
- Offload complexity / content from CAPP

Customizable Hardware Application Accelerator

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

Possible Example: CAPI Attached Flash Optimization



- **Attach flash memory 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**

POWER8 CAPI


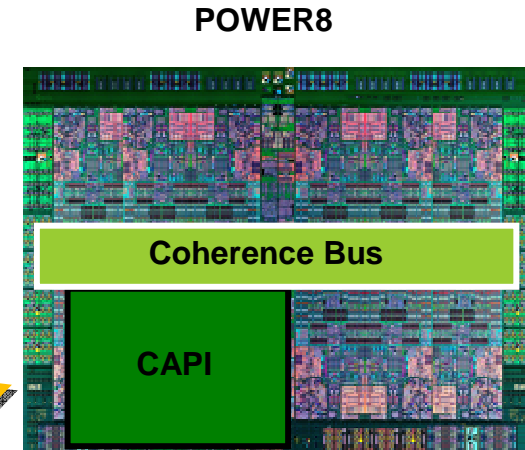
Coherent Accelerator Processor Interface (CAPI)

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



PCIe Gen 3

Transport for encapsulated messages

Processor Service Layer (PSL)

- Present robust, durable interfaces to applications
- Offload complexity / content from CAPI


Customizable Hardware Application Accelerator

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

CAPI enables innovation from the OpenPOWER Foundation

POWER8

April 14



Announce / Availability Plans

	Annc	eConfig	GA
Power S814	28 Apr	28 Apr	10 Jun
Power S822	28 Apr	28 Apr	10 Jun
Power S824	28 Apr	28 Apr	10 Jun
Power S822L	28 Apr	28 Apr	10 Jun
Power S812L	28 Apr	27 May	29 Aug

Power System Roadmap

2013

2014

POWER8

Enhanced HMC Management

KVM

PowerVC

PowerVP



4U Systems: S814 & S824

Power Linux Engines




2U System: S822

Enterprise Pools

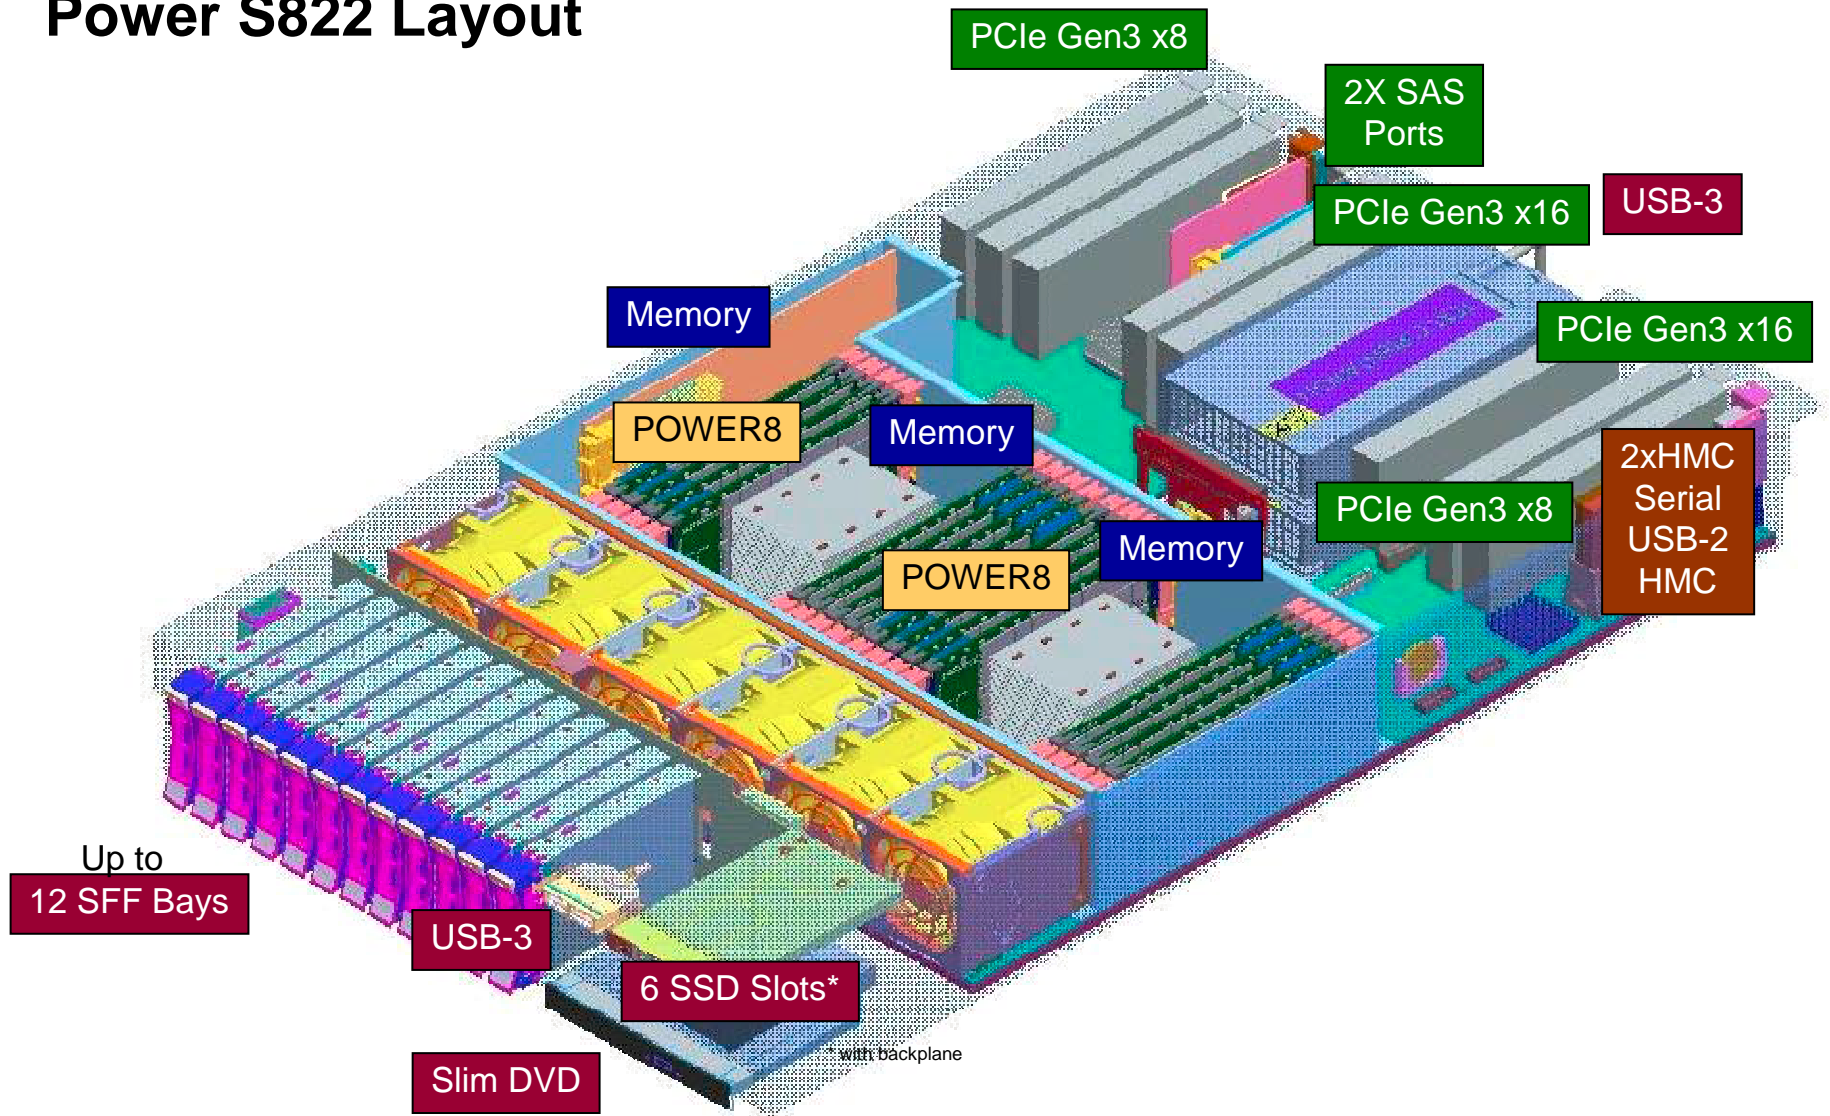


2U Linux: S822L

2U Servers



Power S822 Layout



Power S822 Scale-Out System



Single Socket populated

- Cores: 6 (3.8 GHz) or 10 (3.4 GHz)
- Memory: Up to 512 GB
- PCIe Slots: 6 PCIe Gen3 LP (Hotplug)

Both Sockets populated


- Cores: 12 (3.8 GHz) or 20 (3.4 GHz)
- Memory: Up to 1 TB
- PCIe Slots: 9 PCIe Gen3 LP (Hotplug)

Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4), Serial (2), HMC (2)

Internal Storage

- DVD
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 8 SFF Bays & 6 1.8" SSD Bays with Easy Tier with 7GB write cache



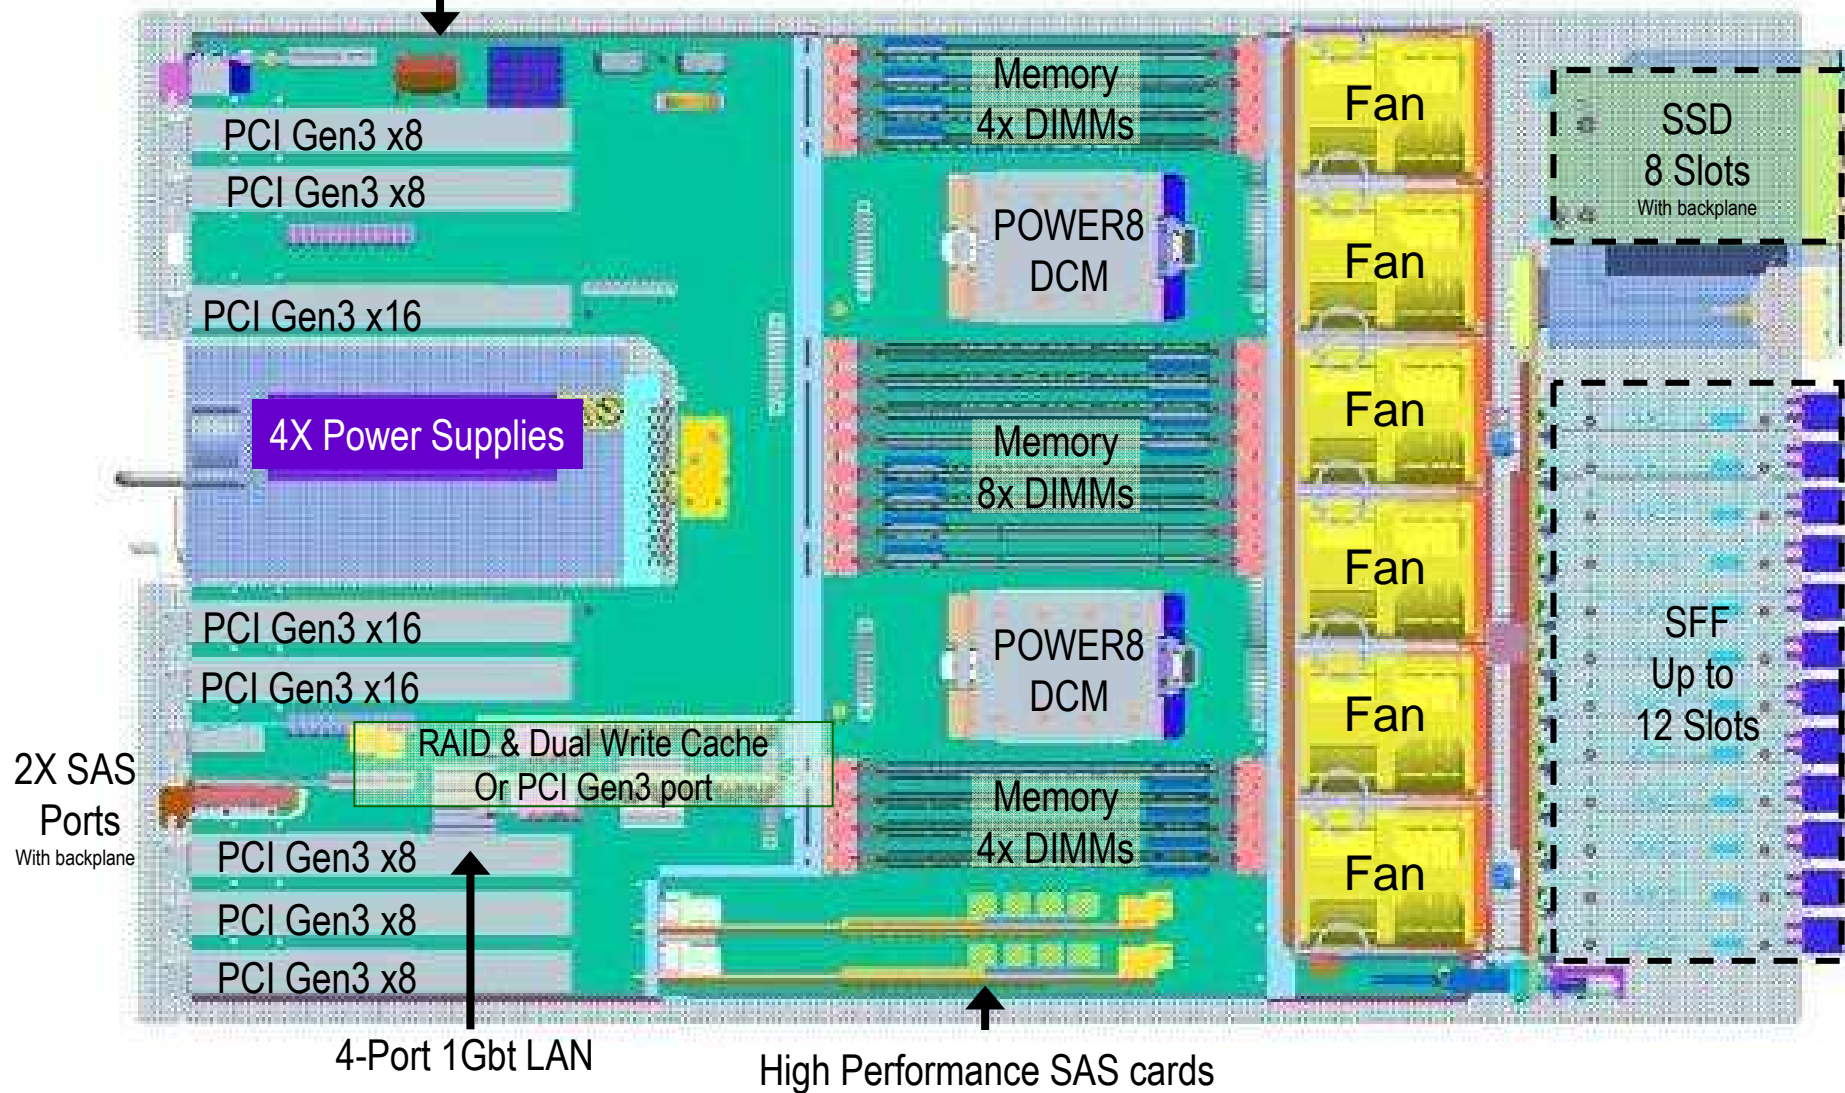
Hypervisor: PowerVM

OS: AIX, Linux **(not IBM i)**


3 years Warranty

Power S822 Layout

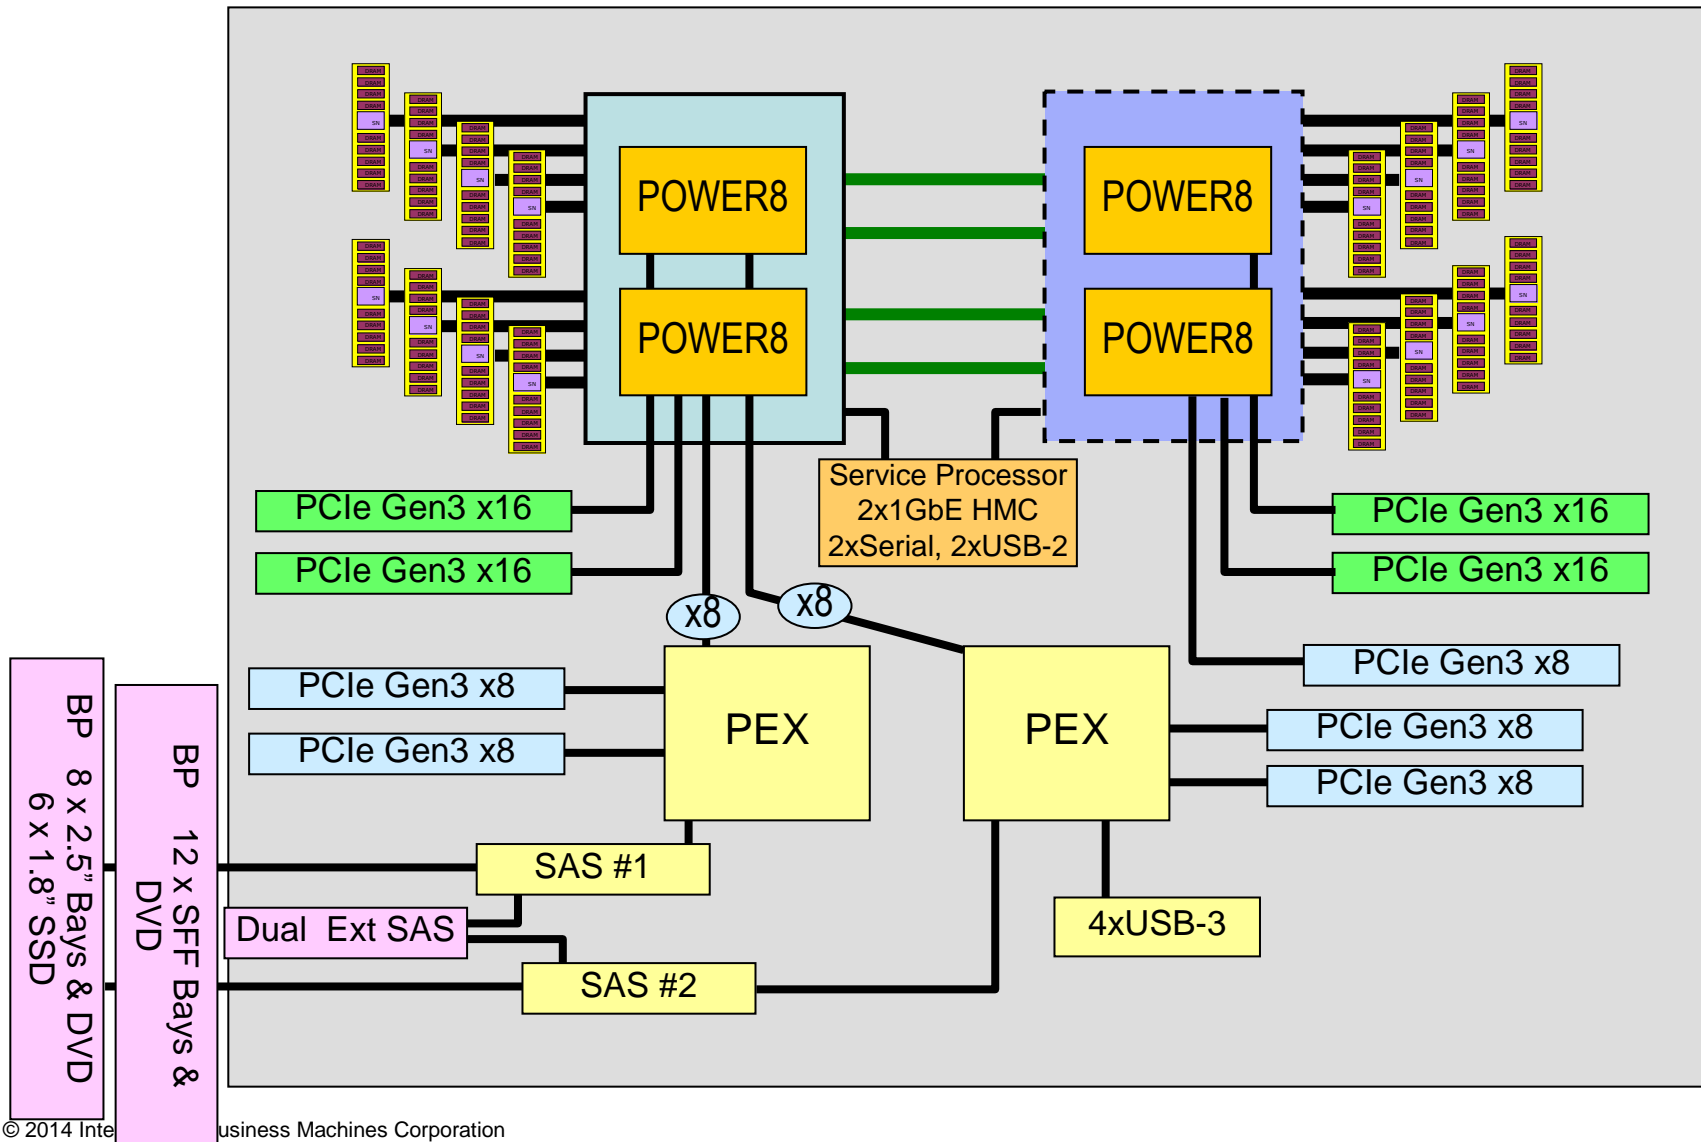
Rear service processor Card
2xSerial / 2x USB-2 / 2x HMC



Power S822 Rear View



Power S822 ... Planar Layout



Power S822 Comparison vs p730

	Power 730	Power System S822
Processor	POWER7+	POWER8
Sockets	2	1 (Upgradeable) / 2
Cores	8 / 12 / 16	6 / 10 or 12 / 20
Maximum Memory	512 GB @ 1066 MHz	256 / 512 GB @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	192 / 384 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 LP	6 / 9 PCIe Gen3 LP
CAPI (Capable slots)	N / A	Yes (one per socket)
PCIe Hot Plug Support	No	Yes
PCIe Expansion Drawers	Optional PCIe Gen1	SoD Gen3
IO bandwidth	40 GB/sec	192 GB/sec
Ethernet ports	Quad 1 Gbt in x4 slot	Quad 1 Gbt in x8 slot
SAS bays in system unit	3 or 6 SFF-1	12 SFF-3 Or 8 SFF-3 + 6 SSD
Integrated write cache	Optional 175MB	Optional effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	No	Yes (6 + 6)
Service Processor	Generation 1	Generation 2

Power S812L Scale-Out System

Single Socket

- Cores: 10 (3.4 GHz) / 12 (3.0 GHz)
- Memory: Up to 512 GB
- Slots: 6 PCIe Gen3 LP (Hotplug)

Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4), Serial (2), HMC (2)


Internal Storage

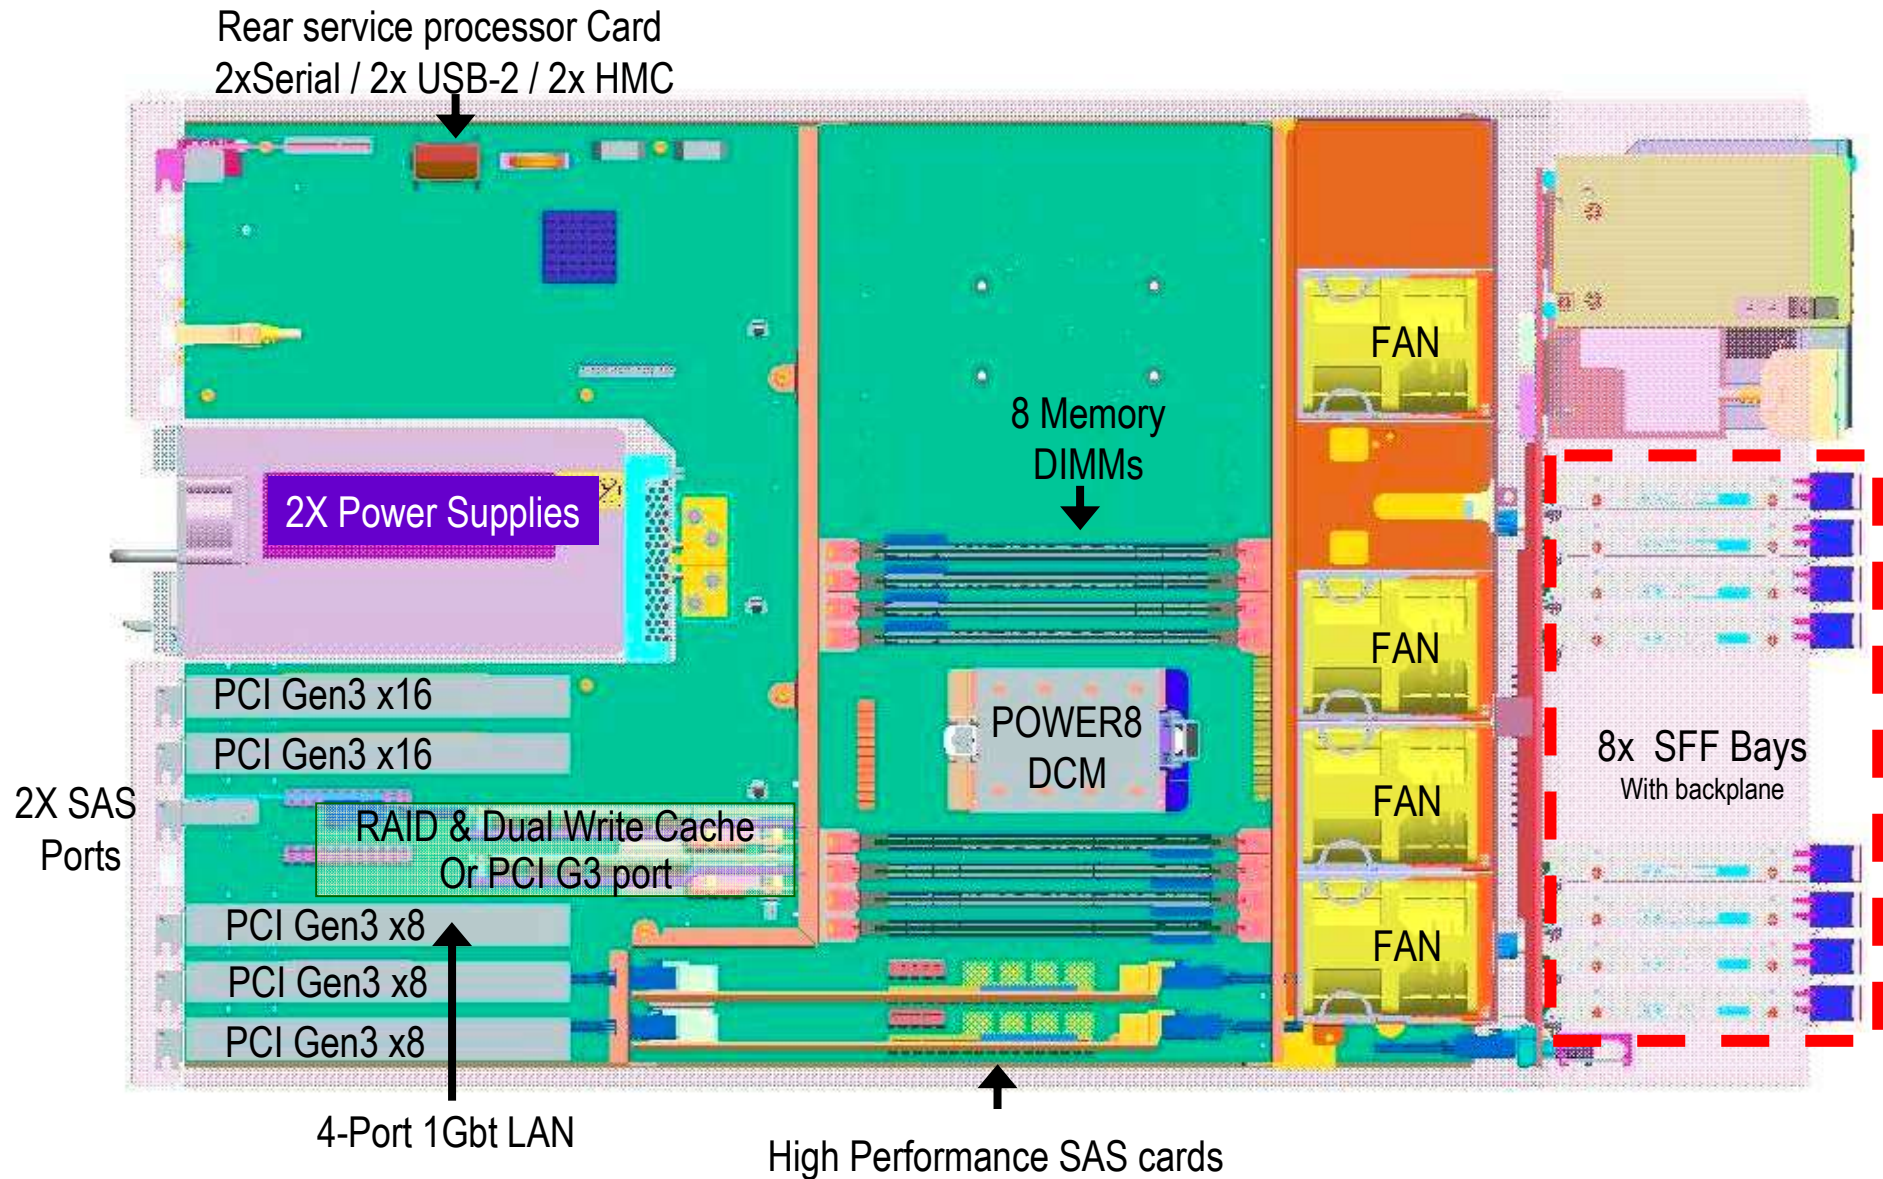
- DVD
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 8 SFF Bays with Easy Tier with 7GB write cache

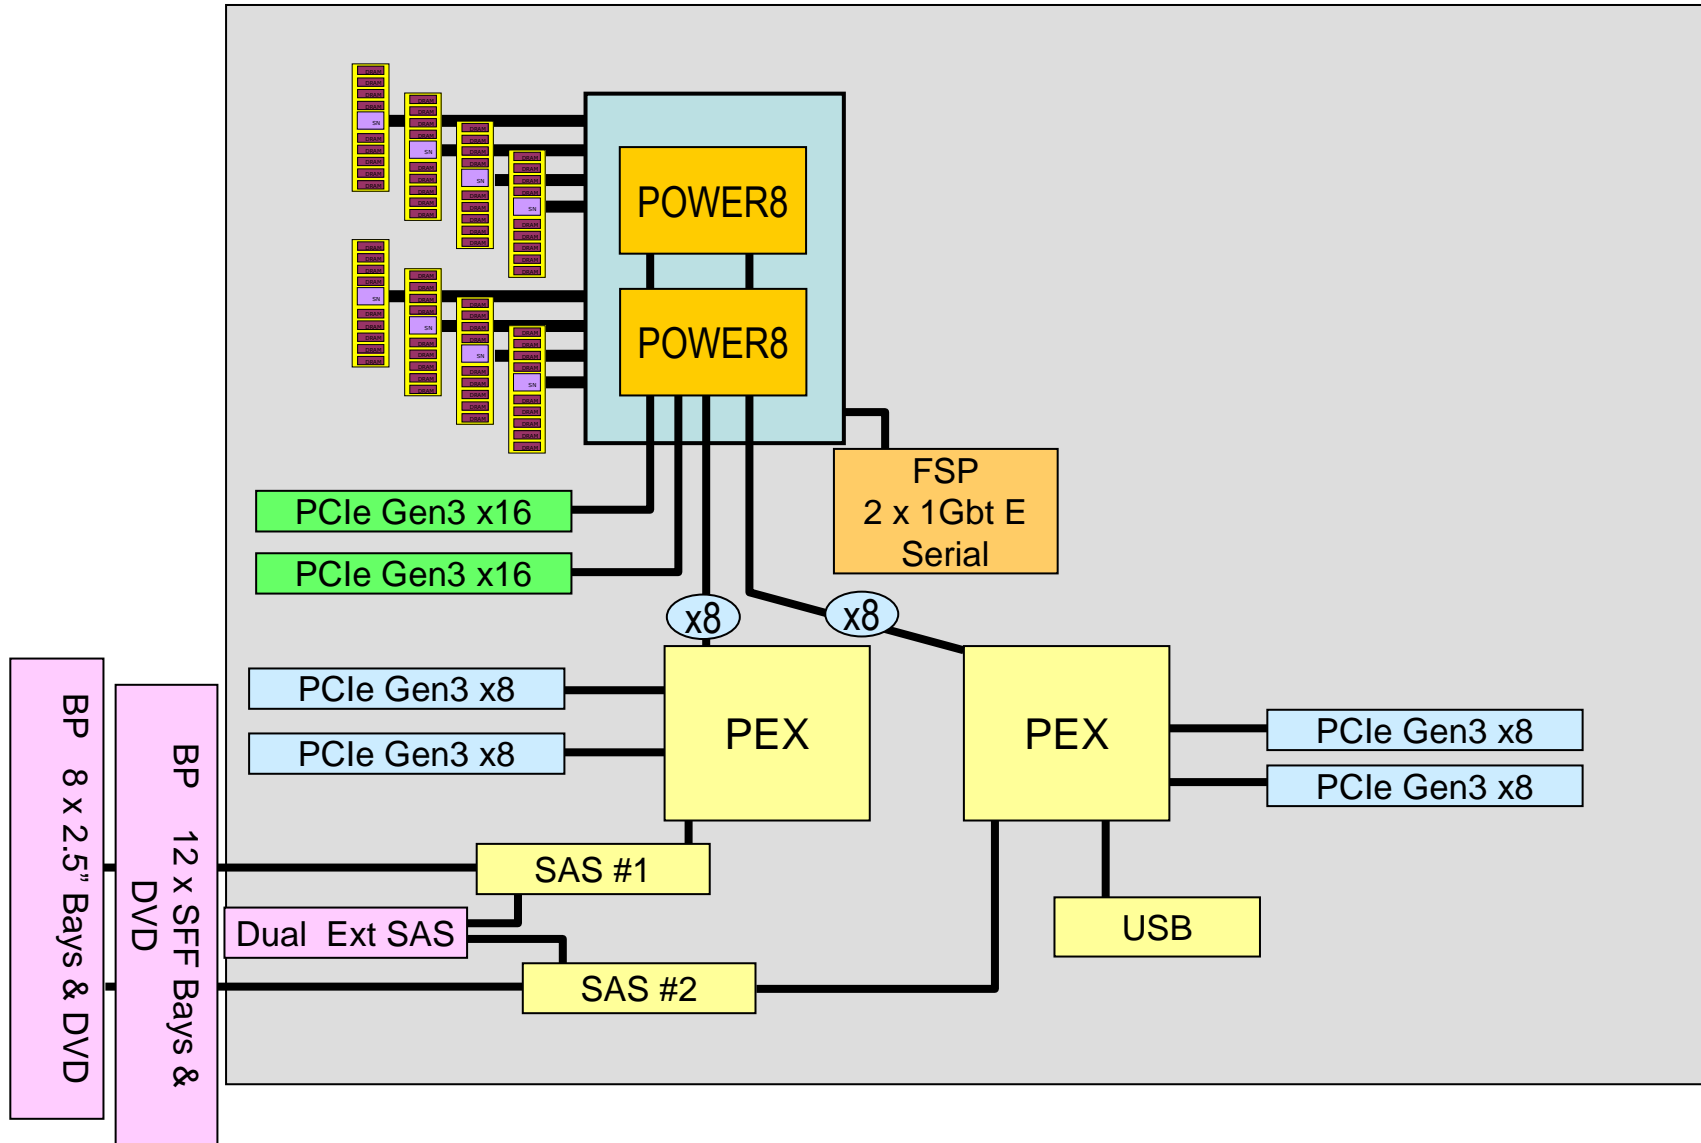
Hypervisor: PowerVM or **PowerKVM**

OS: Linux

3 years Warranty







Power S822L Scale-Out System

Dual Sockets

- Cores: 20 (3.4 GHz) or 24 (3.0 GHz)
- Memory: Up to 1 TB
- Slots: 9 PCIe Gen3 LP (Hotplug)



Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4), Serial (2), HMC (2)


Internal Storage

- DVD
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 8 SFF Bays & 6 1.8" SSD Bays with Easy Tier with 7GB write cache

Hypervisor: PowerVM or **PowerKVM**

OS: Linux


3 Years Warranty




Power S812L Scale-Out Comparison vs p730

	Power 730	Power System S822L
Processor	POWER7+	POWER8
Sockets	2	2
Cores	16	20 @ 3.4 GHz or 24 @ 3 GHz
Maximum Memory	512 GB @ 1066 MHz	512 GB @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	384 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 LP	9 PCIe Gen3 LP
CAPI (Capable slots)	N / A	Yes (one per socket)
PCIe Hot Plug Support	No	Yes
PCIe Expansion Drawers	Optional PCIe Gen1	SoD Gen3
IO bandwidth	40 GB/sec	192 GB/sec
Ethernet ports	Quad 1 Gbt in x4 slot	Quad 1 Gbt in x8 slot
SAS bays in system unit	3 or 6 SFF-1	12 SFF-3 Or 8 SFF-3 + 6 SSD
Integrated write cache	Opt 175MB	Opt effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	No	Yes (6 + 6)
Service Processor	Generation 1	Generation 2

4U Servers



Power S814 Layout




Power S814 Scale-Out System

Form Factor: 4U or Tower

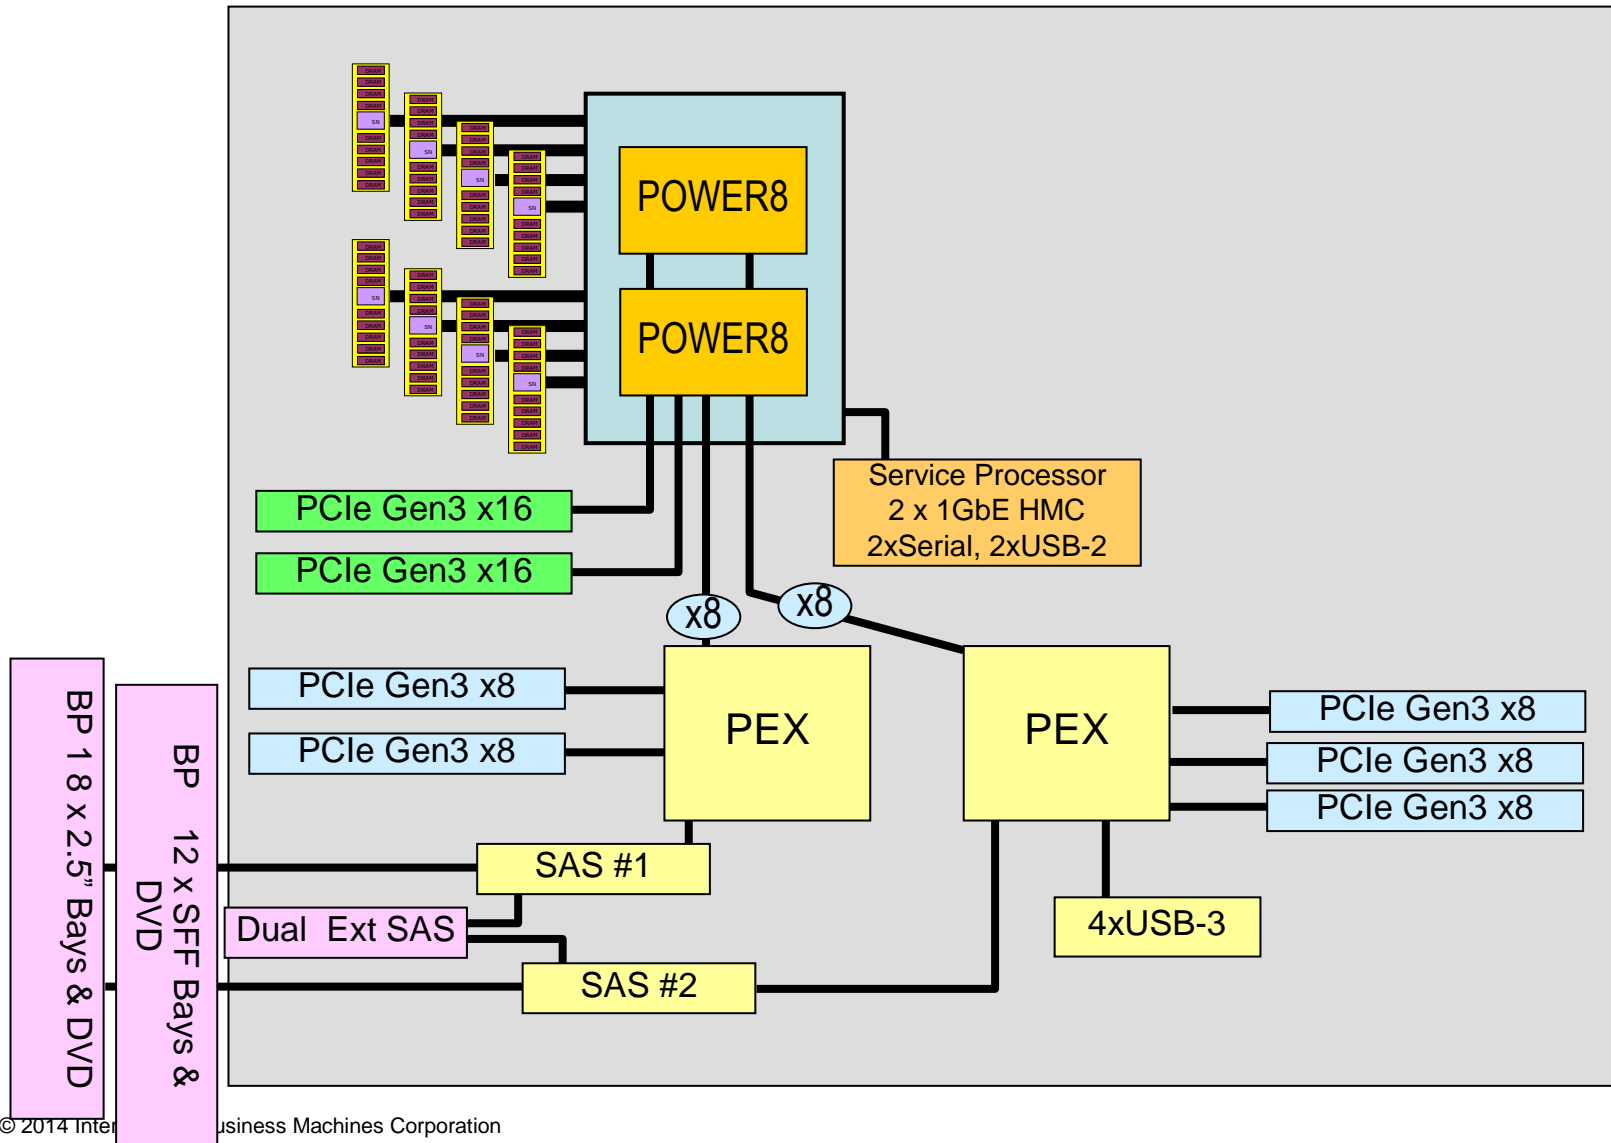
- Single Socket
 - Cores: 6 (3.0 GHz) or 8 (3.7 GHz)
 - Memory: Up to 512 GB
 - Slots: 7 PCIe Gen3 Full-high (Hotplug)

- Ethernet: Quad 1 Gbt / (x8 slot)
- Integrated ports: USB (4/5), Serial (2), HMC (2)
- Internal Storage
 - DVD
 - **No Tape !!**
 - 12 SFF Bays -- Split Backplane: 6 + 6
 - or 18 SFF Bays with Easy Tier with 7GB write cache

- Hypervisor: PowerVM
- OS: AIX, IBM i, Linux
- 3 years Warranty




Power S814 Planar Layout



Power S814 Comparison vs p720

	Power 720	Power System S814
Processor	POWER7+	POWER8
Sockets	1	1
Cores	4 / 6 / 8	6 / 8
Maximum Memory	512 GB @ 1066 MHz	512 GB @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	192 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 FH Opt 4 PCIe Gen2 LP	7 PCIe Gen3 FH
CAPI (Capable slots)	N / A	One
PCIe Hot Plug Support	No	Yes
IO bandwidth	40 GB/sec	96 GB/sec
Ethernet ports	Quad 1 Gbt (x4 slot)	Quad 1 Gbt (x8 Slot)
SAS bays in system unit	6 or 8 SFF-1 bays	12 SFF-3 bays Or 18 SFF-3 bays
Integrated write cache	Optional 175 MB	Optional effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	Yes (3 + 3)	Yes (6 + 6)
Service Processor	Generation 1	Generation 2

Power S824 Layout



Power S824 Scale-Out System



Single Socket populated

- Cores: 6 (3.8 GHz) or 8 (4.1 GHz)
- Memory: Up to 512 GB
- Slots: 7 PCIe Gen3 full-high (Hotplug)

Both Sockets populated

- Cores: 12 (3.8 GHz), 16 (4.1 GHz), or 24 (3.5 GHz)

Memory: Up to 1 TB


- Slots: 11 PCIe Gen3 full-high (Hotplug)

Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4/5), Serial (2), HMC (2)

Internal Storage

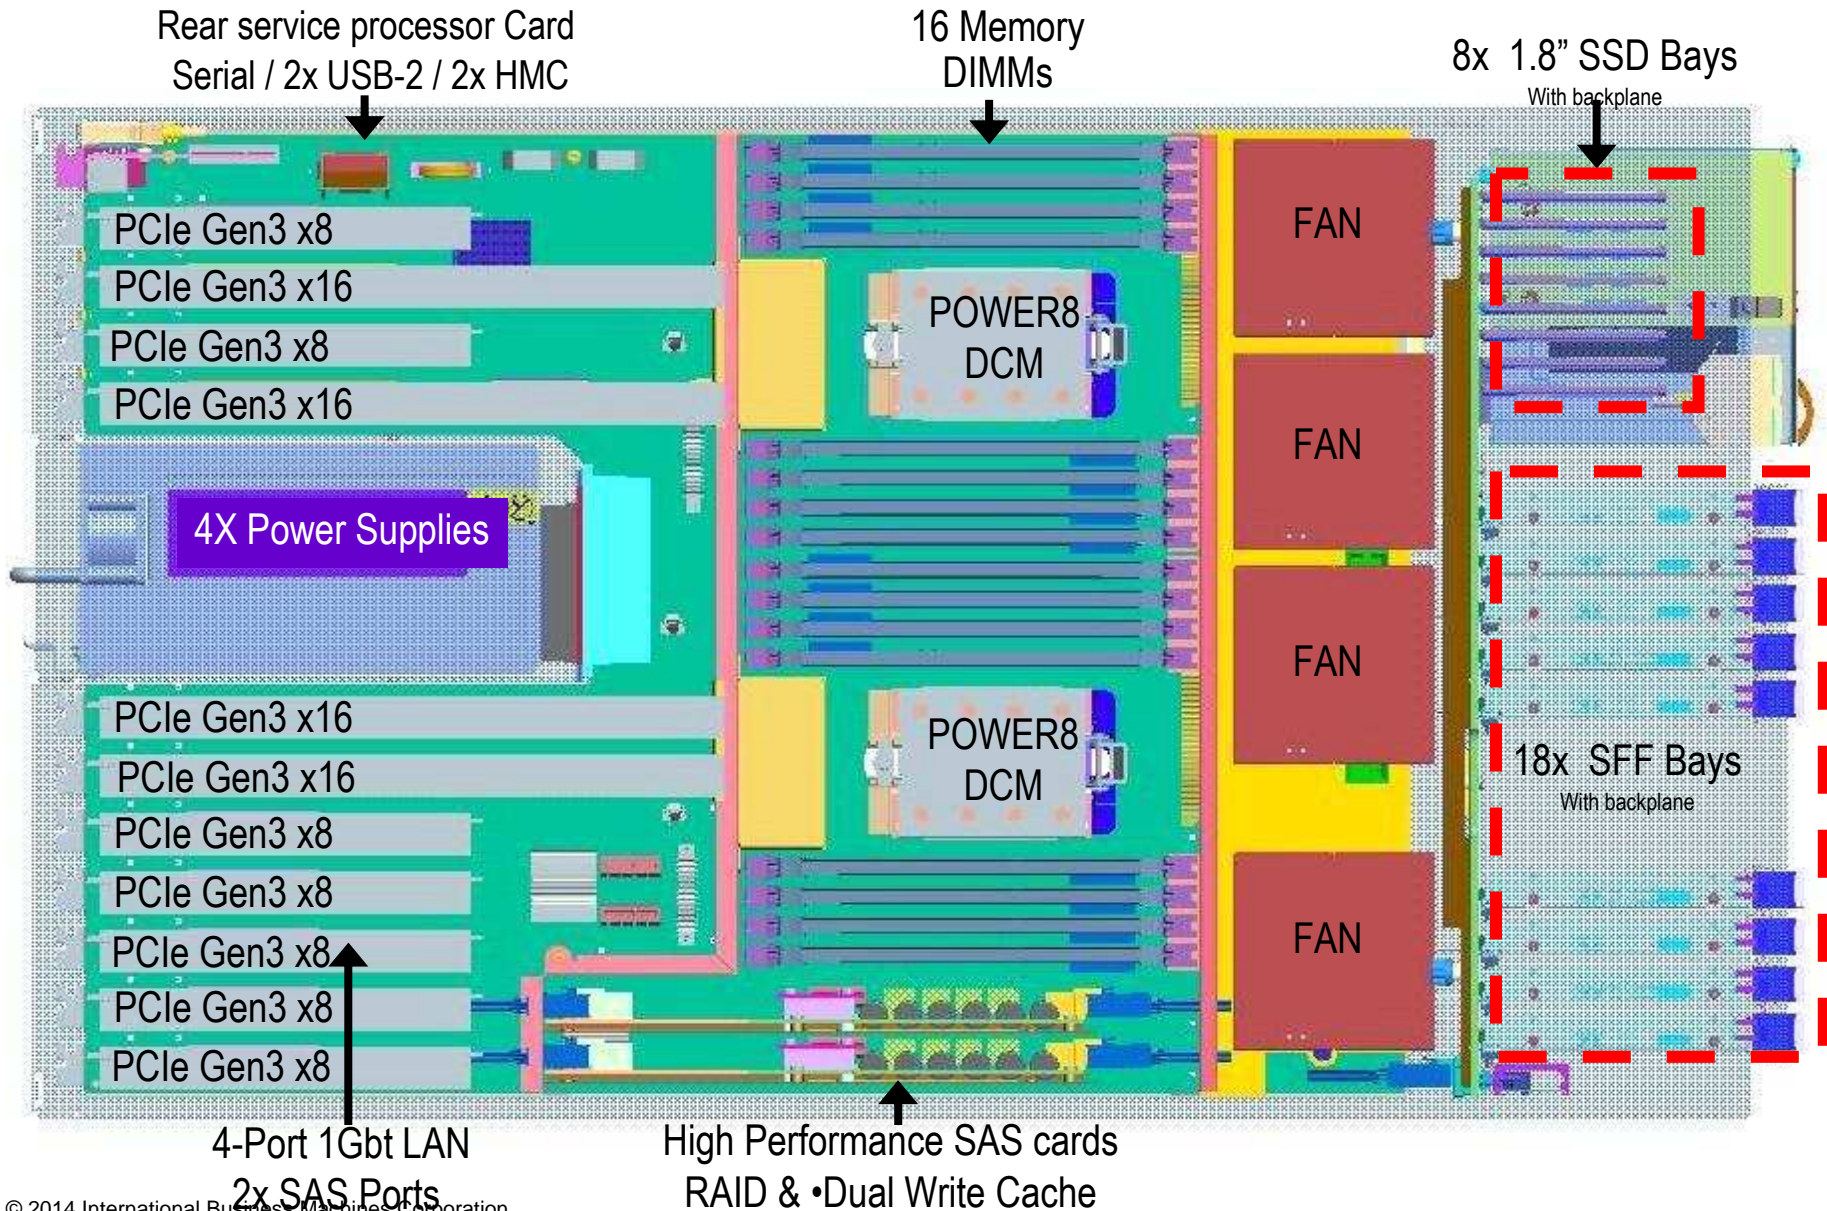
- DVD
- **No Tape !!**
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 18 SFF bays & 8 SSD bays with Easy Tier with 7GB write cache

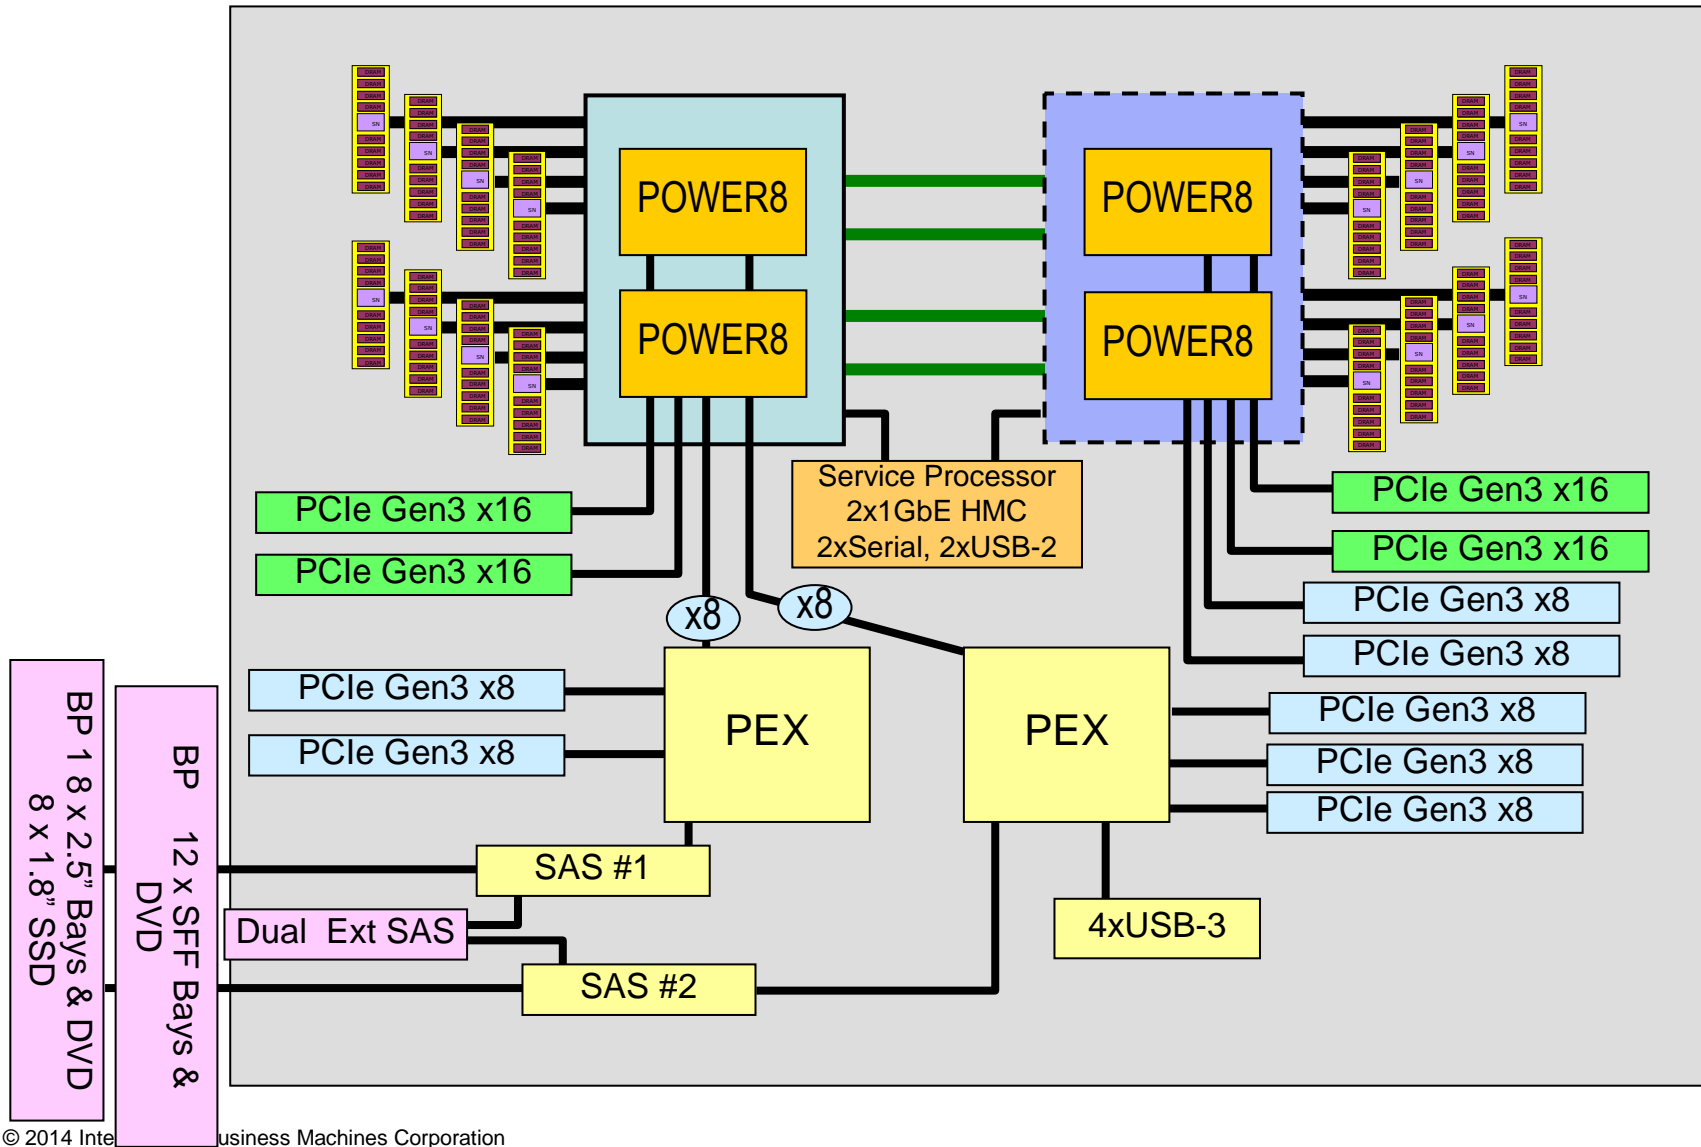


Hypervisor: PowerVM

OS: AIX, IBM i, Linux

3 years Warranty






Power S824 Scale-Out Comparison vs p740

	Power 740	Power System S824
Processor	POWER7+	POWER8
Sockets	1 (upgradeable) / 2	1 (upgradeable) / 2
Max Cores	8 / 16	8 / 24
Maximum Memory	512 GB / 1TB @ 1066 MHz	512GB / 1TB @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	384 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 FH Opt 4 PCIe Gen2 LP	7 / 11 PCIe Gen3 FH
CAPI (Capable slots)	N / A	Two
PCIe Hot Plug Support	No	Yes
IO bandwidth	60 GB/sec	192 GB/sec
Ethernet ports	Quad 1 Gbt (x4 slot)	Quad 1 Gbt (x8 Slot)
SAS bays in system unit	6 or 8 SFF-1	12 SFF-3 bays Or 18 SFF-3 + 8 SSD bays
Integrated write cache	Optional 175 MB	Optional effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	Yes (3 + 3)	Yes (6 + 6)
Service Processor	Generation 1	Generation 2

	Power System 740	Power System S814	Power System S824
Processor Offerings	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	1x or 2x 6 -core 3.89 GHz 1x or 2x 8 -core 4.15 GHz 2x 12 -core 3.52 GHz
DDR3 Memory features	8 / 16 / 32 / 64GB 8GB to 256GB / socket	16 / 32 / 64GB 16GB to 512GB	16 / 32 / 64GB 32GB to 1024GB
OS Support	AIX: 5.3 / 6.1 / 7.1 IBM i: 6.1 & 7.1 & 7.2 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
Internal SAS Bays	Up to 8 HDD or SSD	Up to 18 SFF and 8 SSD	Up to 18 SFF and 8 SSD
PCIe Slots in system unit	Six Gen2 FH Opt four Gen2 LP	Seven Gen3 FH Hot Plug One CAPI capable	Eleven Gen3 FH Hot Plug Two CAPI capable
SAS bays in system unit	6 SFF1 (3+3) or 8 SFF-1	12 SFF-3 bays (6+6) Or 18 SFF-3 bays	12 SFF-3 bays (6+6) Or 18 SFF-3 bays + 8 SSD
Integrated SAS write cache	Opt 175MB	Opt 1.8GB effectively 7GB	Opt 1.8GB effectively 7GB
Integrated SAS/SATA Cntrl	Standard: RAID 0, 1, & 10 Optional: RAID 5 & 6	Standard RAID 0, 1, 5, 6, 10 Optional Easy Tier function	Standard RAID 0, 1, 5, 6, 10 Optional Easy Tier function
Ethernet	Quad 10/100/1000	Quad 10/100/1000	Quad 10/100/1000
Media Bays	1 Slim-line & 1 Half Height (Optional)	1 Slim-line	1 Slim-line
PCIe IO Drawers	PCIe Gen1: 2 on 1S, 4 on 2S	N / A, SOD PCIe Gen3	N / A, SOD PCIe Gen3
Power requirement	100V to 240V AC 1S 200 to 240V AC 2S	100V to 240V AC 200V to 240V AC	200V to 240V AC
Red Power & Cooling	Optional / Standard		Standard
Warranty	3 Years		


Performance




Power 740 / POWER8 S824 ... rPerf Comparisons




Performance per Core



Performance per Socket



Performance per KW



CPW

• **720 POWER7+ (1 socket)**

✓ 4-core	3.6 GHz	28,400
✓ 6-core	3.6 GHz	42,400
✓ 8-core	3.6 GHz	56,300

• **S814 (1 socket)**

✓ 4-core	xxxxx	xxxxxx
✓ 6-core	3.0 GHz	59,500
✓ 8-core	3.7 GHz	85,500

• **740 POWER7+ (1 or 2 socket)**

✓ 6-core	4.2 GHz	49,000
✓ 12-core	4.2 GHz	91,700
✓ 8-core	3.6 GHz	56,300
✓ 16-core	3.6 GHz	106,500
✓ 8-core	4.2 GHz	64,500
✓ 16-core	4.2 GHz	120,000


• **S824 (1 or 2 socket)**

✓ 6-core	3.8 GHz	72,000
✓ 12-core	3.8 GHz	130,000
✓ 8-core	4.1 GHz	94,500
✓ 16-core	4.1 GHz	173,500
✓ 12-core	1-socket not offered	
✓ 24-core	3.5 GHz	230,500


Power 740 / POWER8 S824 ... CPW Comparisons



Performance per Core



Performance Dual Socket




POWER8 delivers Over Twice the Throughput vs Ivy Bridge-EP at 47% Lower Cost



This is an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. The results were obtained under laboratory conditions, and not in an actual customer environment. IBM's internal workload studies are not benchmark applications, nor are they based on any benchmark standard. As such, customer applications, differences in the stack deployed, and other systems variations or testing conditions may produce different results and may vary based on actual configuration, applications, specific queries and other variables in a production environment. Prices, where applicable, are based on published US list prices for both IBM and competitor, and the cost calculation compares the cost per request for the 3yr life of the machine. 3 year total cost of acquisition comparisons are based on similar expected hardware, software, service & support offerings.


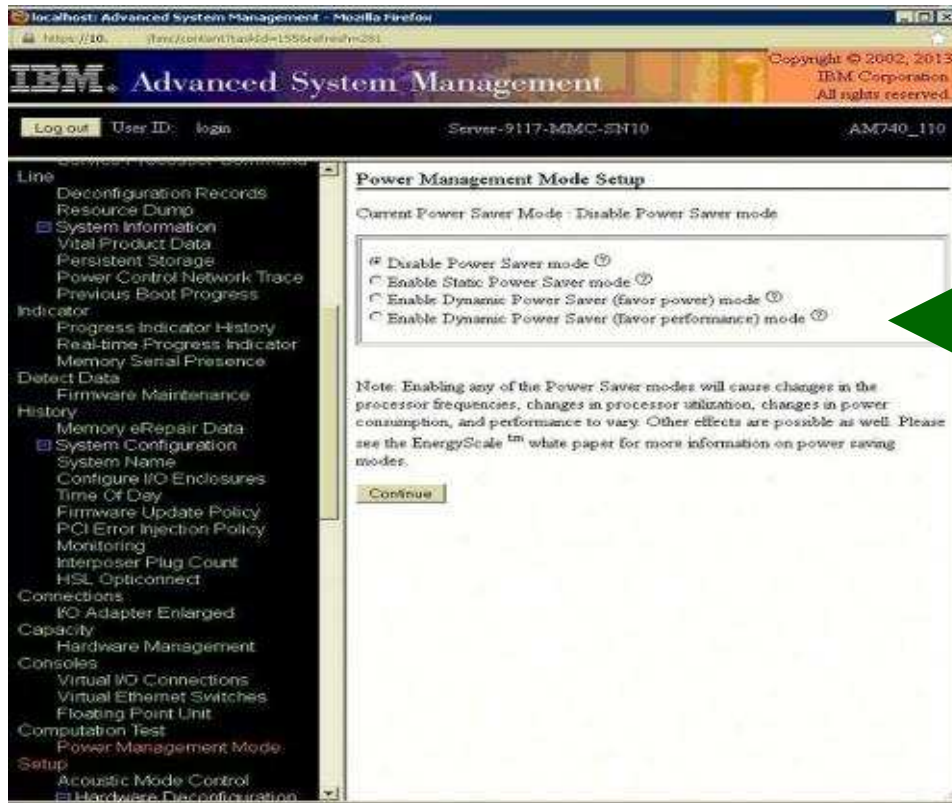
POWER8 Packs Up To 2.9x More Virtual Machines than Intel on Same Number of Cores

Online Banking Workload



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

Operating Systems




Power **KVM**



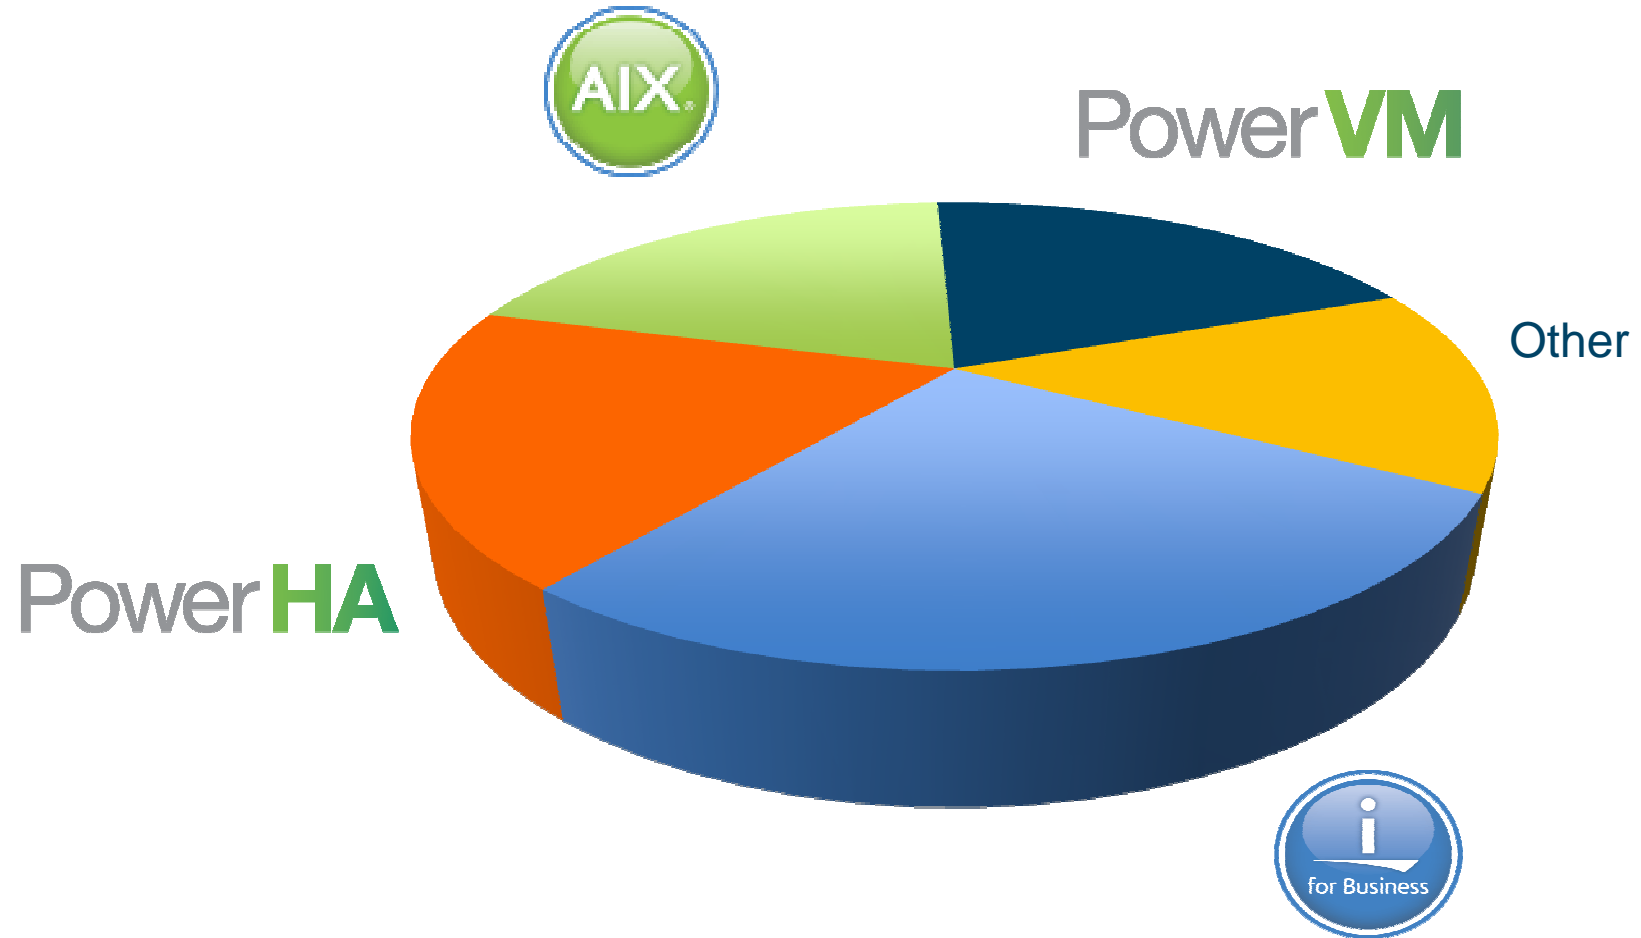
Niveaux AIX supportés

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



IBM i Momentum



%age Contribution to Power Systems Software Revenue in in 2013




- ✓ Mise à niveau intermédiaire de Technology Refresh pour l' IBM i 7.1 par le biais de PTFs ★
- ✓ Avantages des Technology Refreshes intermédiaires ...
 - Amènent le support des nouvelles fonctions et des nouveaux dispositifs E/S
 - Simples à installer sur la version en cours, ne nécessitant pas un arrêt de production.
- ✓ **Pas de nouvelle version ni en 2012, ni en 2013.... mais en 2014 !!**

Les nouvelles fonctionnalités ont été incluses dans les différents Technology Refreshes disponibles depuis Octobre 2010

** All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

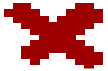
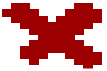













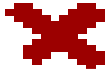









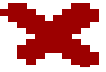
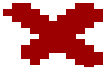
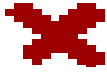
IBM i 7.2 Highlights

- Powerful **new features of DB2® for i** ensures security of the data in a modern environment of mobile, social and network access
- **IBM Navigator for i** extends system management capabilities to manage and monitor performance services
- **PowerHA SystemMirror for i Express Edition** introduces HyperSwap improves resiliency value proposition
- Additional flexibility in configuration for **I/O** with new virtualization: Ethernet through SRIOV and WAN over LAN”
- **Backup and Recovery Management Systems (BRMS)** extends Hub function for easier management
- Support for new **POWER8** systems, and new features for POWER8 and POWER7/7+
- **Integrated Security** SSO application suite extended to include FTP and Telnet authentication with Kerberos
- Many enhancements to the portfolio of IBM i products from **IBM Software Group**
- More highlights extending and enhancing the value proposition of IBM i integration
-



IBM i System Support

<http://www-947.ibm.com/systems/support/i/planning/upgrade/osmapping.html>

Servers	IBM i 5.4 ¹	IBM i 6.1	IBM i 7.1	IBM i 7.2
POWER8			 ⁴	
POWER7/7+ PS700/701/702/730/704, Power 710, 720, 730, 740, 750, 760, 770, 780, 795, Pureflex p260/460		 ^{2 / 3}		
POWER6 JS12, 22, 23/43, 550* 560				 ⁵
POWER6 520, 550*, 570, 595				
POWER5/5+ 515, 520, 525, 550, 570, 595				
800, 810, 825, 870, 890				
270, 820, 830, 840				


- 1 – IBM i V5R4 is no longer marketed or supported other than through extended service contracts
- 2 – POWER7+ 750/760 do not support native I/O.
- 3 - IBM i 6.1 in PureFlex must be client of 7.1 or later
- 4 – Requires Technology Refresh 8
- 5 – no IOP or HSL support 5

Les versions IBM i supportées en POWER8

IBM i 7.1 TR8	
POWER7	Max Scale = 32 cores (SMT4) Max Partition = 64 cores (SMT4) Threads = ST, SMT2, SMT4 up to 256 threads in single partition
POWER8	Max Scale = 32 cores (SMT8) Max Partition = 64 cores (SMT4) Threads = ST, SMT2, SMT4, SMT8 up to 256 threads / single partition

IBM i 7.2	
POWER7	Max Scale = 32 cores (SMT4) Max Partition = 96 cores (SMT4) Threads = ST, SMT2, SMT4 up to 384 threads in single partition
POWER8	Max Scale = 48 cores (SMT8) Max Partition = 96 cores (SMT8) Threads = ST, SMT2, SMT4, SMT8 up to 768 threads / single partition

Les distributions Linux supportées



	1H / 2014	2H / 2014
RHEL6	RHEL 6.5 P7 Mode in P8	RHEL 6.6 P7 mode in P8
RHEL 7	RHEL 7.0 POWER8 Support	
SLES 11	SLES 11 + SP3 P7 Mode in P8	
SLES 12		SLES 12 (LE) POWER8 Support
Ubuntu (LE)	14.04.00/01 P8 Support	14.04.00/02
Debian	LE Introduction POWER8 Support	LE Update
PowerVM	2.2.3.3	2.2.3.X

Power Virtualization Options

PowerKVM



Q2 2014
Initial Offering

PowerKVM provides an [open source choice](#) for Power Virtualization for Linux workloads. Best for clients that aren't familiar with Power and [Linux centric admins](#).

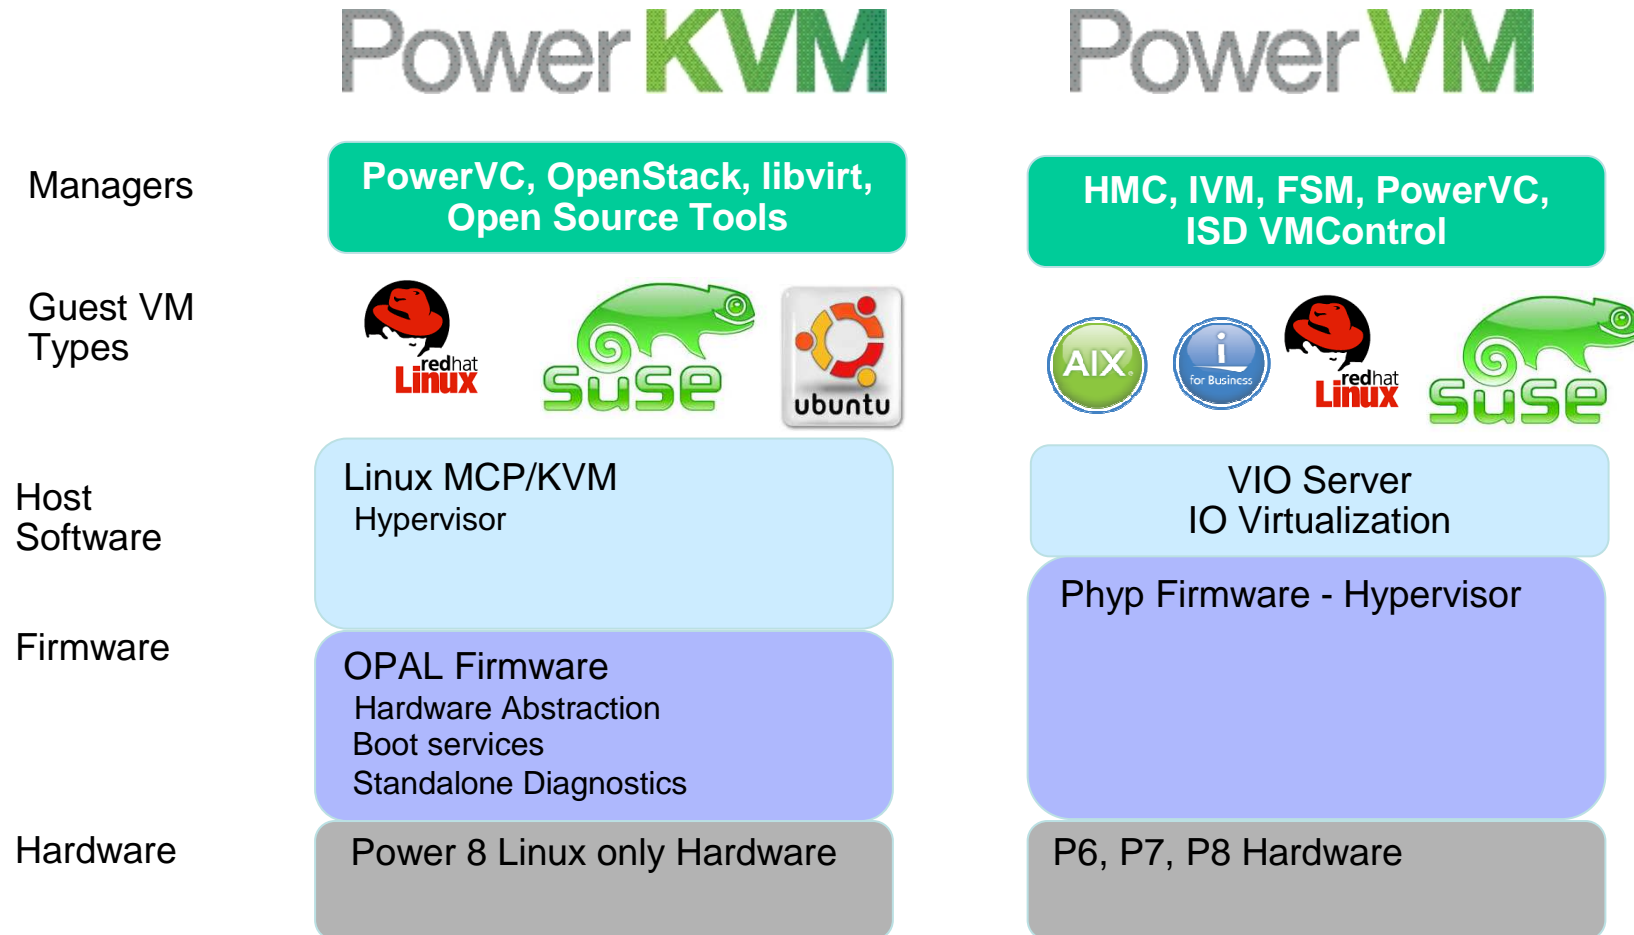
PowerVM



2004
Initial Offering

PowerVM is Power Virtualization that will continue to be enhanced to support [AIX, IBM i Workloads](#) as well as Linux Workloads


Linux on Power enables open source virtualization with PowerKVM



Key Take-Aways for the Next Generation of Power Systems

- Delivering **2X** the performance of x86 Ivy Bridge across a **broad** range of **workloads**
- **Built for data**, delivering **up to 50X** differentiation for data-centric workloads with expansive acceleration capability
- **Optimized for Linux** with simple porting, easy management and better price/performance than x86. Power is ready for the cloud.
- **Open server platform** with an open community innovating across the end-to-end IT stack

Upgrades



Upgrades

No same-serial-number upgrades from Power 7xx to Power S8xx

-> Could not pass financial/accounting rules to do so

Note the Power 770/780 D mdls have an SOD for same-serial-number upgrades and another SOD covering mobile activations

IBM i Entitlement Transfer Offering Group Definitions

Tier	Model	Structure	Processor Group
Large	795	Per core	P50
	780		Group 3
Medium	770	Per core	P30
	760		Group 2
Small	S824 750 / 740 / 730	Per core	P20 Group 1
	PureFlex/Flex: Flex p260 Compute Node 8/16 - core Flex p460 Compute Node 16/32 - core	Per core and user	P10 Group 6
	PureFlex/Flex: Flex System p260 – 7895-23A	Per core and user	P05 Group 6
	S814 720 6/8-core 710 6/8-core PS701/703 8/16 – core PS702/704 16/32 – core	Per core and user	P10 Group 5
	720 4-core 710 4-core PS700 4-core	Per core and user	P05 Group 4


For the 3Q2013 PureFlex/Flex P05 7895-23A:

- Currently, clients can transfer IBM i processor and user entitlements to the 7895-23A via the software-only path in econfig (transfers within groups 4, 5, 6).
- Initial order path: in order to transfer IBM i processor/user entitlements in the initial order path for 7895-23A, econfig support will start in November 2013.
- Transfers from Groups 1/2/3 to the 7895-23A will not be supported.

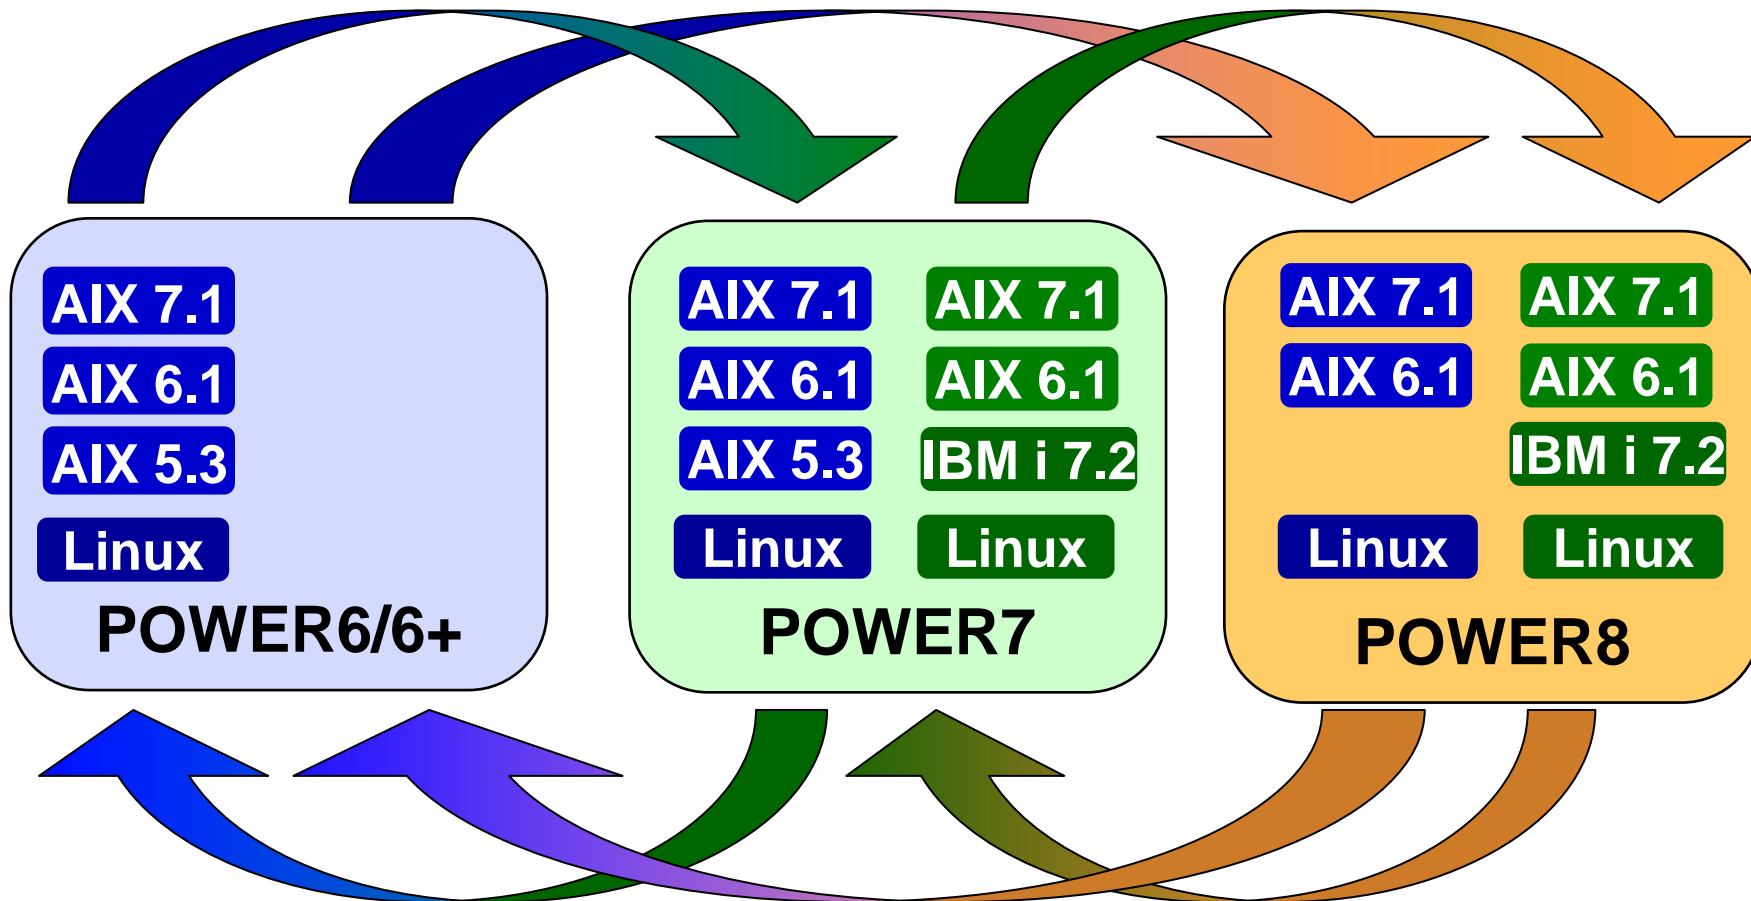
IBM i entitlement transfers can be done...

- within groups 1, 2, and 3
- within groups 4, 5, and 6
- from groups 1, 2, and 3 to group 6 (with the exception of 7895-23A: no transfer from 1/2/3 to the P05 7895-23A)

Mobility with Power8



POWER6 / POWER7 / POWER8 Partition Mobility



Leverage POWER6 / POWER7 Compatibility Modes

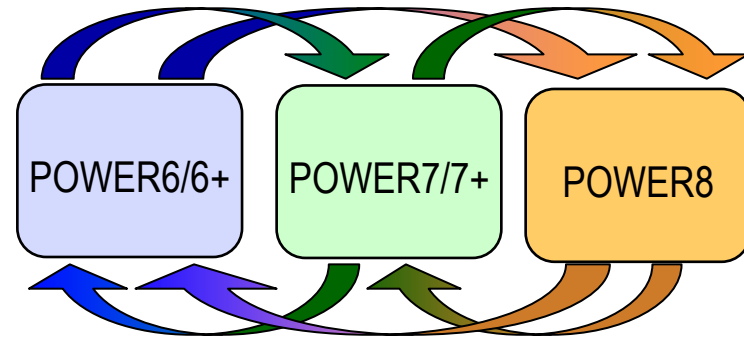
LPAR Migrate between POWER6 / POWER7 / POWER8 Servers

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

POWER6 MODE (and POWER6+ Mode)*	POWER7 MODE (No POWER7+ Mode)	POWER8 MODE
2-Thread SMT	4-Thread SMT, <i>IntelliThreads</i>	8-Thread SMT
8 Protection Keys *(16 in P6+ Mode)	32 Protection Keys User Writeable AMR	32 Protection Keys User Writeable AMR
VMX (Vector Multimedia Extension / AltiVec)	VSX (Vector Scalar Extension)	VSX2, In-Core Encryption Acceleration
Affinity OFF by Default	CPU/Memory Affinity Enhancements ON by Default, HomeNode, 3-tier Memory, MicroPartition Affinity	HW Memory Affinity Tracking Assists, MicroPartition Prefetch, Concurrent LPARs per Core
64-core/128-thread Scaling	64-core / 256-thread Scaling 256-core / 1024-thread Scaling	> 1024-thread Scaling Hybrid Threads Transactional Memory Active System Optimization HW Assists
N/A	Active Memory Expansion	HW Accelerated/Assisted Active Memory Expansion
N/A	P7+ : AME compression acceleration and Encryption acceleration	Coherent Accelerator / FPGA Attach


POWER8 LPM

- Can Utilize 10 Gbt Ethernet
 - Minimize movement time
 - Move more LPARs concurrently
- Used for Migrations
 - POWER6 to POWER8
 - POWER7 to POWER8
- Used for Maintenance Activities...
 - Primary resource for concurrent maintenance functions for POWER8
 - No Application outages
 - Mode status of LPAR dictates LPM movement
- Used for Workload balancing
 - Provide better distribution server resources




PCIe Slots

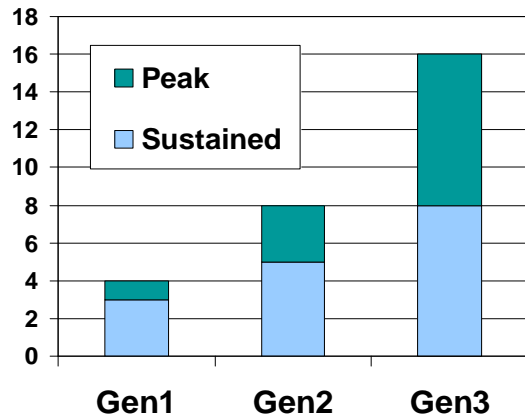
New adapters



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. More lanes = more physical connections = more bandwidth
- A x8 card can be placed in a x16 slot, but only uses half the connections


New I/O Announced with Scale-Out Servers

- **SOD for PCIe3 expansion drawers**
 - **Not planned for 2014 Announce**
 - Will connect to server using two x16 PCIe slots
 - Important planning consideration for clients with lots of I/O.
 - 👉 Emphasize multi-port adapters available to virtualize
 - 👉 Emphasize bandwidth of Gen3 PCIe slots to virtualize
- **Disk/HDD**
 - New capacity - 600GB 15k rpm
 - New 4k block drives
- **SSD new 1.8-inch 387GB for POWER8 servers SSD cage**

New SAS RAID Adapters ...



#EJ0L
Huge cache




#EJ0J
No cache

- ✓ First PCIe Gen3 adapters introduced by Power Systems
- ✓ Most powerful SAS adapters ever created by Power Systems *
 - * IBM Development team not aware of a more powerful SAS adapter in the world
 - Supports 2X more SSD devices than Gen2 SAS adapter
 - Up to 10X more performance than Gen1 SAS Adapter
 - 2X more performance than Gen2 SAS adapter
- ✓ #EJ0L effectively has 6X more write cache (up to 12GB)
- ✓ Better prices than predecessor Gen2 SAS adapters
 - Europe prices per unit: #EJ0J ... 2511,79€ and #EJ0L ... 9976,33€

Prices shown are suggested Europe IBM list prices and are subject to change without notice; reseller prices may vary

#EJ0L Large Cache SAS Adapter Ports

- PCIe x8 Gen-3 with 12GB Cache RAID, SAS-2 Adapter Quad-port 6Gb
- Single wide adapter – PCIe Gen3 full high
 - Supported in Gen1 slots
 - New ASIC processor + IBM VersaRAID firmware
 - OS support: Linux, AIX, VIOS, IBM i
- Mirrored Non-Volatile Write Cache (*No batteries*)
 - 3 GB Physique, up to 12GB compressed
- Protection available: RAID 0, 10, 5, 6, 10T2, 5T2, 6T2, T10-DIF
- Easy Tier Function and Native 4K Sector support
- #EJ0L uses Mini-SAS HD narrow cables
- Up to 48 SSDs and 96 total devices (*per pair of #EJ0L*)
- Bottom two ports (labeled T0 and T1)
 - ☞ Attach high performance I/O drawers to bottom two ports, especially SSD.
 - SSD attachment not supported to top connectors T2 and T3**
 - Top one or two ports used for adapter-to-adapter communications cable (AA) unless I/O drawer attached to those ports




#EJ0L
Huge cache
CCIN: 57CE

Note: cables are not hot add (same as always) – adapter must be powered off to add. If cable is disconnected, power off adapter and then power back on after reconnecting (same as always).

PCIe3 SAS RAID Adapter - #EJ0J (zero cache)

Premier SAS PCIe adapter – IBM technology designed for SSD
(and HDD)

- **Same chip technology as #EJ0L, but zero cache**
 - No cache = physically smaller & lower cost & optional pairing
- **PCIe Gen3 adapter**
 - Up to **4GB/sec** transfer (GB/s limited by PCIe Gen1 Slot Jan 2014)
 - Four 6Gb SAS ports
 - Each port with four x4 channels
- **Supports SAS HDD and/or SSD**
 - All protection options – including RAID-5 and RAID-6
 - RAID, mirroring, etc
 - RAID array sizes from 3 - 32



#EJ0J
Zero cache
CCIN: 57B4


Optional paired adapters for redundancy and performance.
No card-to-card AA cables used when paired.
Cards use cabling to SAS drawer enclosure for cross communication

PCIe3 SAS Tape Adapter - #EJ10

LTO-5/LTO-6 SAS tape adapter

- Supports full bandwidth of LTO-5/LTO-6
- PCIe Gen3 adapter
- Up to 4GB/sec transfer (GB/s limited by PCIe Gen1 Slot)
- 4 SAS ports
 - 6Gbps ports
- DVD and other SAS tape drives supported (*not with #EJ0X*)
- Physically is same card as #EJ0J, but has different feat code to help IBM config tools like eConfig understand card's use and appropriate cabling. eConfig does not know how to swap usage to disk/SSD. Use no-charge RPQ to adjust IBM feature records for any "re-purposing".
- Tape adapter always configured as single card and NOT optionally paired like disk/SSD SAS controllers

Note: Attaching both disk & tape to the same adapter **NOT** supported



#EJ10
SAS Tape adapter

CCIN 57B4

PCIe3 vs PCIe1 SAS Tape Adapters




	#5901/5278 LP	New #EJ10
PCIe card technology	PCIe Gen1	PCIe Gen 3
SAS ports – quantity per card	2	4
SAS ports – Gb/sec	3Gb	6Gb
Max tape drives per card	2	8
SAS/SATA DVD supported	Y	Y
LTO-2, LTO-3 (SCSI vs SAS)	N	N
LTO-4	Y	N
LTO-5, LTO-6	Y, but max bandwidth limited	Y, with full bandwidth
DAT160	Y	N
Attach disk and tape on the same card at the same time	Not supported	Not supported
Models supported on	all	POWER8
Feat code	#5901	#EJ10
Europe List price on POWER8	1008,42€	1077,17€

2X more
2X more
4X more

Prices shown are suggested Europe IBM list prices and are subject to change without notice; reseller prices may vary

PCIe3 SAS Tape Adapter - #EJ0X Cabling

- **SAS Mini-HD Narrow YE1 or AE1 cables**
 - AE1 (#ECBY) is 4 meter long
 - YE1 (#ECBZ) is 3 meter long
- **Up to one cable per port**
- **Max four cables per adapter**
- **One tape drive per AE1 cable**
 - Thus max 4 tape drives per card
- **Up to two tapes per YE1 cable**
 - Thus max 8 tape drives per card
- Same SAS cables can also be ordered as feature codes of the 7226-1U3 Removable Media Drawer
 - AE1 is #5507 or #9850
 - YE1 is #5509



CCIN 57B4

Tape adapter always configured as single card and NOT optionally paired like disk/SSD SAS controllers

Attaching both disk & tape to the same adapter NOT supported

PCIe Bisync Adapter for IBM i




- **PCIe Bisync Adapter (#EN13)**
- IBM i unique, **not used for AIX/Linux**
- Full-high only adapter, no low profile
- One port for Bisync
- For clients stuck back on this older protocol

- Same price as #2893/2894 WAN adapter

- Use lower price #5289/5290 PCIe Async Adapter for non-Bisync communication needs
 - IBM i 7.1 is pre-req for Async card

#5289/#5290 not supported on P8 for the moment !! #2893 is the alternative with IBM i only

- Prior product (#2893/2894) which also had this function will be withdrawn fairly soon due to shortage of components




Port to attach Bisync modem

PCIe Bisync Adapter for IBM i



- PCIe Bisync Adapter (#EN13)
- One RVX port for Bisync provided
- Will also need:
 - Standard modem
 - Cable IBM offers:
 - #0348 - V.24/EIA232 20-Ft PCI Cable
 - #0353 - V.35 20-Ft PCI Cable
 - #0359 - X.21 20-Ft PCI Cable
- **CCIN = 576C**
- IBM i 7.1 TR8 or later supported on POWER8



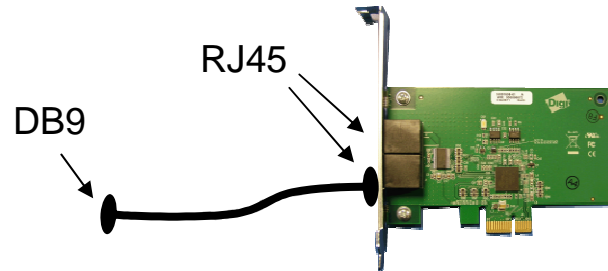
PCIe 2-Port Async EIA-232 Adapter

- 2-Port Async EIA-232 PCIe Adapter (#5289) Full high
- PCIe LP 2-Port Async EIA-232 Adapter (#5290) Low profile

- Used by IBM i 7.1 and AIX/Linux
- Introduced for AIX/Linux Oct 2011
- Choice of PCIe slot config options (full high or low profile)
- Supported on POWER7 servers (not POWER6)
- IBM i 7.1 TR5 support added Oct 2012
- IBM i 7.1 TR7 expands support Oct 2013
- Two RJ45 ports, both need modems

- **Protocols supported include:**
 - Async
 - Async PPP - with IBM i 7.1 TR7

- Fax/400 supported with appropriate external modem
- **CCIN = 57D4**



Cabling insight

- #5289/5290 has RJ45 connectors
- If your device needs a DB9 connector, use an RJ45-to-DB9 converter cable
- For convenience, one converter cable is included with #5289/5290 feature

PCIe Slots - High Level

	4U		2U	
	1S 4U	2S 4U	1S 2U	2S 2U
Total PCIe slots (all hot swap)	7	11	6	9
Required* LAN adapter (available for client use)	1	1	1	1
PCIe slots after required* LAN adapter	6	10	5	8
However if use high performance, expanded function backplane	-1	-1	-1	-1
PCIe slots after required* LAN and if using high performance backplane	5	9	4	7

* 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 Slots - More Detail -- x8 and x16

	4U			2U		
	S814	S824	S824	S812L	S822	S822L S822
	8286-41A	8286-42A	8286-42A	8247-21L	8284-22A	8247-22L 8284-22A
	1S	Only 1S in 2S box	2S	1S	Only 1S in 2S box	2S
Total PCIe slots	7	7	11	6	6	9
x16	2 x16	2 x16	4 x16	2 x16	2 x16	4 x16
x8	5 x8	5 x8	7 x8	4 x8	4 x8	5 x8
Required LAN adapter (available for client use)	1 x8	1 x8	1 x8	1 x8	1 x8	1 x8
PCIe slots after required LAN adapter	6	6	10	5	5	8
	2 x16	2 x16	4 x16	2 x16	2 x16	4 x16
	4 x8	4 x8	6 x8	3 x8	3 x8	4 x8
However if use high performance, expanded function backplane	-1 x8	-1 x8	-1 x8	-1 x8	-1 x8	-1 x8
PCIe slots after required LAN and if using high performance backplane	5	5	9	4	4	7
	2 x16	2 x16	4 x16	2 x16	2 x16	4 x16
	3 x8	3 x8	5 x8	2 x8	2 x8	3 x8

- PCIe slots are all Gen3 slots (Higher MHz used than Gen2 = 2x theoretical bandwidth)
- Some slots are x16 and some are x8. (x16 have 2x theoretical bandwidth)
- In 2014 can ignore which slots are x8 and which are x16 for most adapters ... see more detail

PCIe Slots - Way More Detail

Slot identifier	4U			2U		
	S814	S824	S824	S812L	S822	S822L S822
	8286-41A	8286-42A	8286-42A	8284-21L	8284-22A	8247-22L 8284-22A
	1S	Only 1S in 2S box	2S	1S	Only 1S in 2S box	2S
C2 x8	n/a	n/a	Y	n/a	n/a	Y
C3 x16	n/a	n/a	Y	n/a	n/a	Y
C4 x8	n/a	n/a	Y	n/a	n/a	n/a
C5 x16	n/a	n/a	Y	n/a	n/a	Y
C6 x16	Y	Y	Y	Y	Y	Y
C7 x16	Y	Y	Y	Y	Y	Y
C8 x8	Y	Y	Y	n/a	n/a	n/a
C9 x8	Y	Y	Y	Y ***	Y ***	Y ***
C10 x8	Y *	Y *	Y *	Y *	Y *	Y *
C11 x8	Y	Y	Y	Y	Y	Y
C12 x8	Y **	Y **	Y **	Y	Y	Y
Total PCIe slots	7	7	11	6	6	9

* C10 is where IBM manufacturing places the 4-port Ethernet card.

** Slot not usable by PCIe adapter with expanded function, dual-IOA backplane -- space taken by SAS ports & cabling

*** Slot not usable by PCIe adapter with expanded function, dual-IOA backplane – space taken by power protection for write cache

PCIe Slots

Slot identifier	4U			2U		
	S814	S824	S824	S812L	S822	S822L S822
	8286-41A	8286-42A	8286-42A	8247-21L	8284-22A	8247-22L 8284-22A
	1S	Only 1S in 2S box	2S	1S	Only 1S in 2S box	2S
C2 x8	n/a	n/a	direct	n/a	n/a	direct
C3 x16	n/a	n/a	direct	n/a	n/a	direct
C4 x8	n/a	n/a	direct	n/a	n/a	n/a
C5 x16	n/a	n/a	direct	n/a	n/a	direct
C6 x16	direct	direct	direct	direct	direct	direct
C7 x16	direct	direct	direct	direct	direct	direct
C8 x8	PEX	PEX	PEX	n/a	n/a	n/a
C9 x8	PEX	PEX	PEX	PEX ***	PEX ***	PEX ***
C10 x8	PEX *	PEX *	PEX *	PEX *	PEX *	PEX *
C11 x8	PEX	PEX	PEX	PEX	PEX	PEX
C12 x8	PEX **	PEX **	PEX **	PEX	PEX	PEX
Total PCIe slots	7	7	11	6	6	9

Direct --- straight to processor DCM, not shared bandwidth

PEX --- shared bandwidth - PLX PEX8748 Switch – two PEX sets per system indicated by boxes

Note, internal SAS controllers for system unit SAS bays & DVD attach to PEX – one controller per PEX

Comment – compared to POWER7+ 710-740 with this is much better. There are no 7xx direct slots and PCIe slots share bandwidth over GX++ and P7IOC I/O Controller Hub. POWER8 architecture has fewer PCI slots sharing and much bigger bandwidth to start with.

* C10 is where IBM manufacturing places the 4-port Ethernet card.

** Slot not usable by PCIe adapter with high function, dual-IOA backplane -- space taken by SAS ports & cabling

*** Slot not usable by PCIe adapter with high function, dual-IOA backplane – space taken by power protection for write cache

PCIe x16 and x8 Slot Usage Considerations

- **For most cards it doesn't matter** which slot used (read all chart)
 - All low profile slots in 2U box, all full-high slots in 4U box
 - All slots are Gen3
 - Bandwidth is so big on the Gen3 x8 there is no significant difference where the x16 will noticeably make a difference with existing adapters
 - No x16 cards announced as of Apr 2014
- **It does matter for these specific cards**
 - SOD future CAPI cards -- **must use the x16**
 - 2-port 40Gb Ethernet and IB cards Nice to put on x16 ... helps DMA capability (more TCEs)
 - These older adapters* only supported in the x16 slots
 - #5901/#5278(LP)/#EL10(LP) PCIe Dual-x4 SAS Adapter
 - #5287(LP)/#5288 PCIe2 2-port 10GbE SR Adapter

* Why these adapters: This is an idiosyncrasy of their then-existing standards interpretation/implementation as applied to today's more "evolved/reviced" PCIe standards

PCIe Special Placement Considerations



A few adapters have special placement needs

- IBM manufacturing will place an Ethernet adapter in C10 slot.
- These adapters need to be in specific slots: **Non PEX**
 - slots C6, C7, C11 on any 1-socket or 2-socket server or also in 2-socket servers slots C2, C3, C5 or also in 2S4U server C4 slot (nonPEX)
 - #2893, 2894 PCIe 2-Line WAN w/Modem Adapter
 - #5269, 5748 PCIe POWER GXT145 Graphics Accelerator
 - #EN13, EN14 PCIe Bisync Adapter
- **These adapters need to be in x16 PCIe slots**
 - #5901/#5278/EL10 PCIe Dual-x4 SAS Adapter
 - #5287(LP)/#5288 PCIe2 2-port 10GbE SR Adapter

POWER8 PCIe Adapter Migration



POWER7-supported PCIe adapters supported on POWER8

- Initial GA: robust list
- Modest set of PCIe adapters not supported on POWER8 Scale-out Servers -- *mostly elderly or low volume*

Remember, No PCI-X adapters on POWER8

- This means no SCSI
- For IBM i, means no IOPs (no SDLC)



Robust List PCIe Adapters Supported (1/2)

Ethernet NIC	4-port 1GbE RJ45	#5899, #5260	
Ethernet NIC	2-port 10GbE 10GBase-T RJ45	#EN0W, #EN0X, #EL3Z	NEW
Ethernet NIC & FCoE (CNA)	4-port 10GbE+1GbE SR+RJ45 (SR-IOV SOD)	#EN0H, #EN0J, #EL3B	
Ethernet NIC	4-port 10GbE+1GbE SR optical	#EN0S, #EN0T	NEW
Ethernet NIC	4-port 10GbE+1GbE Copper twinax	#EN0U, #EN0V	NEW
Ethernet NIC & RoCE	2-port 10GbE SR optical	#EC29, #EC30, #EL2Z	
Ethernet NIC	2-port 10GbE SR optical iSCSI TOE	#5744, #5280	
Ethernet NIC	2-port 10GbE SR optical	#5284, #5287, #EL2P	
Ethernet NIC	1-port 10GbE LR optical (IBM i native)	#5772	
Ethernet NIC & OpenOnload	2-port 10GbE Copper twinax	#EL39, #EC2J, #EC2G	
Ethernet NIC & RoCE	2-port 40GbE QSFP+	#EC3A, #EC3B	NEW
Fibre Channel	2-port 8Gb	#5735, #5273, #EL2N	
Fibre Channel	4-port 8Gb (FH)	#5729	
Fibre Channel	4-port 8Gb (LP)	#EN0Y	
Fibre Channel	2-port 16Gb	#EN0A, #EN0B	
Communications	2-port Async RS232	#5289, #5290	
Communications	1-port Bisync (IBM i)	#EN13, #EN14	
Communications	2-port Async/Bisync (support only) (IBM i)	#2893, #2894	



Robust List PCIe Adapters Supported (2/2)

SAS RAID	4-port no-cache PCIe3 for SSD/HDD	#EJ0J, #EJ0M, #EL3B	
SAS Tape/DVD	4-port tape/DVD PCIe3	#EJ10, #EJ11, #EL60	NEW
SAS RAID	4-port huge-cache PCIe3 for SSD/HDD	#EJ0L	
SAS RAID/Tape/DVD	2-port no-cache PCIe1 for HDD	#5901, #5278, #EL10	
Infiniband (IB)	2-port QDR IB SR optical	#5285, #5283	
Graphics	2D graphics for general use	#5748, #5269	
Graphics	3D graphics for RHEL7 (RHEL7 in beta)	#EC42, #EC41	New soon*
Encryption	Crypto Coprocessor 4765-001	#4807	
USB	4-port USB-3	#EC45, #EC46	NEW
Programmable FPGA	PCIe3 FPGA Accelerator	#EJ12, #EJ13	NEW

* IBM Confidential until announced. Will not announce in April.



PCIe Adapters NOT Supported: April 2014

Ethernet NIC	2-port 1GbE RJ45	#5767, #5281	Old, use 4-port
Ethernet NIC	2-port 1GbE SX optical	#5768, #5274	Old, use 4-port RJ45
Ethernet NIC	4-port 1Gbe RJ45	#5717, #5271	Old, use newer
Ethernet NIC	1-port 10GbE SR optical	#5769, #5275	old
Ethernet NIC & FCoE	2-port 10GbE SR optical	#5708, #5270	Use newer
Ethernet NIC & FCoE	4-port 10GbE+1GbE Copper Twinax+RJ45	#EN0K, #EN0L, #EL3C	Use SR optical
Ethernet NIC	2-port 10GbE copper twinax	#5288, #5286	old
Ethernet NIC	1-port 10GbE CX4	#5732, #5272	Old
Ethernet NIC	4-port 10GbE+1GbE Copper twinax+ RJ45	#5745, #5279	Old
Ethernet NIC	2-port 10GbE Copper twinax	#EL3A, #EC2K, #EC2H	Use OpenOnload
Ethernet NIC & RoCE	2-port 10GbE Copper twinax	#EC28, #EC27, #EL27	Use SR optical
Fibre Channel	1-port 4Gb	#5773	Old, use 8Gb
Fibre Channel	2-port 4Gb	#5744, #5276, #EL09	Old, use 8Gb
Communications	4-port Async RS232	#5785, #5277	Old, use 2-port
SAS RAID	2-port 380MB cache PCIe1 HDD/SSD	#5805, #5903	Use PCIe3
SAS RAID	3-port large cache PCIe2 HDD/SSD	#5913, #ESA3	Use PCIe3
SAS RAID	2-port zero cache PCIe2 SSD	#ESA1, #ESA2, #EL2K	Use PCIe3
USB	4-port USB-2	#2728,	Old, use USB-3
SSD on PCIe card	4-slot 1.8-inch SSD on SAS adapter	#2053, #2055, 2055	Use controller in storage backplane
Flash on card	Flash Adapter 90	#ES09	



PCIe Adapters NOT Supported: April 2014


Ethernet NIC	2-port 1GbE	RJ45	#5767, #5281	Old, use 4-port
Ethernet NIC	2-port 1GbE	SX optical	#5768, #5274	Old, use 4-port RJ45
Ethernet NIC	4-port 1Gbe	RJ45	#5717, #5271	Old, use newer
Ethernet NIC	1-port 10GbE	SR optical	#5769, #5275	old
Ethernet NIC & FC	2-port			
Ethernet NIC & FC	4-port			cal
Ethernet NIC	2-port			
Ethernet NIC	1-port			
Ethernet NIC	4-port			
Ethernet NIC	2-port			nload
Ethernet NIC & RoCE	2-port			cal
Fibre Channel	1-port			b
Fibre Channel	2-port			b
Communications	4-port			ort
SAS RAID	2-port			
SAS RAID	3-port			
SAS RAID	2-port	zero cache PCIe2	SSD	#ESA1, #ESA2, #EL2K
USB	2-port	USB-2	#2728,	Old, use USB-3
SSD on PCIe card	4-slot	1.8-inch	SSD on SAS adapter	#2053, #2055, 2055
Flash on card	Flash Adapter	90	#ES09	Use controller in storage backplane

Important comment – do not assume that many of the adapters shown on this list as of April 2014 will never be supported on the POWER8 servers.

If supported, do assume most or all of these would never be sold/ordered as new features on a POWER8 server, but supported if a client already had a card.

Do assume the ones with a black star will probably never be supported. (preliminary list)

Storage Backplanes



Storage Backplanes

2U server **example**
12 SAS bays base & split



Backplanes provide

- High performance integrated SAS controller(s) built on IBM industry leading PCIe Gen3 SAS adapter technology
 - All include RAID 0, 1, 5, 6, 10. Plus hot spare capability
 - Split backplane option with zero-write cache controllers
 - Easy Tier® function*
- 8-18** SAS bays for 2.5-inch (SFF) HDD or SSD
- 6-8 SAS bays for 1.8-inch SSD*
- One DVD bay
- Option for attaching one EXP24S drawer of HDD or SSD*

* With dual IOA, expanded function backplane with write cache

** number varies based on 2U/4U and backplane option selected

POWER8 2U Storage Backplane Options



Base: 1x SAS controller

- 12 SFF HDD/SSD Bays
- RAID 0,5,6,10 / JBoD
- No split backplane

Optional: 2x SAS Controller

- Split Disk Backplane = 6 + 6
- RAID 0,5,6,10 / JBoD



Optional: High Performance RAID 0,5,6,10,5T2,6T2,10T2

- 8 SFF Bays (HDD / SSD) and 6 SSD
- No split backplane

POWER8 2U Storage Backplane Options (for AIX/Linux)

• Must select one → #EJ0T #EJ0T+EJ0V #EJ0U *

12 SFF SAS bays 1 SAS controller No write cache DVD bay <i>Staged availability</i>	6+6 SFF SAS bays 2 SAS controllers No write cache DVD bay	8 SFF SAS bays Dual SAS controllers 7.2**GB cache DVD bay 6-bay SSD cage*** External SAS ports
--	--	---

AIX / IBM i / Linux	Yes		Yes		Yes
<i>Easy Tier Function</i>	no		no		Yes
JBOD	Yes		yes		no
RAID 0	Yes		yes		yes
RAID 10	Yes		Yes		yes
RAID 5/6	Yes		Yes		yes
Split backplane	no		Yes		no

* Uses one x8 PCIe slot (space taken up by cache protection hardware)

** 1.8GB physical write cache provides up to 7.2GB effectively **with compression**

*** 6-bay 1.8-inch SSD cage #EJTL --- required with #EJ0U on 8S22

POWER8 4U Front View

Base ... RAID 0,5,6,10 Feature - 12 SFF HDD/SSD disks
Split disk .. 6 + 6 feature (optional)



High Performance ... RAID 0,5,6,10,5T2,6T2,10T2 Feature (optional) + External SAS



POWER8 4U Storage Backplane Options

• Must select one → #EJON #EJON+EJOS #EJOP *

12 SFF SAS bays
1 SAS controller
No write cache
DVD bay

Staged
availability

6+6 SFF SAS bays
2 SAS controllers
No write cache
DVD bay

18 SFF SAS bays
Dual SAS controllers
7.2**GB cache
DVD bay
8-bay SSD cage***
External SAS ports

Note that no HH tape bay is provided – different than POWER7 720/740

AIX / IBM i / Linux	yes		yes		yes
<i>Easy Tier Function</i>	no		no		yes
JBOD	yes		yes		no
RAID 0	yes		yes		yes
RAID 10	yes		Yes		yes
RAID 5/6	Yes		Yes		yes
Split backplane	no		yes		no

* Uses one x8 PCIe slot (space taken up by #EJ0Z SAS ports/cabling – EJ0Z is mandatory with EJ0P)

** 1.8GB physical write cache provides up to 7.2GB effectively **with compression**

*** 8-bay 1.8-inch SSD cage #EJTM NOTE: Not available on mdl 41A, Required on mdl 42A with #EJ0P

POWER8 2U Storage Backplane Options (for Linux only)

Lower cost features for Linux only

- Must select one → #EL3T #EL3T+EL3V #EL3U *

12 SFF SAS bays
1 SAS controller
No write cache
DVD bay

Staged availability

6+6 SFF SAS bays
2 SAS controllers
No write cache
DVD bay

8 SFF SAS bays
Dual SAS controllers
7.2**GB cache
DVD bay
6-bay SSD cage***
External SAS ports

Easy Tier Function	no		no		Yes
JBOD	Yes		yes		no
RAID 0	Yes		yes		yes
RAID 10	Yes		Yes		yes
RAID 5/6	Yes		Yes		yes
Split backplane	no		Yes		no

* Uses one x8 PCIe slot (space taken up by cache protection hardware)

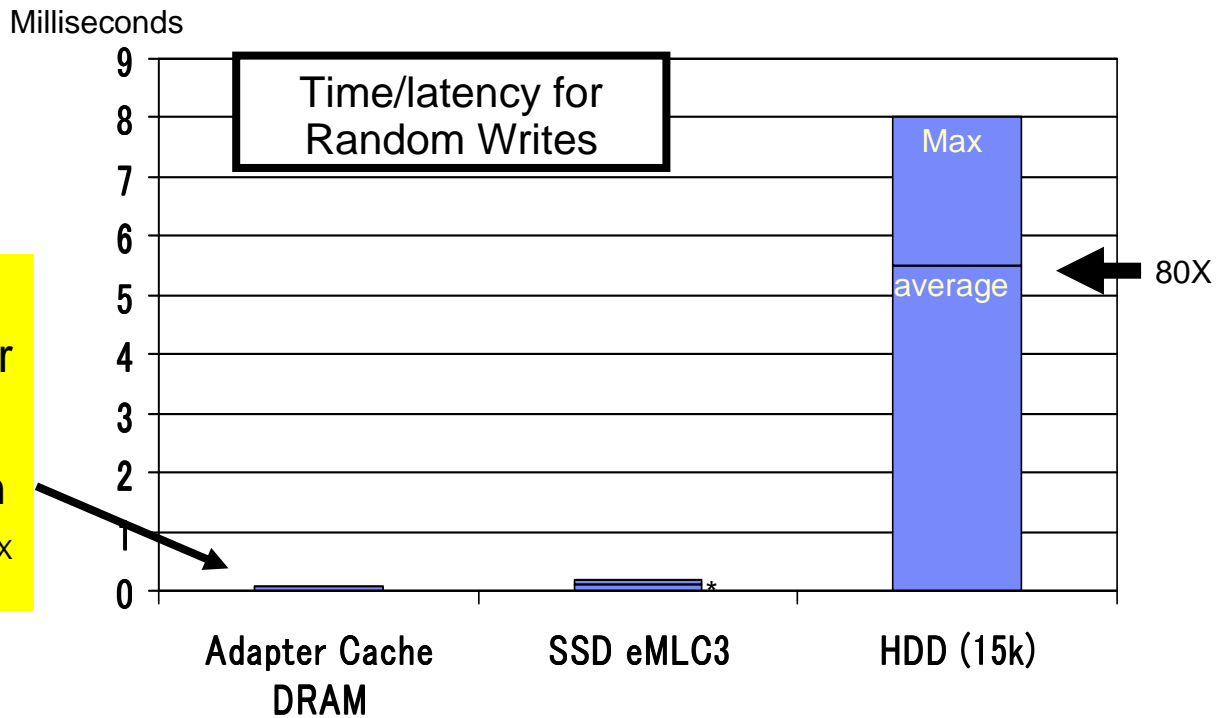
** 1.8GB physical write cache provides up to 7.2GB effectively with compression

*** 6-bay 1.8-inch SSD cage #EJTL -- required feature with #EL3U on 8S22L, but not supported on 8S12L

Controller Write Cache Value

Controller cache is

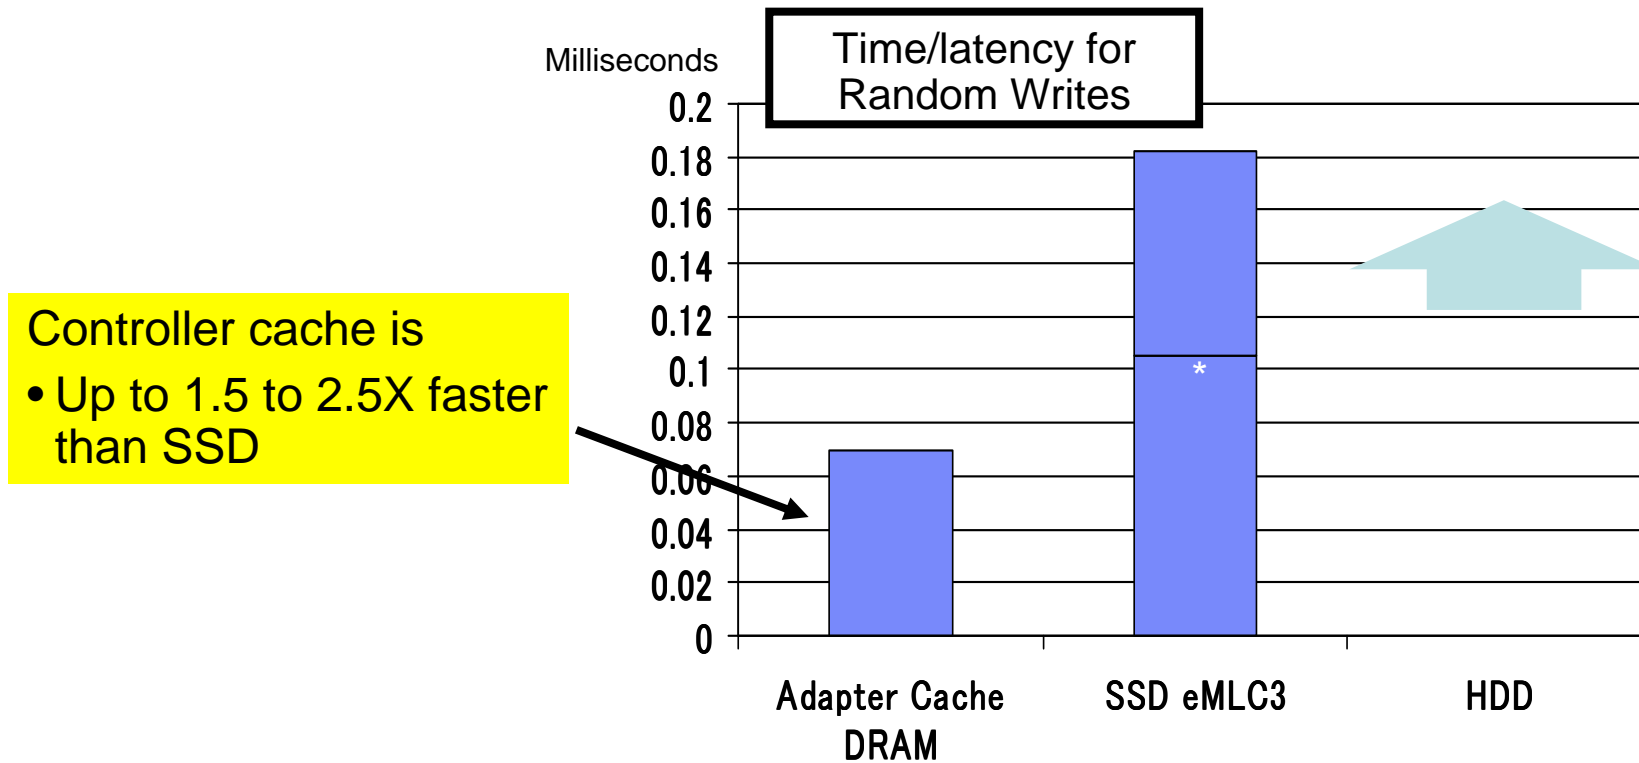
- Up to 1.5 to 2.5X faster than SSD
- Up to 80X faster than HDD average (Up to 115 X HDD max)



- Value depends on the amount/percentage of Writes
- Valuable for SSD, even more valuable for HDD
- Can even improve “reads” if application using recently written data still sitting in cache

Graph is a simplification. All performance discussions start with the words “it depends”. HDD 15k Max ms shows typical maximum rotational delay and arm movement. 10k HDD is about 1 ms slower. Non-random work will have better HDD measurements. Actual HDD performance varies from HDD to HDD. Adapter write cache can also speed reads, but value of write cache for reads is highly application dependent. The bottom line (*) of the SSD is obtained when the DRAM write cache integrated into SSD can handle the write and with a low queue depth. The higher SSD value is with a higher queue depth and/or when the SSD write cache is not able to keep up with a stream of writes and the write is occurring to the NAND flash memory.

Controller Write Cache Value (Other Scale)



Controller cache is

- Up to 1.5 to 2.5X faster than SSD

- Value depends on the amount/percentage of Writes
- Valuable for SSD, even more valuable for HDD
- Can even improve “reads” if application using recently written data still sitting in cache

Graph is a simplification. All performance discussions start with the words “it depends”. HDD 15k Max ms shows typical maximum rotational delay and arm movement. 10k HDD is about 1 ms slower. Non-random work will have better HDD measurements. Actual HDD performance varies from HDD to HDD. Adapter write cache can also speed reads, but value of write cache for reads is highly application dependent. The bottom line (*) of the SSD is obtained when the DRAM write cache is integrated into SSD can handle the write and with a low queue depth. The higher SSD value is with a higher queue depth and/or when the SSD write cache is not able to keep up with a stream of writes and the write is occurring to the NAND flash memory.

Backplane Performance Considerations



If you have applications with write-sensitive performance characteristics, be cautious about zero write cache configurations, especially for HDD

- Boot drives typically not a concern

Zero write cache controllers support the use on RAID-5 and RAID-6 on POWER8 servers

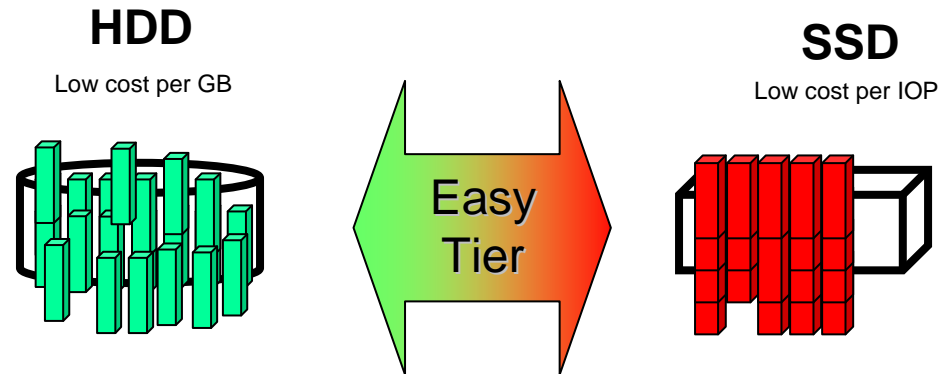
- Very cost effective per GB
- HOWEVER, increases the number of writes significantly

Note IBM i supports the use of all three backplane options

Be VERY, VERY careful of HDD workloads on base and split backplanes for IBM i.


Easy Tier Function Advantage

- Optional function with expanded-function, high-performance storage backplane



- Automatically moves high activity (hot) data to SSD and low activity (cold) data to HDD
- Function handled totally by POWER8 integrated SAS controllers. No application coding. No SAN, just internal SAS drives.
- For AIX/Linux/VIOS. Just configure as a new type RAID array
- IBM i essentially already has same kind of hot/cold function in OS for all SAS adapters, but could use this function via VIOS

Easy Tier w/ Power SAS RAID



Tiers of Hot and Cold data

Adapter dynamically swaps data between SSD/HDD tiers

- Based on Read / Write statistics
- Hottest data moved to SSD
- Cool/Cold data moved to SFF
- 1 MB or 2MB bands of data moved

Real-time updating (Seconds/minutes)


AIX and Linux Support

- **IBM i requires VIOS to use** (i alternatively has built in OS function)

Supports Boot and Hot Spares

Function also support with #EJ0L SAS adapter in POWER7+ and POWER8 server

Improves cost / performance tradeoff between HDDs/SSDs



Easy Tier® Function Configuration

- Function available with expanded-function, high-performance backplane. (#EJ0P, #EJ0U, #EL3U)
- For AIX/Linux/VIOS
 - AIX V7.1 TL3 SP3 or later or AIX V6.1 TL9 SP3 or later
 - VIOS 2.2.3.3 or later
 - Linux: RHEL 6.5 or later and SLES 11 SP3 or later
 - IBM i essentially already has same hot/cold function in OS. Could also use function via VIOS.
Not natively supported by IBM i.
- Configured as new RAID arrays:
 - RAID-5T2 (2-tiered RAID-5), RAID-6T2 (2-tiered RAID-6), RAID-10T2 (2-tiered RAID-10)
- SSD located in SSF-3 bays and/or 1.8-inch SSD bays and/or EXP24S drawer attached to the SAS ports
- Must combine HDD/SSD in the same array
 - HDD and SSD must be same block size ... all 5xx or all 4k (No 4k SSD announced, thus use 512/528 block size)
 - Can use different capacity HDD and SSD
 - If multiple HDD capacities (example 300GB and 600GB) then only 300 GB of the 600GB drives will be used
 - If multiple SSD capacities (example 387GB and 775GB) then only 387 GB of the 775GB drives will be used
 - No capacity ignored in HDD or SSD if only one size capacity in HDD and one size capacity in SSD, for example 387GB SSD and 600GB HDD.

SAS Bays for POWER8 Servers



SFF-3 used in 2U and 4U POWER8 servers

- 3rd generation carrier/tray on which 2.5-inch (SFF) is mounted
- Different tray/carrier than existing SFF-1 and SFF-2
 - SFF-1 in POWER6/POWER7 system unit or #5802 I/O drawer
 - SFF-2 in EXP24S Storage drawer (#5887 or #EL1S)
 - Can not put wrong tray/carrier in the wrong SAS bay
- For migrations to POWER8, conversions for SFF-1 to SFF-2 on POWER7 servers is planned. Protects client investment in HDD and SSD. Note is not a conversion to SFF-3, requires space in an EXP24S drawer to place converted SFF drives

SAS SFF-3 Options as of April 2014

SFF-3 HDD	Block size	Formatted with 512 or 4096 byte sectors	Linux only with 512 or 4096 byte sectors	Formatted with 528 or 4224 byte sectors
10k	5xx	300 GB #ESDR *	300 GB #ELDR	283 GB #ESDS *
10k	5xx	600 GB #ESD5	600 GB #ELD5	571 GB #ESD4
10k	5xx	1.2 TB #ESD9	1.2 TB #ELD9	1.1 TB #ESD8
15k	5xx	146 GB #ESDT *	146 GB #ELDT	139 GB #ESDU *
15k	5xx	300 GB #ESDB	300 GB #ELDB	283 GB #ESDA
15k	4k	300 GB #ESFB	300 GB #ELFB	283 GB #ESFA
15k	4k	600 GB #ESFF	600 GB #ELFF	571 GB #ESFE



* Staged eConfig support (27 May) and GA (25 July)


SFF-3 SSD	AIX/Linux/VIOS (528 byte sectors)	Linux only (528 byte sectors)	IBM i formatted (528 byte sectors)
eMLC Gen3	387 GB #ES0L	387 GB #EL14	387 GB #ES0M
eMLC Gen3	775 GB #ES0N	775 GB #EL13	775 GB #ES0P

Nota : The above features use SFF-3 or Gen3 carriers/trays and can be placed in the SFF SAS bays of the POWER8 servers

512/528 (5xx) byte sectors vs 4k byte sectors

- Same performance based on early results and predictions
- ??Same price?? **initial assumption assuming cost is the same**
- 4k is something the industry will be moving to over the next 5-10 years. It allows for bigger capacity disk drives.
- Can mix 4K and 5xx byte sector drives on the same SAS controller/adaptor, but can NOT mix in the same array.
- Suggestion from a pessimist... Don't mix 300GB 4k and 5xx block **SFF-3** on the same server. Chance of confusion in the future. "SFF-3" and "300GB" are key words in this statement.
- Same advice for EXP24S drawers... Don't mix 300GB 4k and 5xx block **SFF-2** on the same server.

SAS 1.8-Inch SSD for System Unit Cage



SSD

	AIX/Linux/VIOS (528 byte sectors)	Linux only (528 byte sectors)	IBM i formatted (528 byte sectors)
eMLC Gen3	387 GB #ES16	387 GB #ES46	387 GB #ES17

The above features placed in the SSD cage which is ordered with the dual IOA backplane feature on the 2U and the 2S4U server.

The cage features are:

For 2U: 6-bay cage #EJTL (Linux only #EL0H)

For 4U/S824: 8-bay cage #EJTM

Note: cage not supported on 2U 1-Socket or on 4U 1-Socket server

EXP24S Ports with Expanded Function Backplane

- Two SAS ports on rear of server with Expanded Function backplane
- Internal cabling connects ports to SAS expanders/controllers in the server. No separate feature code required to order internal cabling. 4U server uses #EJ0Z for communicating SAS ports to IBM Manufacturing. No feature code used for 2U server.
- Two YO cables (separately ordered) connect the two SAS ports to the EXP24S drawer.
- Drawer must be in mode 1
- Drawer can hold up to 24 SSD or up to 24 HDD, but SSD and HDD can not be mixed in the drawer
- Additional EXP24S drawers attached to SAS adapters in PCIe slots

18 SFF SAS bays
Dual SAS controllers
7.2**GB cache
RAID-0,1,5,6,10
DVD bay
8-bay SSD cage***
External SAS ports
Easy Tier function

Expanded
Function
Backplane

Max number of EXP24S drawers & disk/SSD

	4U		2U	
	1S 4U	2S 4U	1S 2U	2S 2U
Max EXP24S supported	14	14	14	14
Total SFF-2 disk drives supported in EXP24S	336	336	336	336
Total SFF-2 SSD drives supported in EXP24S	168	168	168	168

The above max calculated using SAS adapters which can attach up to 4 drawers per adapter (#EJ0J/#EJ0M/#EL3B) and then reducing it in consideration of limitation of SAS cables contained in cable management arm.

Max of 14 includes both EXP24S attached via PCIe adapters and/or SAS ports on the rear of the server with expanded function backplane

EXP24S ordered with #5887 for mdl 22A, 41A, 42A; ordered as #EL1S for mdl 22L

SAS SFF-2 Drives (for EXP24S Drawer)

SFF-2 or SAS drives on carriers/trays used in #5887 EXP24S drawer (Linux only #EL1S EXP24S drawer)

HDD SFF-2	Block Size	512-byte or 4096-byte sectors (JBOD)	512-byte or 4096-byte sectors (JBOD)	528-byte or 4224-byte sectors (RAID)	USA List price for model 720
		"AIX/Linux" formatted	"Linux" formatted Linux only	"IBM i" formatted	
10k	5xx	300 GB #1925	300 GB #EL1N	283 GB #1956	\$ 798
10k	5xx	600 GB #1964	600 GB #EL1Q	571 GB #1962	\$1,200
10k	5xx	900 GB #1752	900 GB #EL1R	856 GB #1738	\$1,500
10k	5xx	1.2 TB #ESD3	1.2 TB #ELD3	1.1 TB #ESD2	\$1,500
15k	5xx	146 GB #1917	146 GB #EL1M	139 GB #1947	\$ 498
15k	5xx	300 GB #1953	300 GB #EL1P	283 GB #1948	\$ 950
NEW 15k	4k	300 GB #ESFB	300 GB #ELFB	283 GB #ESFA	TBD
15k	4k	600 GB #ESFP	600 GB #ELFP	571 GB #ESFE	TBD

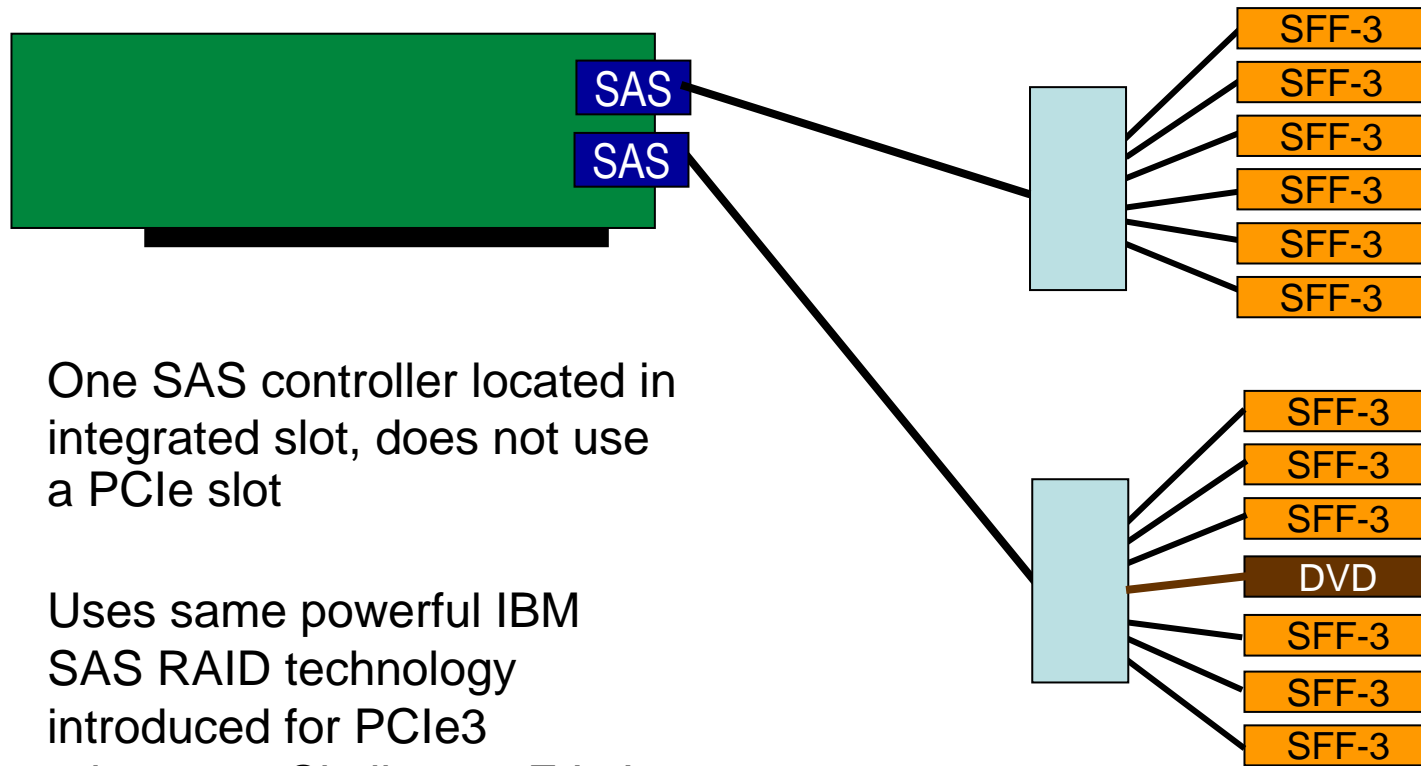
SSD SFF-2	Block Size	528-byte sectors	528-byte sectors	528-byte sectors	USA List price for model 720
		For AIX/Linux rules	For Linux rules	For "IBM i" rules	
eMLC2	5xx	387 GB #ES0C	387 GB #EL1L	387 GB #ES0D	
eMLC3	5xx	387 GB #ES19	387 GB #EL19	387 GB #ES1A	
eMLC3	5xx	775 GB #ES0G	775 GB #EL3G	775 GB #ES0H	
eMLC1	5xx	177 GB #1793	177 GB #EL1K	177 GB #1794	withdrawn

IBM USA suggested list prices as of Oct, 2013, for Power 720/740. Reseller prices may vary. Prices subject to change without notice.

10k vs 15k RPM Disk Drive Performance Insights

- KEY CAVEAT ... “it depends, your mileage will vary”. Use tools like WLE to help guide decisions.
- 10k drive around 20% fewer IOPS (I/O Operations Per Second) and/or 15k drive around 25% more IOPS. 10k is about 1 millisecond slower on average than 15k.
- Assuming adequate arms and the same protection (RAID vs mirroring); there is probably very modest performance impact of using 10k vs 15k arms in many scenarios. But there is an impact. How much impact is heavily application dependent. If this is an important business application with significant disk I/O performance dependency on a busy server, it leans the decision toward 15k. If the server is not that busy and disk I/O is not being stressed, then it leans the decision toward 10k.
- Mixing 10k and 15k drives: Supported even in the same array if the same capacity (assuming same block size (5xx or 4k)). Probably works fine for environments not being stressed. But avoid mixing if I/O workload is heavy and I/O performance is critical to the business. Keeping 10k arrays and 15k arrays separate on the system can allow different workloads to be assigned to each disk type and avoid this performance concern.
- Also consider mixing with SSD (not in same array). The mix of SSD plus 10k drives can be a better financial mix than SSD and 15k drives assuming hot data on SSD. In a SSD/HDD mix, the performance impact of 10k drives can be considerably less.

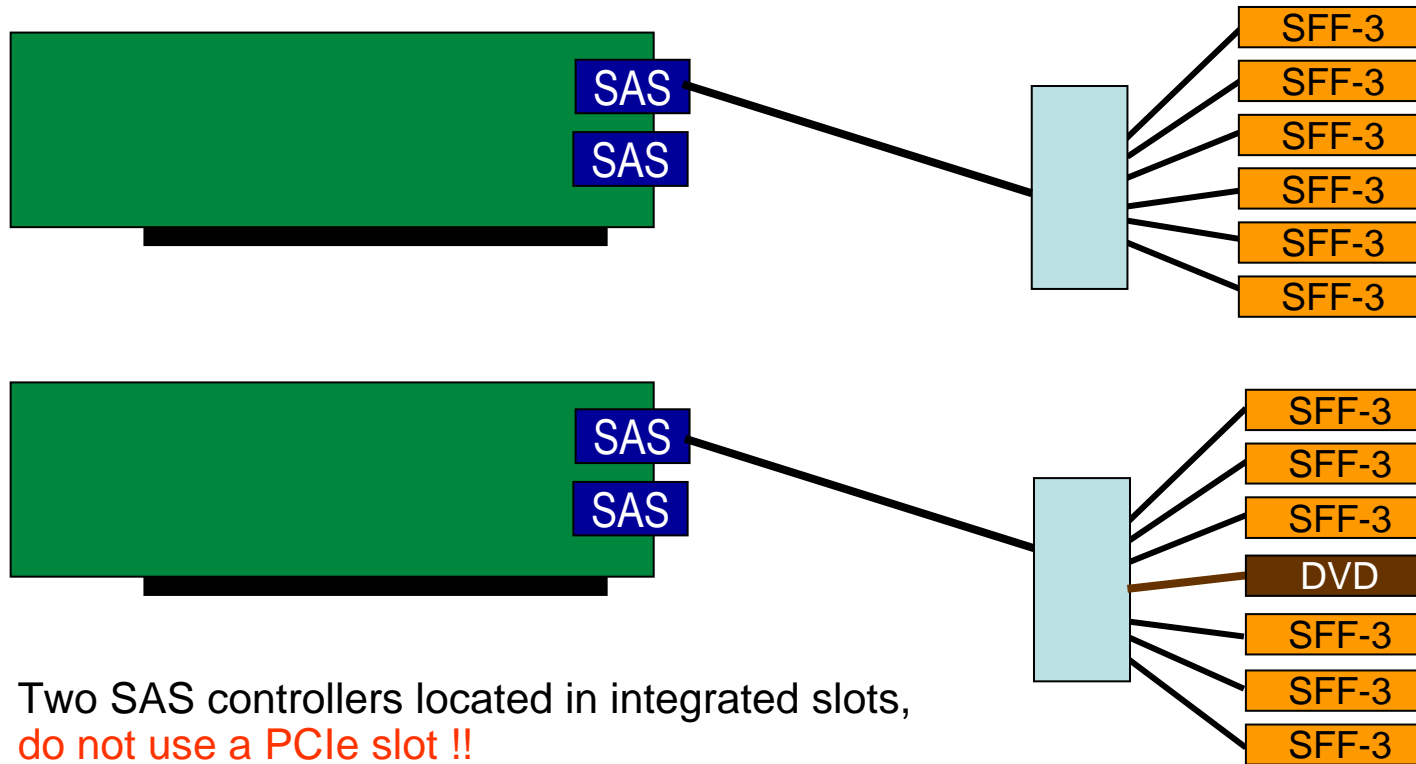
Base Configuration



One SAS controller located in integrated slot, does not use a PCIe slot

Uses same powerful IBM SAS RAID technology introduced for PCIe3 adapters. Similar to #EJ0J or #EJ0M or #EL3B SAS adapter.

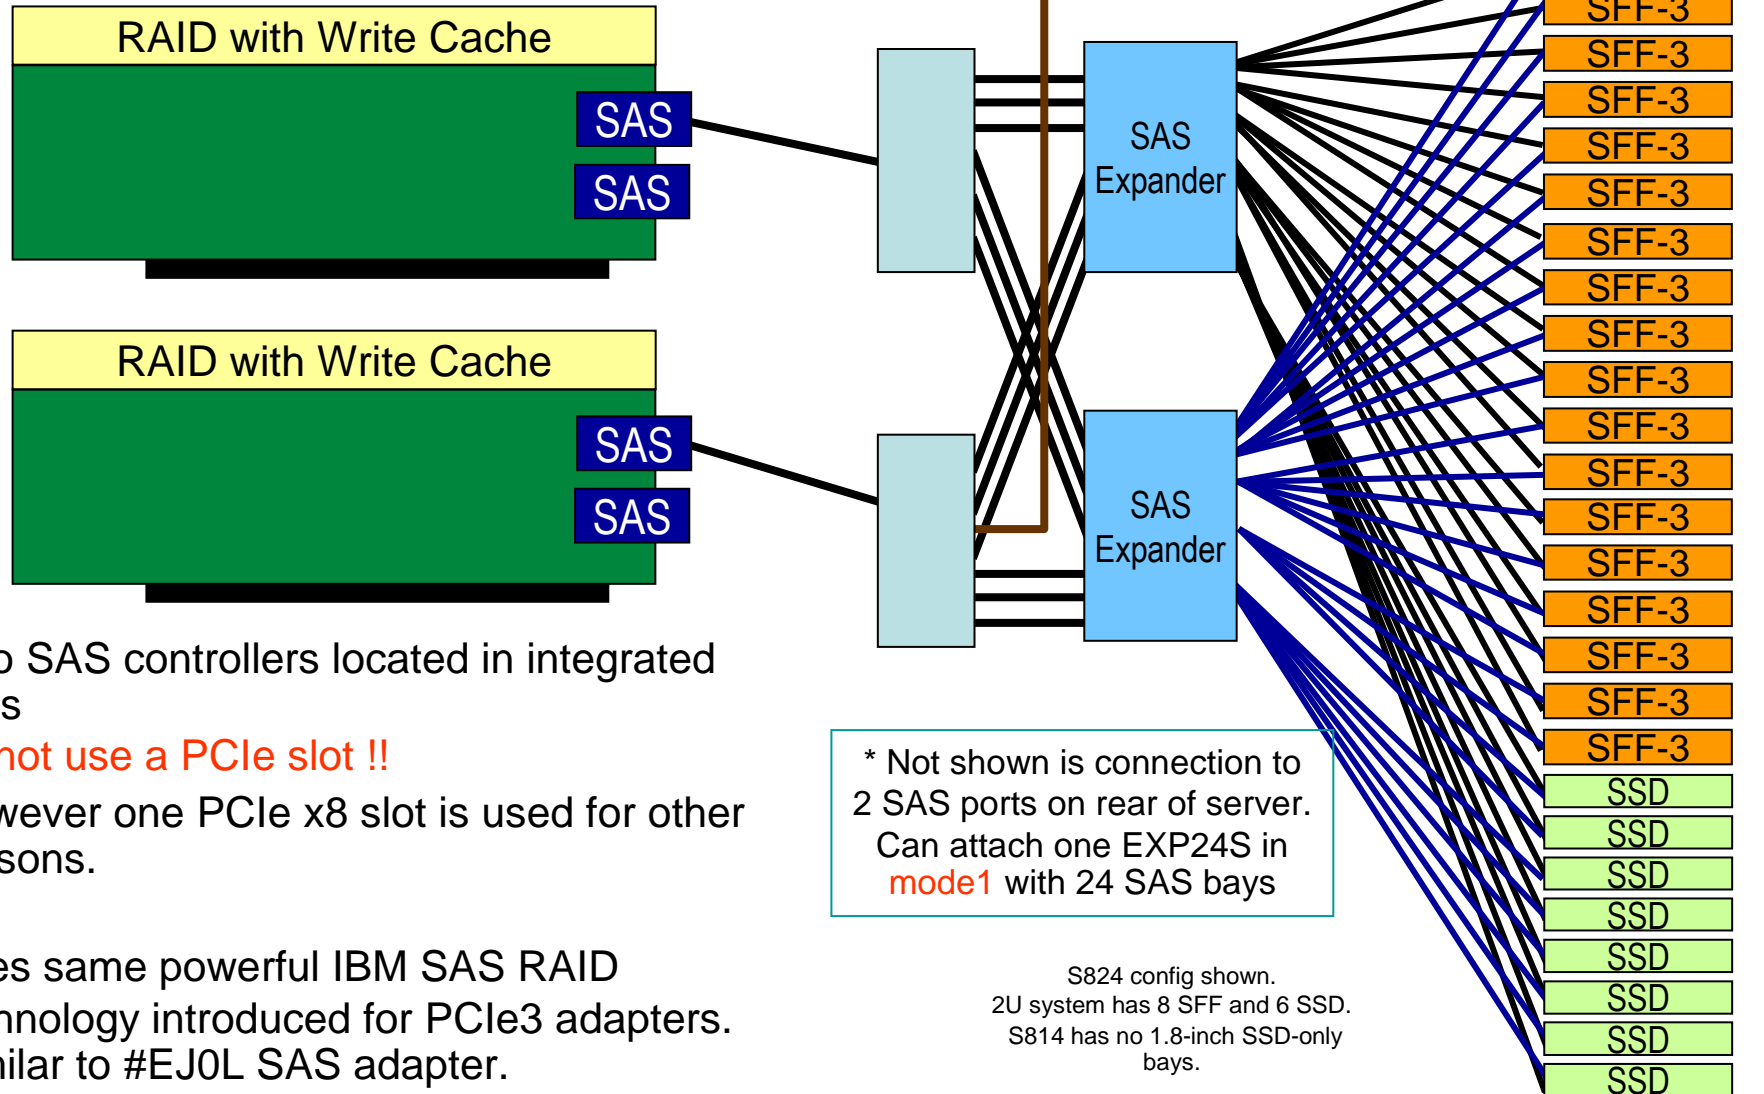
Split Disk Backplane Feature



Two SAS controllers located in integrated slots,
do not use a PCIe slot !!

Uses same powerful IBM SAS RAID technology introduced for PCIe3 adapters. Similar to two #EJ0J or #EJ0M or #EL3B SAS adapters.

High Performance RAID Feature




Two SAS controllers located in integrated slots
do not use a PCIe slot !!
 However one PCIe x8 slot is used for other reasons.

Uses same powerful IBM SAS RAID technology introduced for PCIe3 adapters. Similar to #EJ0L SAS adapter.

* Not shown is connection to 2 SAS ports on rear of server. Can attach one EXP24S in mode1 with 24 SAS bays

S824 config shown.
 2U system has 8 SFF and 6 SSD.
 S814 has no 1.8-inch SSD-only bays.

Integrated Ports




Integrated Ports on the System Unit

No port feature codes (except 4U SAS ports)	4U	2U
USB-3 with full client usage	2 front 2 rear	2 front, 2 rear
USB-2 primarily IBM use, limited client use supported*. Ports off Service Processor card	2* rear	2 rear
HMC ports - 1Gb Ethernet RJ45	2 rear	2 rear
Serial port - RJ45 **	2 rear	2 rear
2 SAS ports with high performance, dual IOA controller storage backplane – can attach one EXP24S drawer	2 rear <small>Mandatory #EJ0Z feature. Ports take a PCIe slot</small>	2 rear <small>No feat code. Ports do not take a PCIe slot</small>


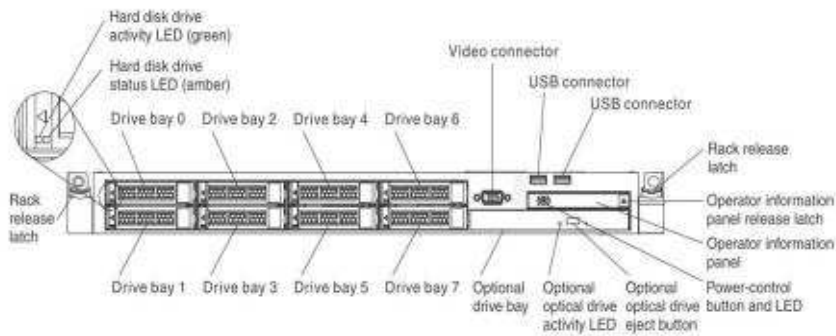
* For IBM i Unlike POWER7 or POWER6, IBM i communication to UPS is NOT via the serial port on POWER8 server. POWER8 connection to UPS for IBM i is over the USB-2 ports off the service processor card. A converter cable #ECCF should be used to convert USB port to 9-pin D-Shell connector.

** For AIX/Linux For attaching serial devices like a Async console which doesn't have a RJ45 interface. Just like on POWER7/POWER7+, use #3930 converter cable to provide a 9-pin D-Shell connector.

HMC



HMC Model 7042-CR8



- 7042-CR8 is System x model 7914 Ref3
 - Based on 7042-CR7
 - 2GHz Ivy Bridge Intel Proc*
 - 8GB DDR3 Mem (1x8GB DIMM)
 - Dual 500GB HDD (RAID1)
 - Quad GigE (First eth port can be shared IMM)
 - Dedicated IMM port
- UEFI, IMM, DSA is similar to CR7 (same for all 7914)
- Planned GA with HMC: 2Q 2014
- FW Level: 8R810

Enhanced HMC Management Capabilities

Virtual Environment Deployment

Create System Template that meets specific configuration requirements

§ Edit System Template & Capture Configuration as Template

Deploy Virtualization Configuration via “Deploy System from Template”

Create Partition Template that meets specific workload requirements

Deploy Partition

§ Create Partition from Template and Install Operating System

Virtual Environment Management

1. Configure Virtual Network via “Manage PowerVM”
2. View Virtual Storage Adapter Mappings via “Manage PowerVM”
3. Add storage to a partition via “Manage Partition: Virtual Storage”
4. Connect partition to a virtual network via “Manage Partition: Virtual Network”

Performance Monitoring

1. View System Performance Dashboard
2. View Resource Utilization Trends (Processor, Memory, Network, Storage)

New Server-level Tasks

Hardware Management Console Workplace (V7R7.6.0.0) - Mozilla Firefox: IBM Edition

https://.austin.ibm.com/hmc/connects/mainuiFrameset.jsp

Hardware Management Console

hscroot | Help | Logoff

Systems Management > Servers View: Table

Filter Tasks Views

Select	Name	Status	Available Process... Units	Available Memory (...)	Reference Code
<input type="checkbox"/>	...-sflp	Power Off	0	0	
<input checked="" type="checkbox"/>	...	Operating	0.3	21.875	
<input type="checkbox"/>	...	No Connection	0	0	Connecting 0000-0000-000000

Entered: 3 Selected: 0

- Properties
- Manage PowerVM
- Templates
- Operations
- Configuration
- Connections
- Hardware Information
- Updates
- Serviceability
- Capacity On Demand (CoD)
- Performance

Tasks: ilex.austi... C

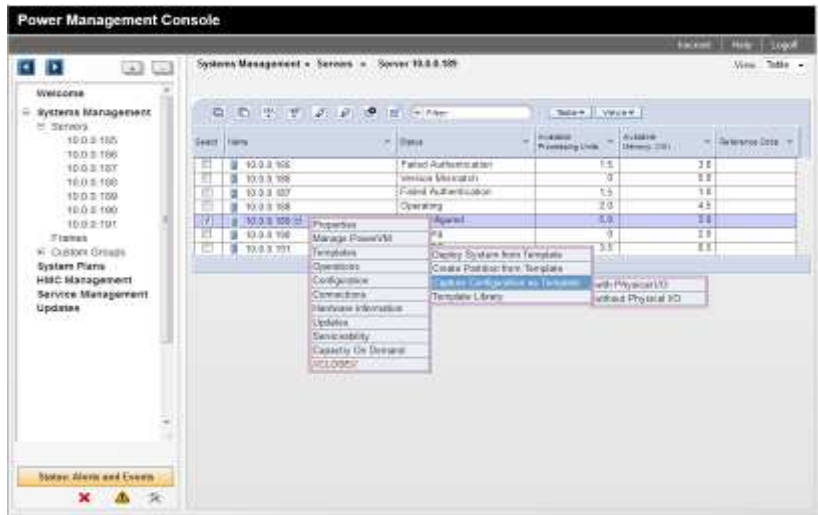
- Properties
- Operations
- Configuration
- Hardware Information
- Updates
- Serviceability
- Capacity On Demand (CoD)
- Performance

Status: Attention and Events

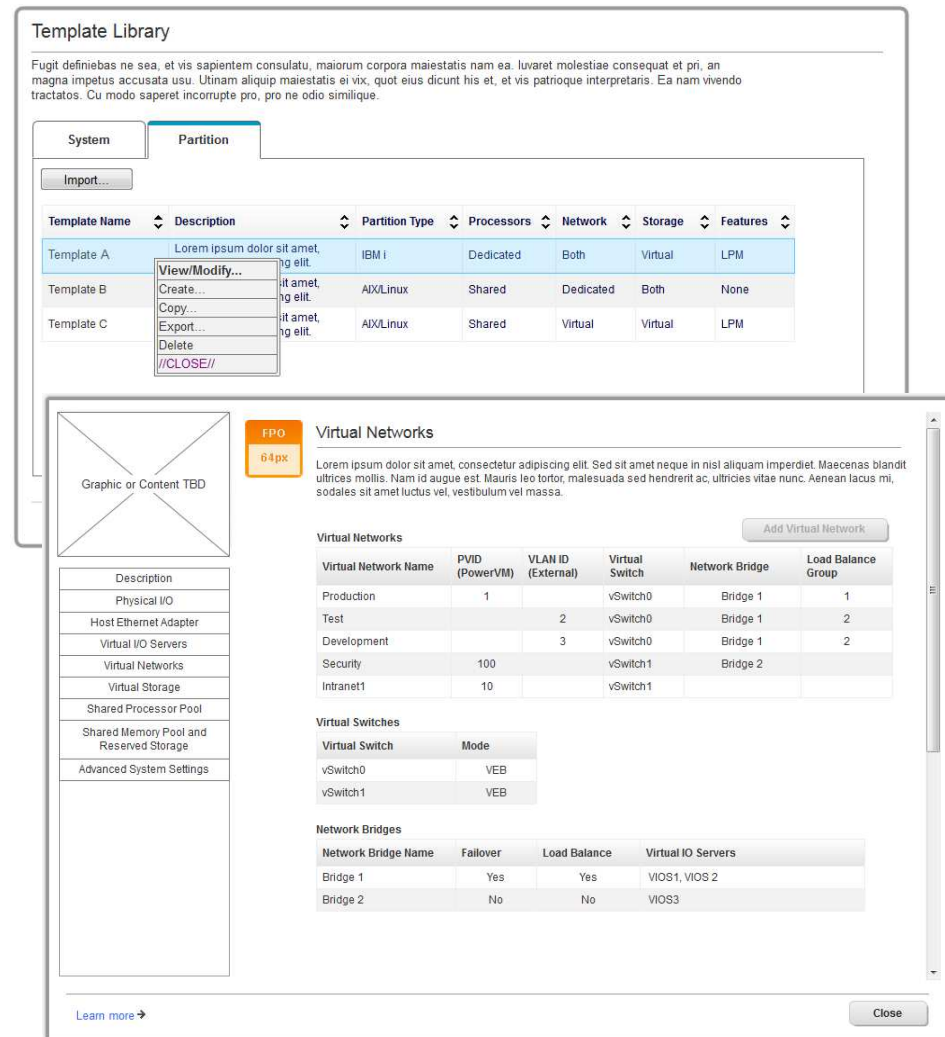
Transferring data from .austin.ibm.com...

System Configuration / Deployment – Using Templates

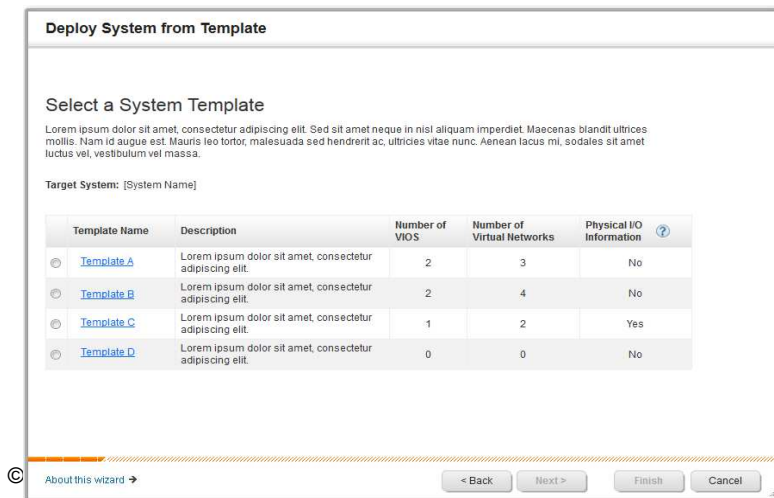
1) Capture Configuration as Template



2) View & Edit Templates

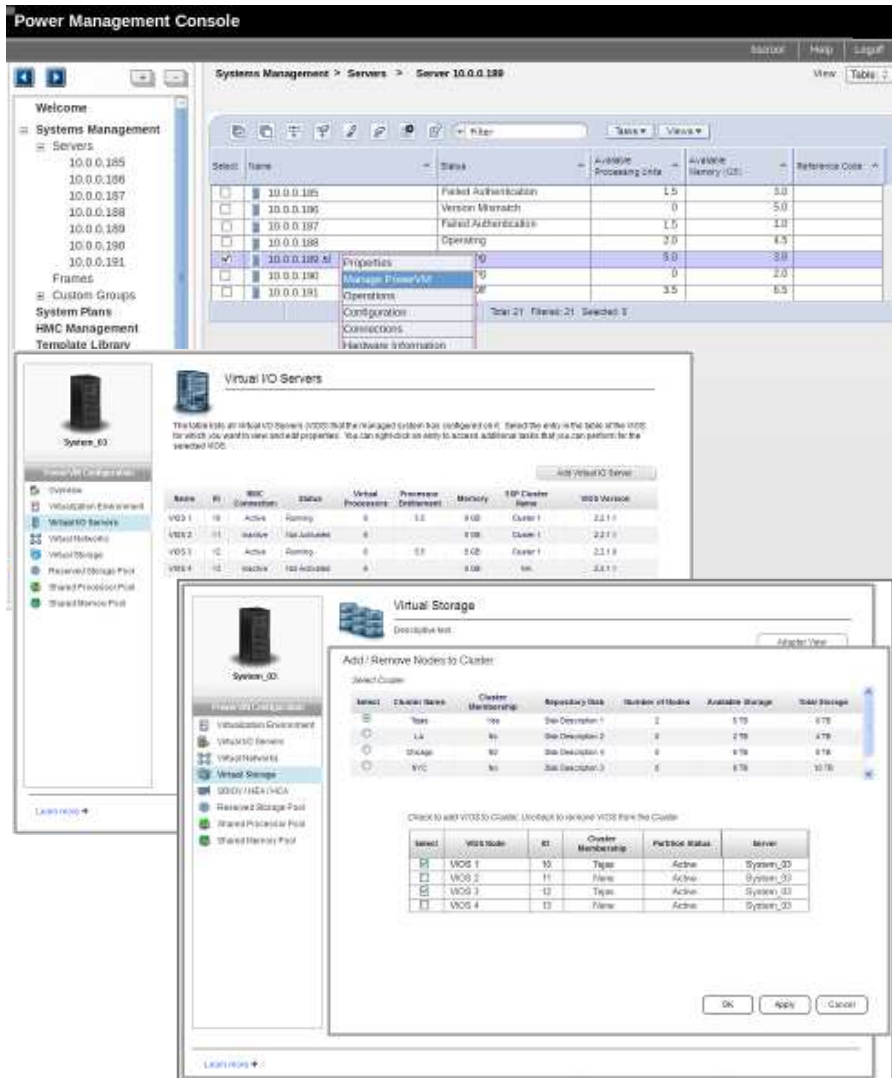


3) Deploy Template



PowerVM Management GUI

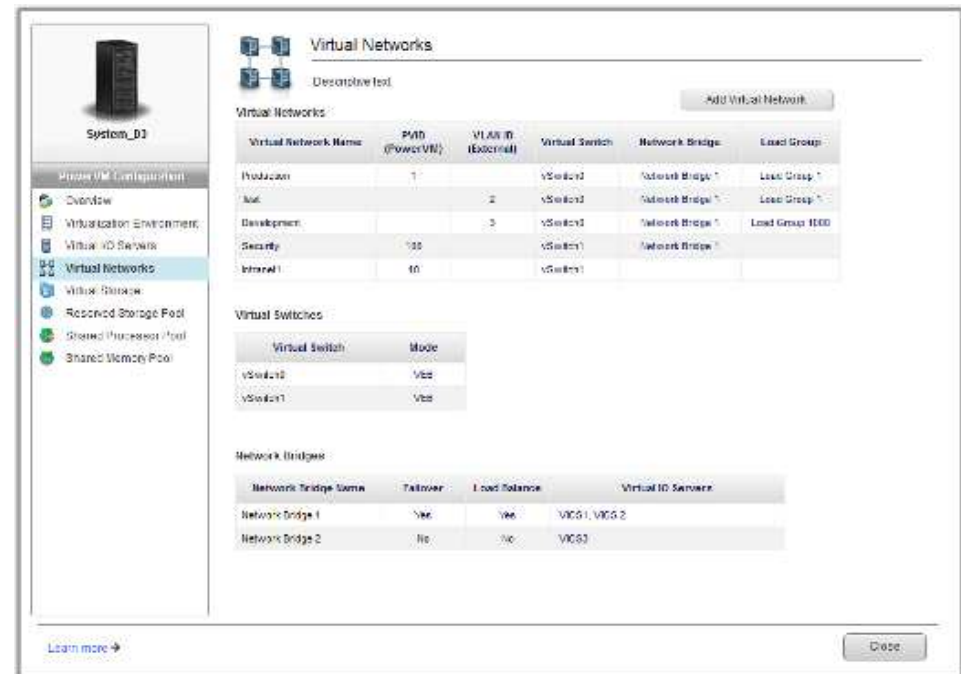
- Single launch point for all PowerVM configuration
- Management of all VIOS function
- Fast, scalable management of VIOS via a hardened HMC<->VIOS API
- Simplified virtual network & storage model
- Manage VIOS, Networks, and Storage



The screenshot displays the Power Management Console interface. The top section shows a navigation pane on the left with 'Systems Management' selected. The main area is titled 'Virtual IO Servers' and contains a table of VIOS instances. Below this, a 'Virtual Storage' dialog box is open, showing a table for adding or removing nodes to a cluster.

Name	IP	WCC Connection	Status	Virtual Processors	Preemptive Dispatches	Memory	BP Cluster Name	WCC Version
VIO1	10	Active	Running	8	12	8 GB	Cluster 1	2.2.1.1
VIO2	11	Inactive	Not Attached	8	12	8 GB	Cluster 1	2.2.1.1
VIO3	12	Active	Running	8	12	8 GB	Cluster 1	2.2.1.1
VIO4	13	Inactive	Not Attached	8	12	8 GB	NA	2.2.1.1

Select	Cluster Name	Cluster Membership	Responsible Disk	Number of Nodes	Available Storage	Total Storage
<input type="radio"/>	Test	NA	300 (Diskette) 1	2	1 TB	4 TB
<input type="radio"/>	LA	NA	300 (Diskette) 2	2	2 TB	4 TB
<input type="radio"/>	SHAGE	NA	300 (Diskette) 3	2	1 TB	4 TB
<input type="radio"/>	VVC	NA	300 (Diskette) 3	4	1 TB	16 TB



The screenshot displays the 'Virtual Networks' configuration page. It includes a table for Virtual Networks and a table for Network Bridges.

Virtual Network Name	PVID (PowerVM)	VLAN ID (External)	Virtual Switch	Network Bridge	Load Group
Production	1		vSw02/2	Network Bridge 1	Load Group 1
Test		2	vSw02/2	Network Bridge 1	Load Group 1
Development		3	vSw02/2	Network Bridge 1	Load Group 1000
Security	100		vSw05/1	Network Bridge 1	
Intranet 1	10		vSw05/1	Network Bridge 1	

Network Bridge Name	Failover	Load Balance	Virtual IO Servers
Network Bridge 1	Yes	Yes	VIO1, VIO2
Network Bridge 2	No	No	VIO3

Performance Monitoring – Metrics & Dashboard

Performance metric indicators & utilization dashboard

- Processor, memory & I/O

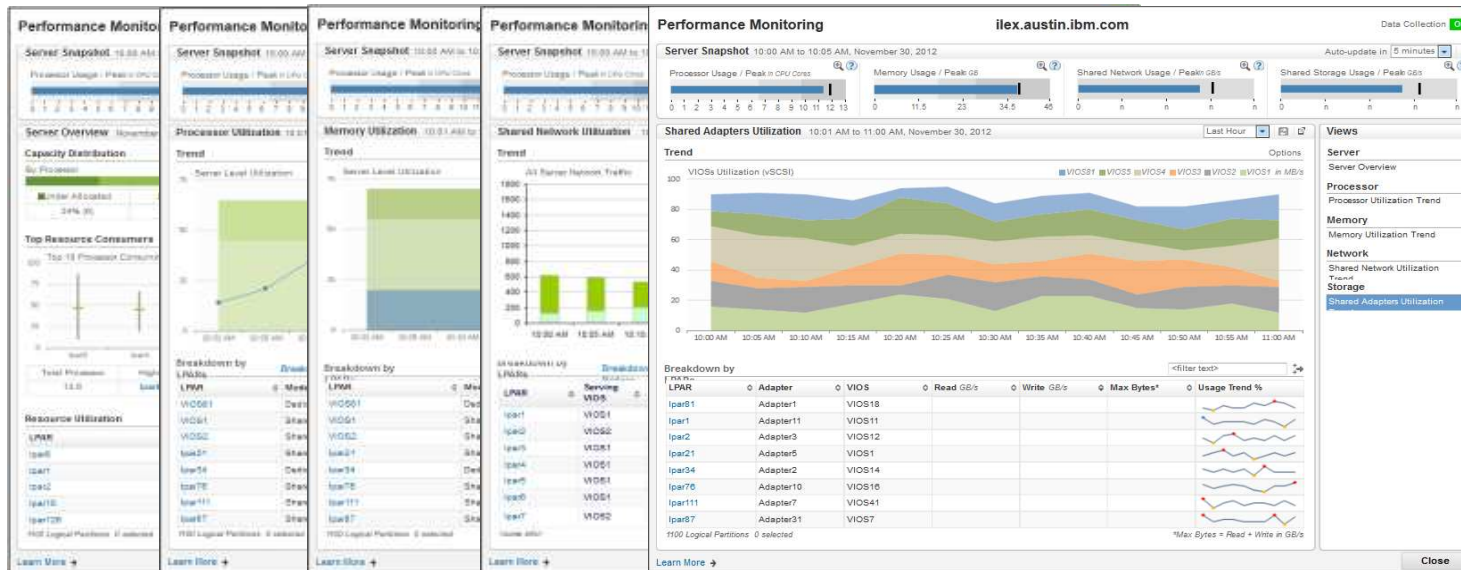
Server & LPAR level information

Basic trend data collection and visualization

- Identify bottlenecks
- Early problem detection

REST based API to access:

- All platform (PHYP & VIOS) metrics for Tivoli
- Third Party tools

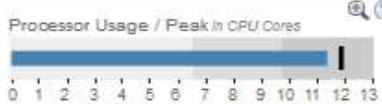


Performance Monitoring – Processor Utilization

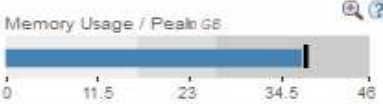
Performance Monitoring
ilex.austin.ibm.com
Data Collection On

Server Snapshot 10:00 AM to 10:05 AM, November 30, 2012 Auto-update in 5 minutes

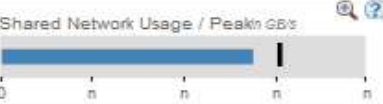
Processor Usage / Peak In CPU Cores



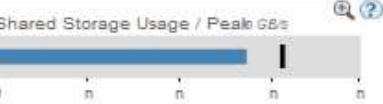
Memory Usage / Peak GB



Shared Network Usage / Peak In GB/s




Shared Storage Usage / Peak GB/s











Processor Utilization 10:01 AM to 11:00 AM, November 30, 2012 Last Hour

Trend Options

Server Level Utilization Total 65% Allocated 42% Overall Usage 37%



Breakdown by LPARs **Breakdown by Pools** <filter text>

LPAR	Mode	Pool	Entitled	Used	Max Usage	Usage Trend %
VIOS81	Dedicated			2	2	
VIOS1	Shared	Pool1		2	2	
VIOS2	Shared	Pool2		2	2	
Ipar21	Shared	Pool4		2	2	
Ipar34	Dedicated			2	2	
Ipar78	Shared	Pool3		2	2	
Ipar111	Shared	Pool4		2	2	
Ipar87	Shared	Pool0		2	2	

1100 Logical Partitions 0 selected All in CPU Cores


[Learn More](#) →
Close

Power Systems Management Requirements

Minimum Requirement:

HMC FW	V8.R1.0 or Later
HMC HW	7042-CR5 / 7042-C08 or Later
HMC Memory	2GB or More
Power HW	POWER6 and Later
PowerVM/VIOS	V2.2.3

Ordering.... **Processor Cards, Activations, Editions**



Processor DCM Feature Codes

	S814	S824	S822	S812L	S822L
	8286-41A	8286-42A	8284-22A	8247-21L	8247-22L
Number Sockets	1S	2S (Opt 1S)	2S (Opt 1S)	1S	2S
Processor DCM	#EPX0 6-core 3.02 GHz	One or Two #EPXE 6-core 3.89 GHz	One or Two #EPX1 6-core 3.89 GHz	One #ELPD 10-core 3.42 GHz	Two #ELPD 10-core 3.42 GHz
Processor DCM	#EPX6 8-core 3.72 GHz	One or Two #EPXF 8-core 4.15 GHz	One or Two #EPXD 10-core 3.42 GHz	One #ELP3 12-core 3.02 GHz	Two #ELP4 12-core 3.02 GHz
Processor DCM		Two #EPXH 12-core 3.52 GHz			

Processor Activation Features

	S814	S824	S822	S812L	S822L
	8286-41A	8286-42A	8284-22A	8247-21L	8247-22L
Processor DCM	#EPX0 6c 3.02 GHz	#EPXE 6c 3.89 GHz	#EPX1 6c 3.89 GHz	#ELPD 10c 3.42 GHz	#ELPD 10c 3.42 GHz
Charge activate	#EPY0	#EPYE	#EPY1	--	--
No-charge activate	#EPZ0	#EPZE	#EPZ1	#ELAD	#ELAD
Processor DCM	#EPX6 8c 3.72 GHz	#EPXF 8c 4.15 GHz	#EPXD 10c 3.42 GHz	#ELP3 12c 3.02 GHz	#ELP4 12c 3.02 GHz
Charge activate	#EPY6	#EPYF	#EPYD	--	--
No-charge activate	#EPZ6	#EPZF	#EPZD	#ELA3	#ELA4
Processor DCM		#EPXH 12c 3.52 GHz			
Charge activate		#EPYH			
No-charge activate		#EPZH			

100% of processor cores are activated.

As of April 2014, the only use of no-charge activations is with IBM i solution editions or with Linux-only models

#2319 factory deconfiguration is supported.

Editions/Vouchers/CBU/Factory Deconfiguration

- Unlike Power 710/720/730/740, no IBM editions, no IBM i editions, no AIX solution editions
 - Simpler – no mix of priced & no-charge activations
 - Simpler – no memory & I/O minimums to qualify for no-charge activations
 - ??vouchers TBD?? – similar to what used to be associated with above editions
- Like Power 720 IBM i Solution editions offered
 - Same structure/deal, same registration requirement, same edition feature codes (#4927), same website, same voucher
 - However unlike Power 740, not offered for Power S824
- CBU for IBM i yes ... basically same as 720/740, however the “primary” server must be a POWER7 or later.
- #2319 Factory deconfiguration -- same rules as Power 7xx

IBM i Solution Editions

For qualified sales of new servers with participating ISVs

Minimum size deal required

Registration/Approval of deal required prior to installation

#4927

6-core or 8-core Power 8S14

- 5 no charge processor core activations
- IBM i users @ only \$70
- 1 IBM Service voucher
- Unlimited users for IBM i Access Family
- Reduced price development tools

IBM i Solution Editions ISVs



Worldwide ISVs

IBM PowerHA SystemMirror for i
Coglin Mill
FIS
Help/Systems
IBS
Infor
Lawson
Maxava

menten
MISys
Oracle JD Edwards
Rocket
SAP AG
SugarCRM
Trader's
Vision Solutions

See Solution Edition Web page for latest listing
www.ibm.com/systems/power/hardware/editions/solutions.html

IBM i Solution Edition Service vouchers

- **In most countries, service voucher may be redeemed for:**
- **Migration Assistance**
An IBM representative or Business Partner will assist you in migrating from older IBM hardware to new IBM hardware by providing 1 day (8 hours) of service.
- **SAP on IBM i Installation Jump-start**
Provides for up to a day and a half of on-site installation and skills transfer services from one of our certified SAP on IBM i Basis experts.
- **Oracle JD Edwards EnterpriseOne Rapid Installation**
 - Provides optimized installation services, using rapid installation tools, to help get you up and running quickly.
- **Along with many other available services**
 - <http://www.ibm.com/systems/power/hardware/vouchers/index.html>

DVD Reminder

- DVD support is like the Power 710/720/730/740
- A SAS/SATA controller runs both DVD and HDD/SSD
- This has sharing and virtualization considerations

- This is not like POWER7+ 750/760/770/780 with a separate controller

Tape Reminder

- **No Tape available in the CEC !!!**

IBM i Solution Editions ISVs - Europe

<p>Akra Informatica Coop V AMAN Ametras Informatik AG ANDEP, S.L ANTARA Appligen ARCADE CONSULTORES, S.L ASERSA ATLAS Atos Avenue Software BALANKU BATE BITON Coheris Comarch COMTEK D. R. MIRKAM DEISTER Desadi10 Divalto DPI DRP Software EASI ECINSA EKAMART Eugenio Branco EURO-SYS EXECUTECH</p>	<p>FMR Fusion Cia Informática Gabinete AK Gabinete de SW Profesional Generix Gestion 400 GIAS Gilsar Grupo QS Halcyon Software HARDIS I2S IMTECH ICT In Concept Inforges Iris J2C Jiway LOGINAR GIAM - GROUPE DL NEGOCE Lusodata MAGIC SOFTWARE ENTERPRISES MONITIN Net@Work NSI, S.L. ORDIROPE OXAION AG Pact Sotraig Pantheon Automatisering</p>	<p>PinkRocade PORTEXICTOS QSI Registro REMICHSUJ Resolution Informatique Retail Assist Limited ReWa Warenwirtschaft RPO Automatisering BV Serig SERVICIOS INFORMATICOS SISTEMAS AUTOMATICOS SL Sistemas de Datos, S. L SN Axilog SOFT CATALUNYA SOFTGASA INFORMATICA SL SOPRA Tango/04 Computing Group Trentisa UNIT4 CODA</p>	<p>Italian Smart Solutions (720 4-core only) ACG AGOMIR S.p.A. C.A.T.A. INFORMATICA SRL COSMIC BLUE TEAM SPA DATASYS SRL DELTA SYSTEM SPA Divalto EDM SRL GEI SOFT SRL GL Italia Srl INFOMAX GROUP ITALIA Srl MULTI CONSULT MILANO S.p.a. QSL ITALIA srl Raz Lee Security (2001) LTD S.I. Soluzioni Informatiche srl S.INFO srl Sanmarco Informatica S.p.A. Sinapsi Informatica S.r.l. SIRIO INFORMATICA E SISTEMI SPA. SME UP S.r.l. SORMA SpA VAR GROUP SPA Vision Solutions WSS ITALIA</p>
--	---	--	--

See Solution Edition Web page for latest listing
www.ibm.com/systems/power/hardware/editions/solutions.html

IBM i Solution Editions ISVs

Worldwide ISVs

IBM PowerHA SystemMirror for i
Coglin Mill
FIS
Help/Systems
IBS
Infor
Lawson
Maxava
menten
MISys
Oracle JD Edwards
Rocket
SAP AG
SugarCRM
Trader's
Vision Solutions

Plus European ISVS and China

See Solution Edition Web page for latest listing
www.ibm.com/systems/power/hardware/editions/solutions.html

USA, Canada & Latin America ISVs

Agilisys
ANDEP, S. L.
BCS Group
Bellamy Software
CAPP Associates
Computer Guidance Corp.
Distribution Service Technologies, Inc (DST)
DMS Equipment Corporation
DPS, Inc (Data Processing Services, Inc)
Educational Data Resources, LLC
Epicor Software
Fiserv
Friedman Corporation
Gilsar
HarrisData (Harris Business Group)
Healthcare Management Systems
IS2
Jack Henry & Associates
JDA Software
Lavender and Wyatt Systems (LWSI)
Mincron
N2N Global (formerly KPG)
New Generation Software
PFW Systems
Retail Assist Limited
Retalix SA
Rippe & Kingston
Serti Dealership System (SDS)
Service Information Access, Inc (SIA)
Software Concepts Inc
Syntax Distribution Mgmt (SDM)
Tango/04 Computing Group
TMW Systems
UNIT4 CODA
United Computer Group, Inc
VAI
Xperia Solutions

How to Order Solution Edition

- Available via econfig – no unique channel terms
 - Select required solution edition:
 - Power 8S14 #4971/#4975????,**
 - Power 8S24 #4972????**
 - Discounted features automatically added to order
- All editions must be registered via web form
- Requirements for ordering a solution edition:
 - Must be with a authorized ISV
 - The server must be ordered with a Solution Edition feature code
 - The server must be a new purchase, not an MES upgrade
 - The offering must include new and/or upgrade software licenses and/or software maintenance from the ISV for the qualifying IBM server. Services and/or training for the qualifying server can also be provided.
 - The combined value (software, maintenance, services and training) of the ISV purchase must meet specific minimums which varies by model
 - Proof of purchase of the solution with a participating ISV must be provided to IBM on request. The proof must be dated within 90 days before or after the date of order of the qualifying server

Registration web form

<http://www.ibm.com/systems/power/hardware/editions/solutions.html>



Registering a Solution Edition

The fields indicated with an asterisk (*) are required to complete this transaction, other fields are optional. If you do not want to provide us with the required information, please use the "Back" button on your browser to return to the previous page, or close the window or browser session that is displaying this page.

Submitter information

Name*

E-mail*

Country*

Phone number* (including country code)

Fax number* (including country code)

Customer information

Customer Name*

IBM Customer Number*

Hardware information

Server model*

Serial #* (or Plant Order # if on order)

Software information

ISV*

4-Core Power S814 *Not in April Announce*

- The Power S814 offered in 6-core or 8-core
 - Like Power 720 6-core/8-core uses IBM i P10 software tier
- The Power 720 4-core offers a lower price/cost entry config
 - 4-core config has lower I/O and memory maximums
 - It uses IBM i P05 software tier
- No 4-core S814 announced in April, nor is an SOD published
- However, FYI: (IBM Confidential)
 - Plan* to Annc/GA a 4-core offering 4Q 2014*
 - Investigating an RPQ prior to 4Q
- **NOTE: The Planned 4-core Power 8S14 WON'T reduce the hardware pricing compared to a 6-core 8S14**



Merci !!!



ibm.com/systems/fr/power/