April 14, 2014

POWER8 Hardware Technical

Joe Armstrong & Mickey Sparks

With credit for charts to
Pat O'Rourke & Mark Olson

#powersystems, #powerlinux, #bigdata, #IBMWatson, #OpenPOWER

Agenda

- P8 Technology
- New Servers
- Performance
- PCIe detail
- Storage
- Memory
- OS Levels
- I/O

Power Processor Technology Roadmap

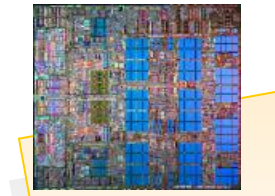

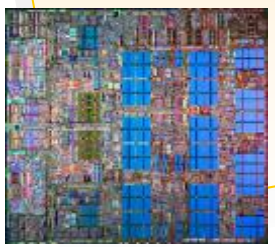

POWER5/5+ 130/90 nm

- ✓Dual Core
- ✓Enhanced Scaling
- ✓SMT
- ✓Distributed Switch +
- ✓Core Parallelism +
- ✓FP Performance +
- ✓Memory Bandwidth +
- ✓Virtualization

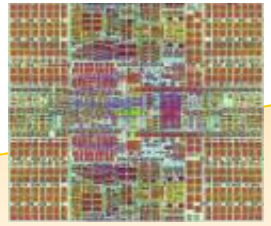

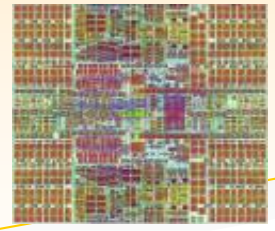

POWER6/6+ 65/65 nm

- ✓Dual Core
- ✓High Frequencies
- ✓Virtualization +
- ✓Memory Subsystem +
- ✓Altivec
- ✓Instruction Retry
- ✓Dynamic Energy Mgmt
- ✓SMT +
- ✓Protection Keys

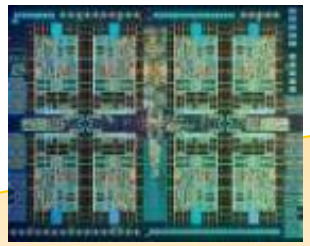

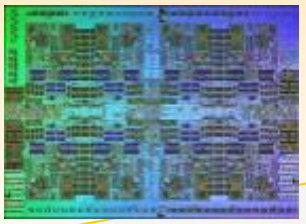

POWER7/7+ 45/32 nm

- ✓Eight Cores
- ✓On-Chip eDRAM
- ✓Power-Optimized Cores
- ✓Memory Subsystem ++
- ✓SMT++
- ✓Reliability +
- ✓VSM & VSX
- ✓Protection Keys+

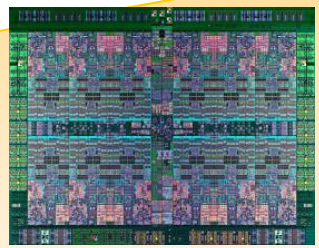

POWER8

- ✓More Cores
- ✓SMT+++
- ✓Reliability ++
- ✓FPGA Support
- ✓Transactional Memory
- ✓PCIe Acceleration

✓ 200+ systems in test

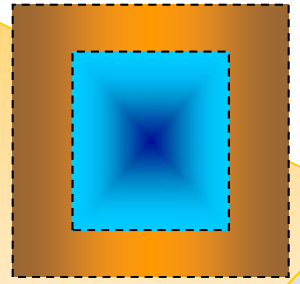

POWER9

- ✓Extreme Analytics Optimization
- ✓Extreme Big Data Optimization
- ✓On-chip accelerators

2004 2007 2010 2014

POWER8 Processor

POWER8 Processor

Technology
22nm SOI, eDRAM, 15 ML 650mm2

Cores
- 12 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

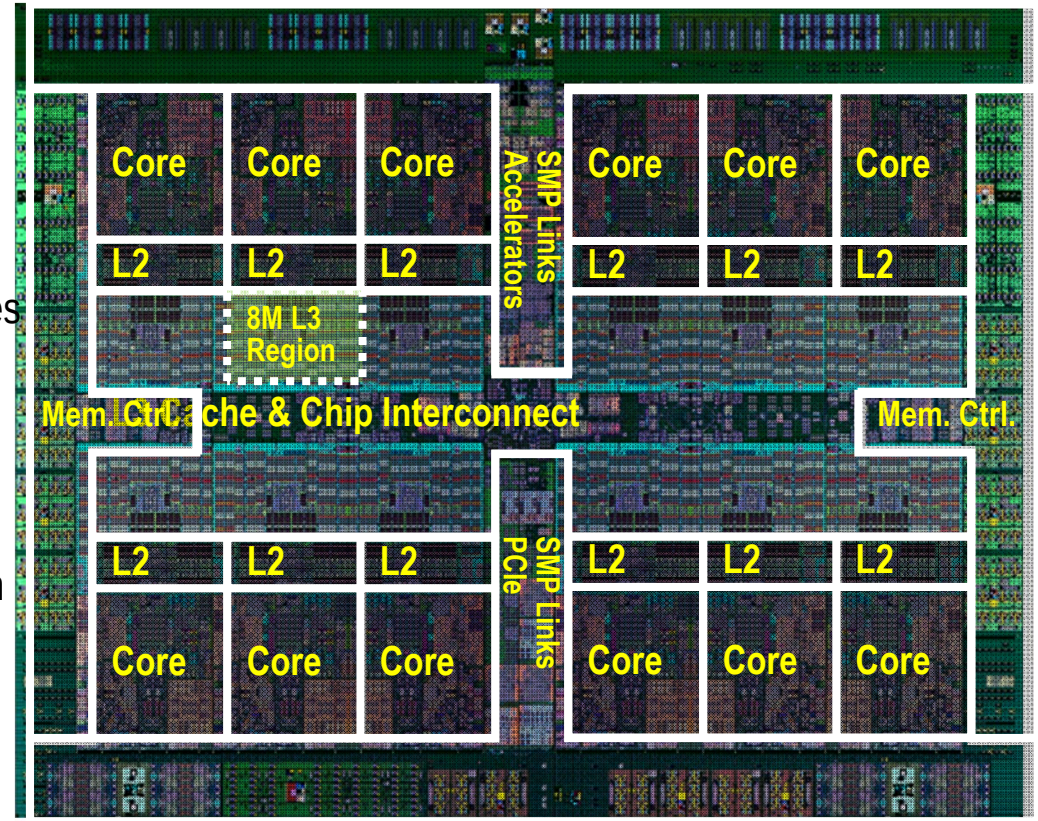

Caches
- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)

Memory
- Up to 230 GB/s sustained bandwidth

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

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

POWER8 Core

Execution Improvement vs. POWER7

- SMT4 → SMT8
- 8 dispatch
- 10 issue
- 16 execution pipes:
  - 2 FXU, 2 LSU, 2 LU, 4 FPU, 2 VMX, 1 Crypto, 1 DFU, 1 CR, 1 BR
- Larger Issue queues (4 x 16-entry)
- Larger global completion, Load/Store reorder
- Improved branch prediction
- Improved unaligned storage access

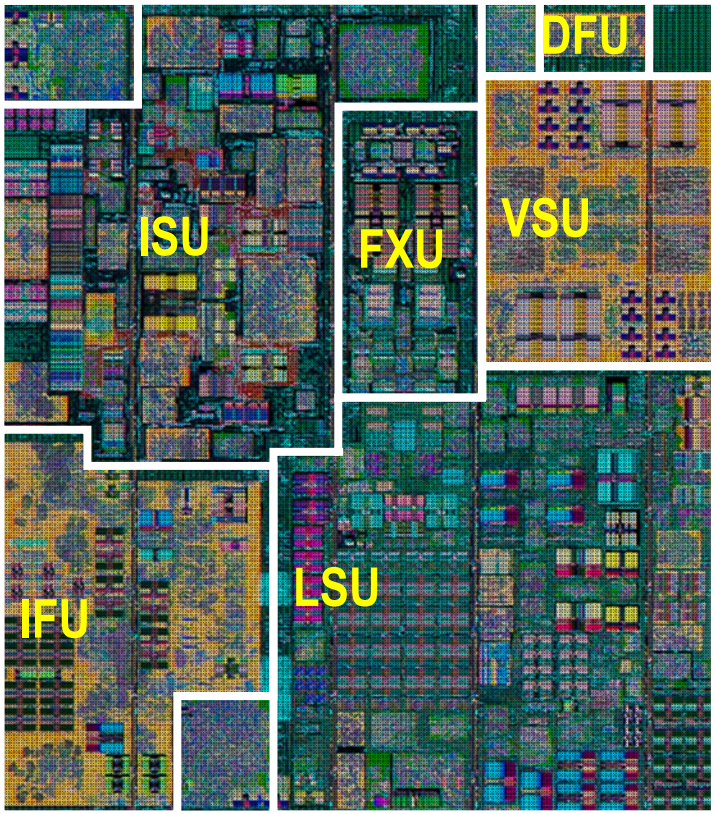

Larger Caching Structures vs. POWER7

- 2x L1 data cache (64 KB)
- 2x outstanding data cache misses
- 4x translation Cache

Wider Load/Store

- 32B → 64B L2 to L1 data bus
- 2x data cache to execution dataflow

Enhanced Prefetch

- Instruction speculation awareness
- Data prefetch depth awareness
- Adaptive bandwidth awareness
- Topology awareness

Core Performance vs . POWER7

~1.6x Thread

~2x Max SMT

Scale Out Systems - DCMs and POWER8 Chips

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

- 1 DCM fills 1 socket .... Similar to POWER7+ 750 / 760
- 1 DCM has two Scale Out POWER8 chips
- 1 DCM can provide 6-core, 8-core, 10-core or 12-core sockets

6-core Processor Chip
- 362 mm2
- 22nm SOI w/ eDRAM

Strengthen Cores
- 8 Threads per Core

Caches
- D Cache: 64KB
- L2: 512KB
- L3: 8 MB per Region Total: 48MB

Fine Grained Power Management
- On Chip power management

Excellent I/O bandwidth per socket

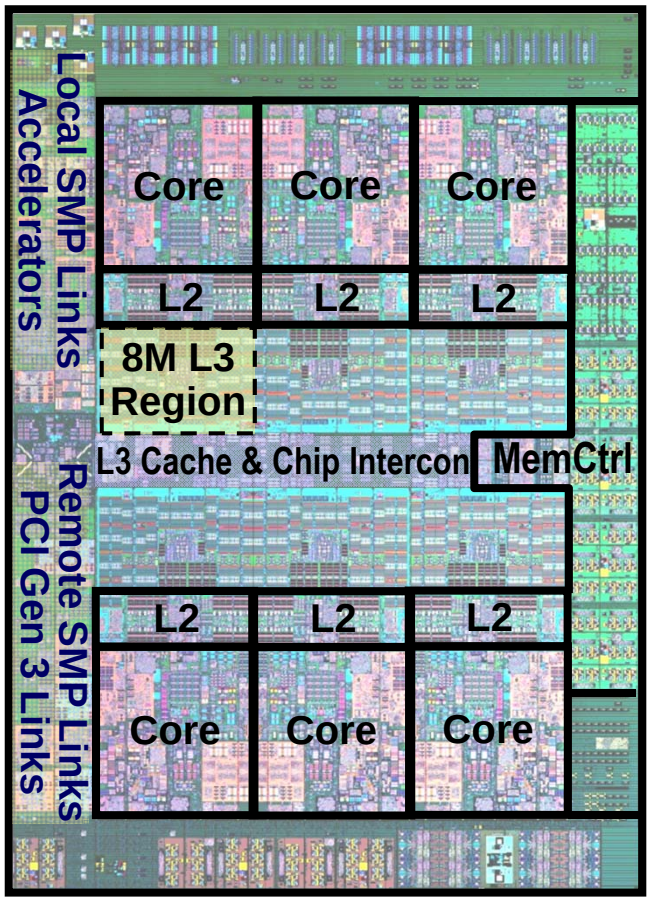

POWER8 Split Core Modes (Linux only)

Two modes:

- Big core mode
- Small core / Split Core mode)

Four small cores:

- Each core can run ST or SMT2 modes independent of the other
- Physical threads tied to each partition (i.e. LPAR0 = T0 and T1…)

Architecturally small-core is identical to the big-core

- Separate architected registers
- Separate TLBIE handling
- Recovery is across all 4 cores (Small core mode)
  - ❖ Restart all 4 partitions on error

Fairness between cores

- D-Cache: Split 8 sets 4 ways
- Dispatch priority: All 4 LPARS have equal priority

POWER8 Memory Buffer Chip

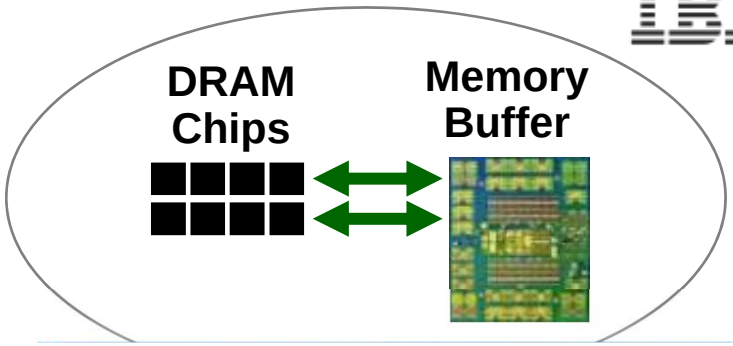

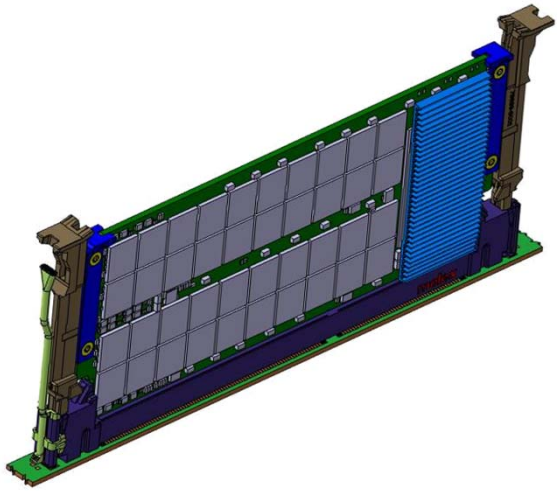

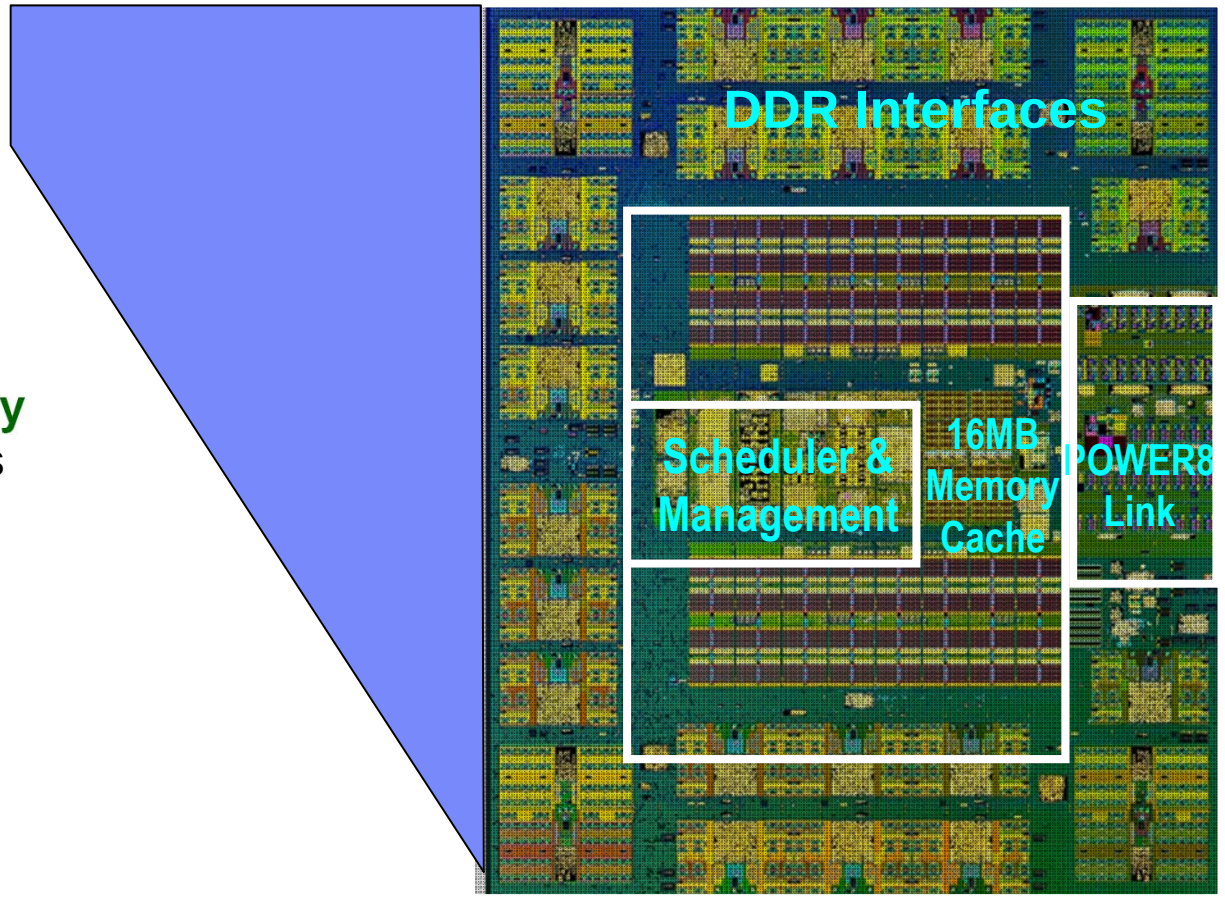

Intelligence Moved into Memory

- 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

Performance Value

- End-to-end fastpath and data retry (latency)
- Cache → latency/bandwidth, partial updates
- Cache → write scheduling, prefetch, energy

POWER8 Memory Organization (Max Config shown)

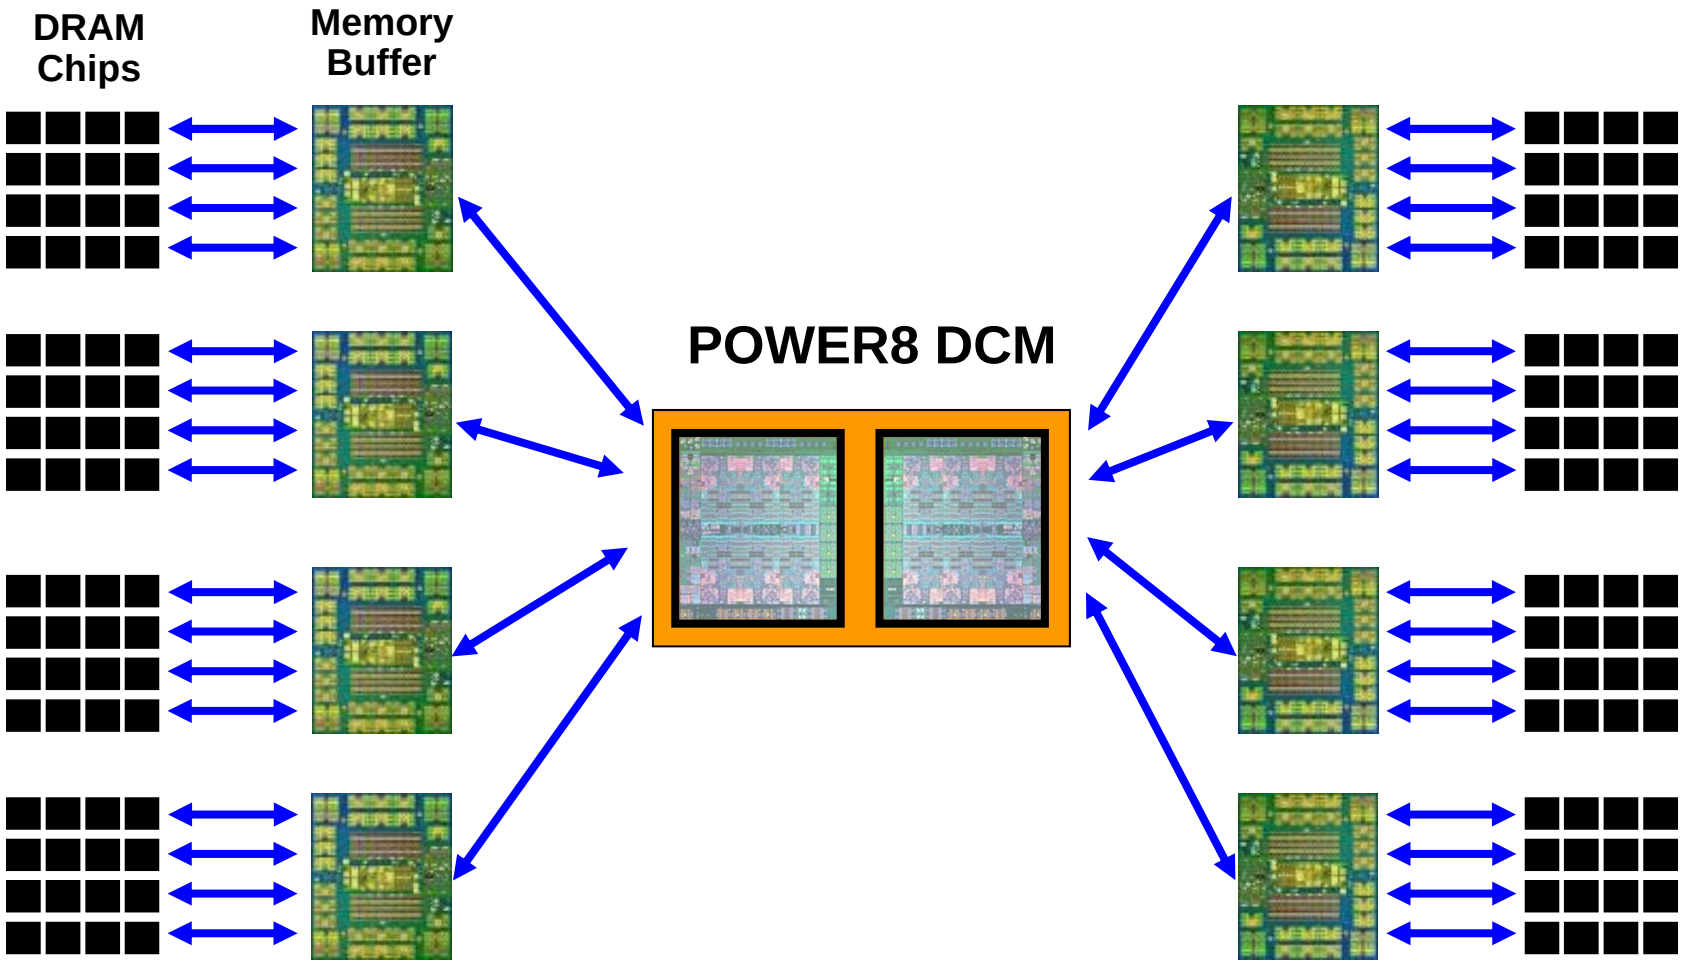

POWER8 Memory Organization

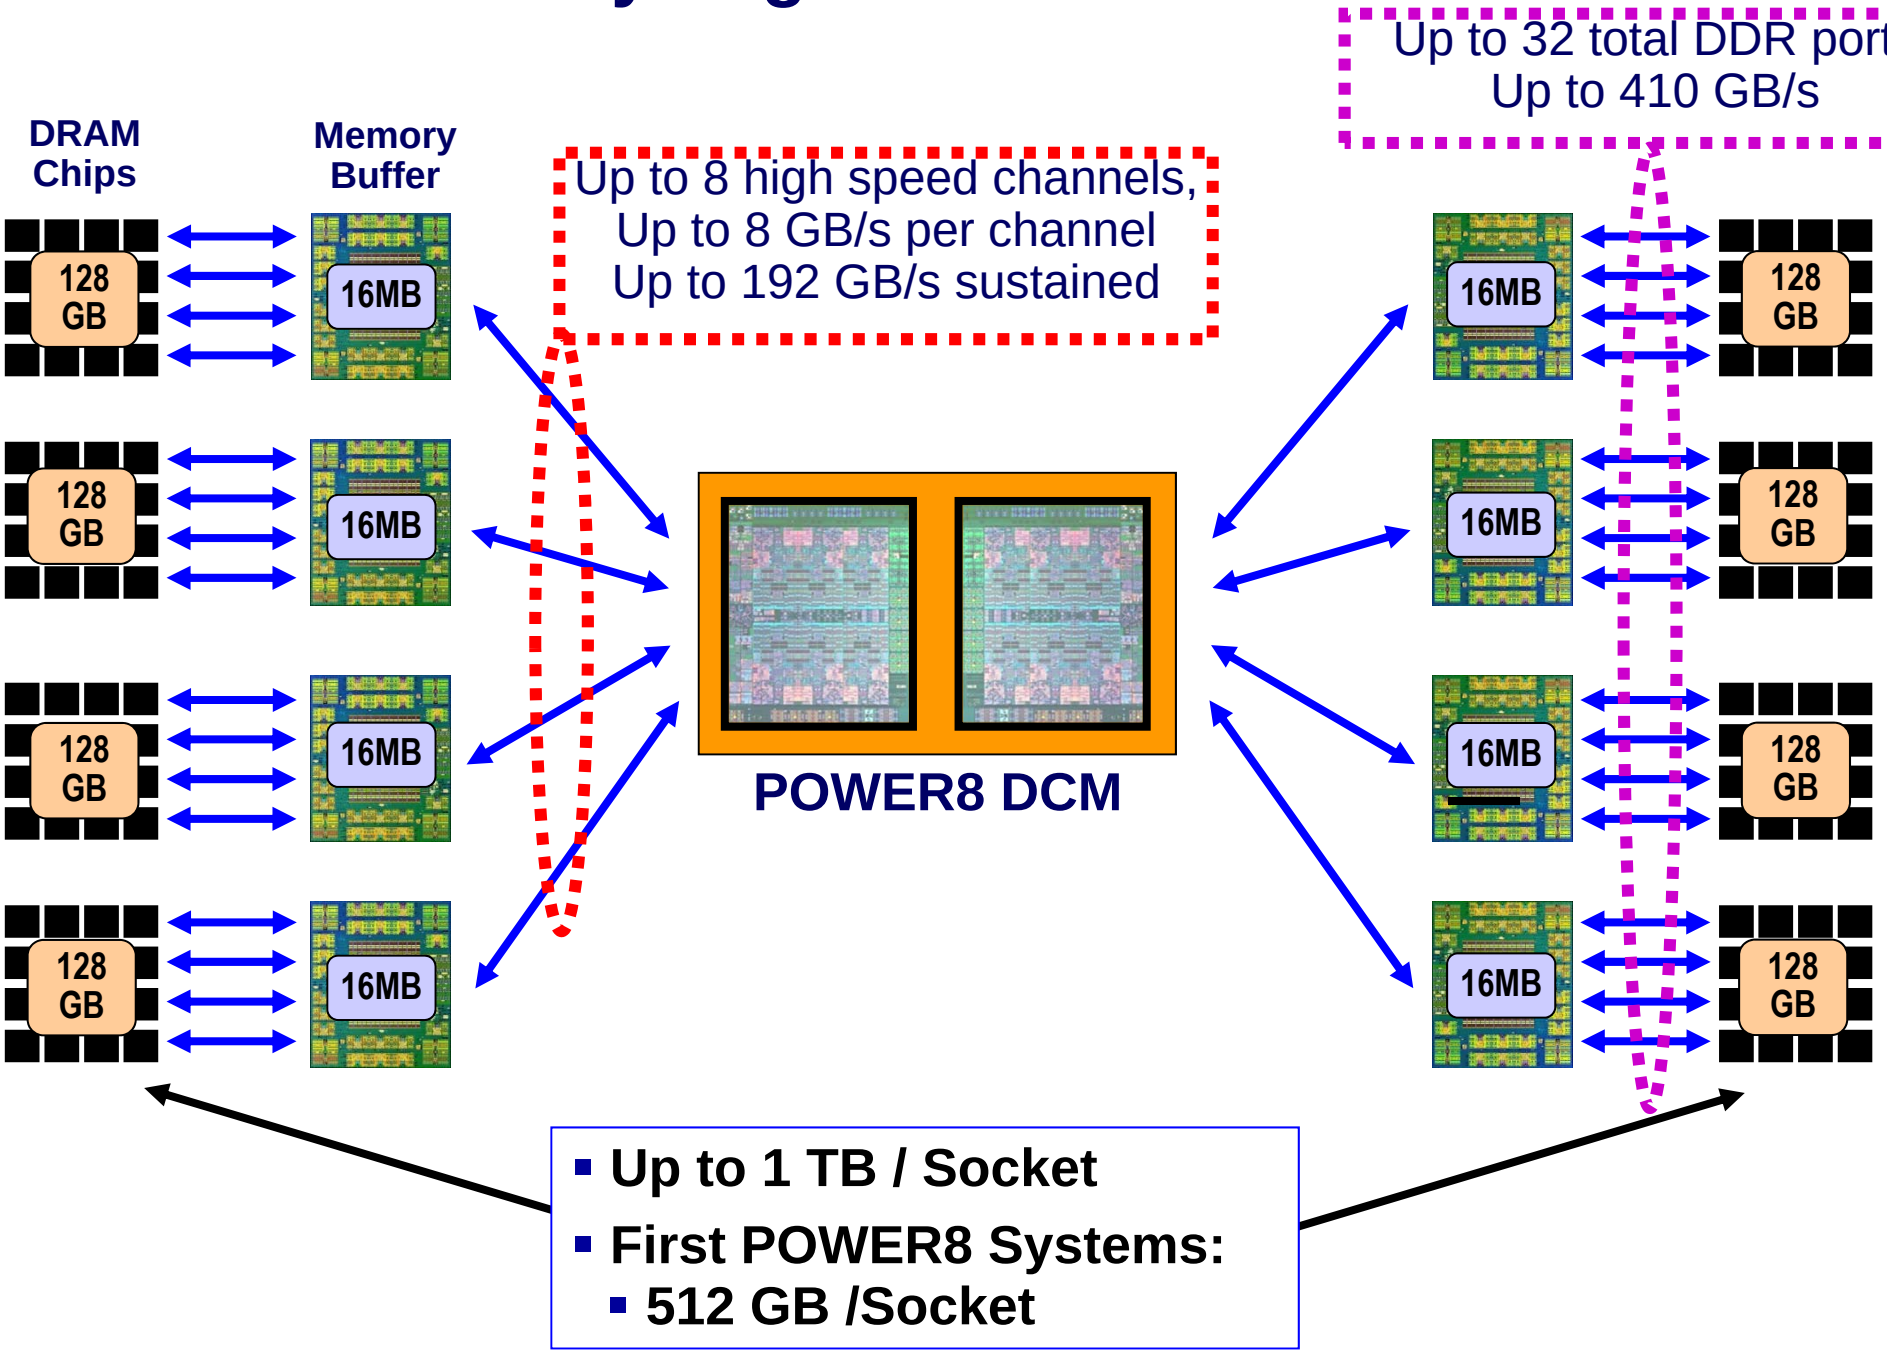

POWER8 Integrated PCI Gen 3

POWER7

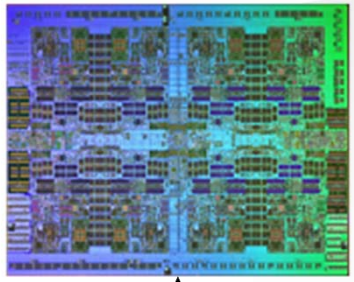

GX Bus

I/O Bridge

PCIe Gen2

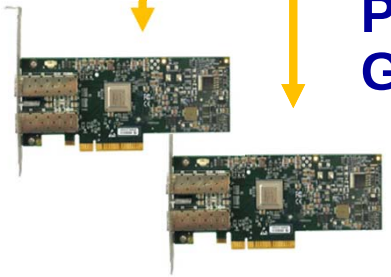

PCI Devices

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

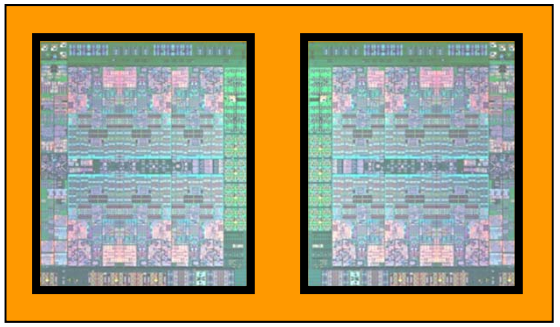

PCIe Gen3

x16

x8

x8

x16

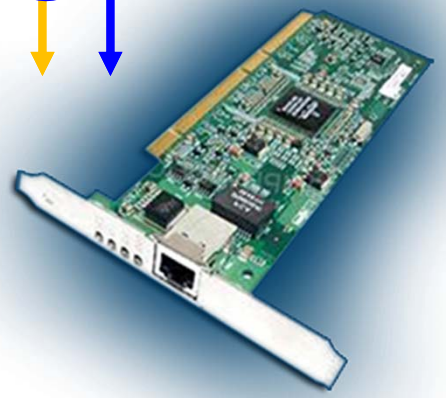

PCI Devices

POWER8 Bandwidths

IO Bandwidth

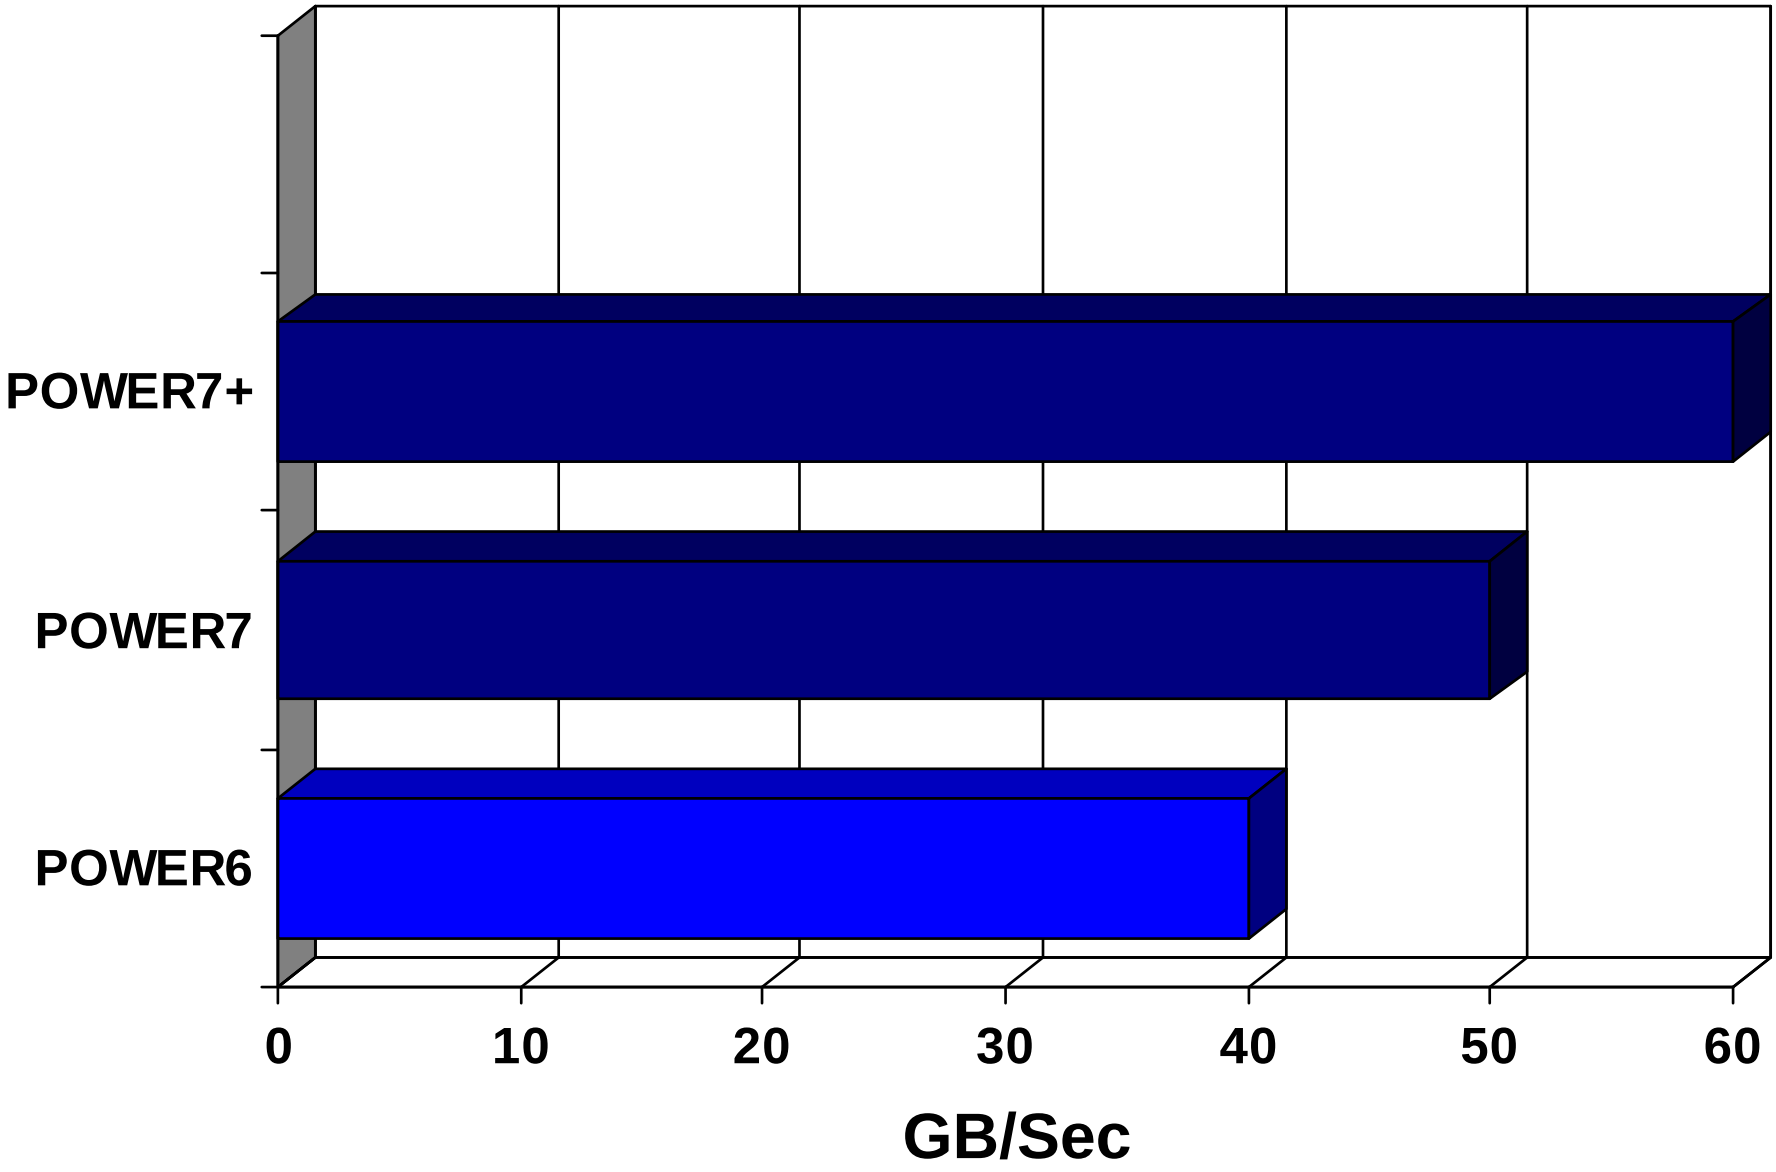

IO Bandwidth

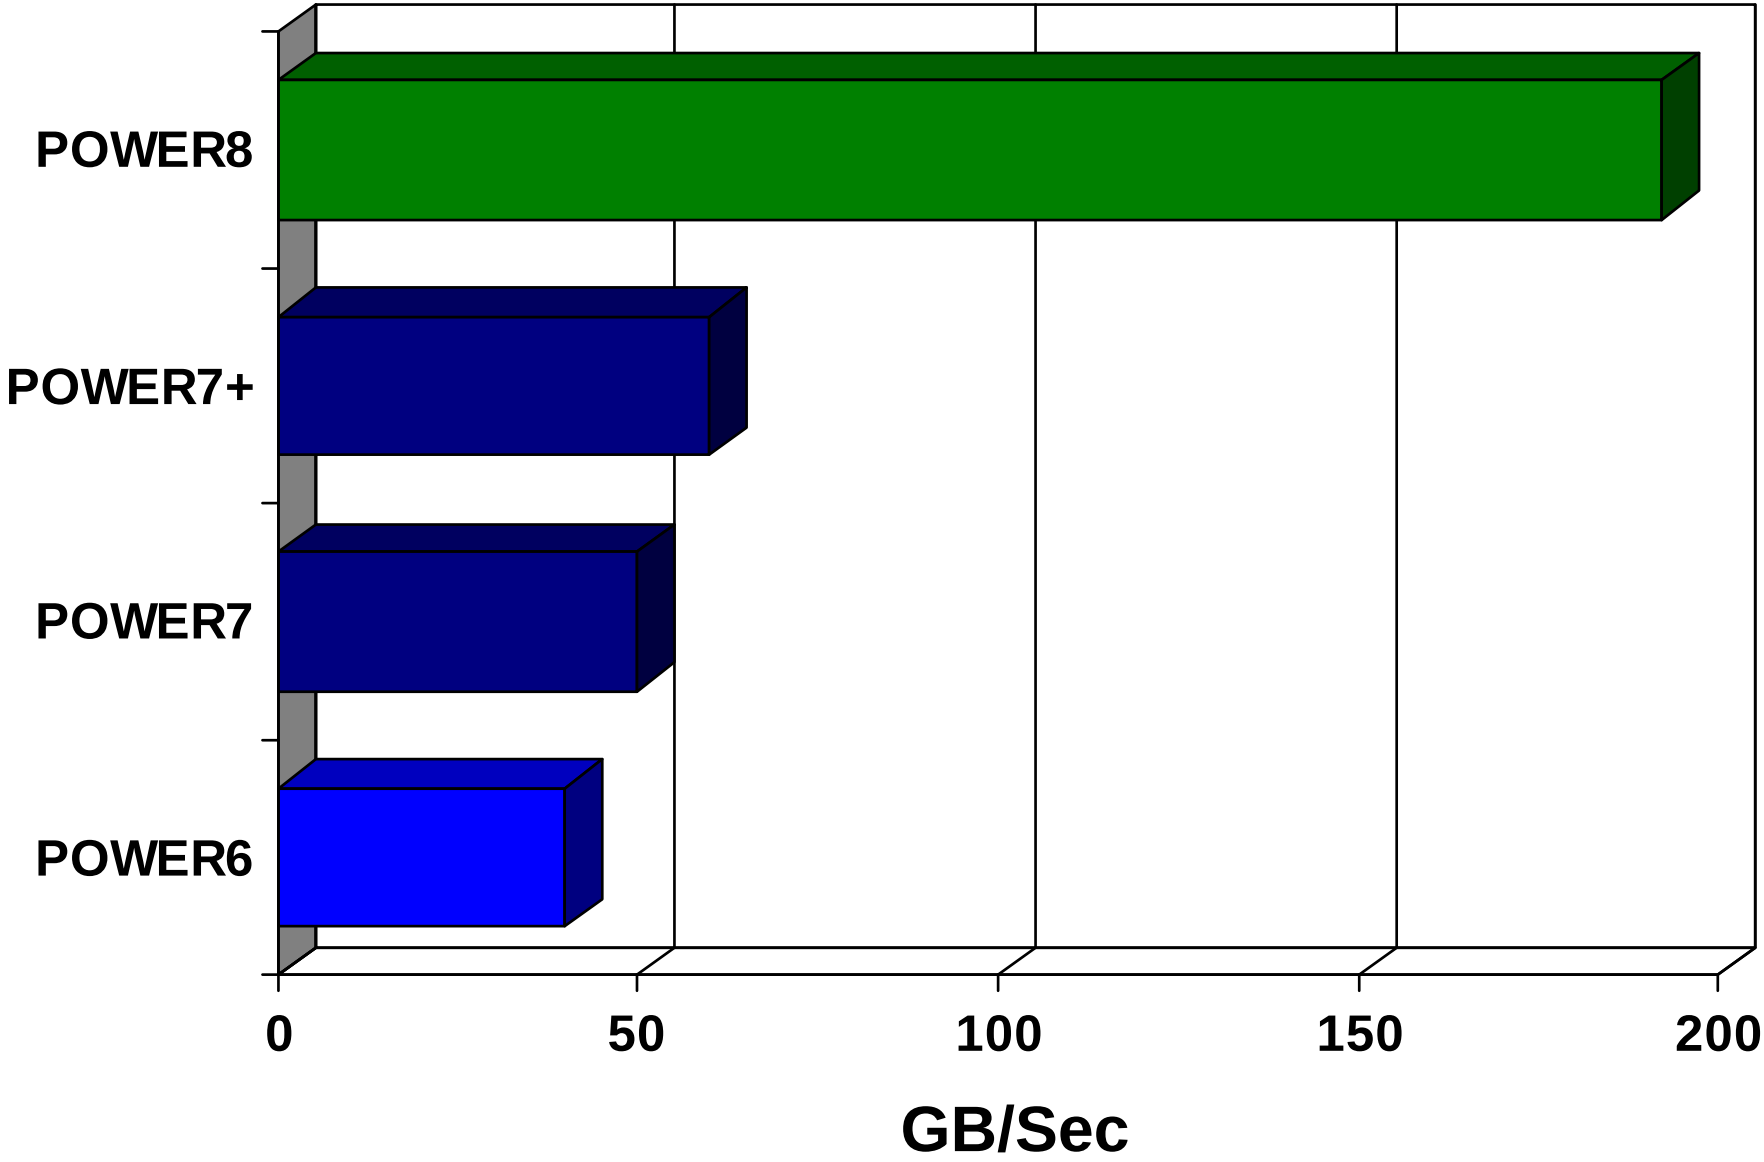

Memory Bandwidth per Socket

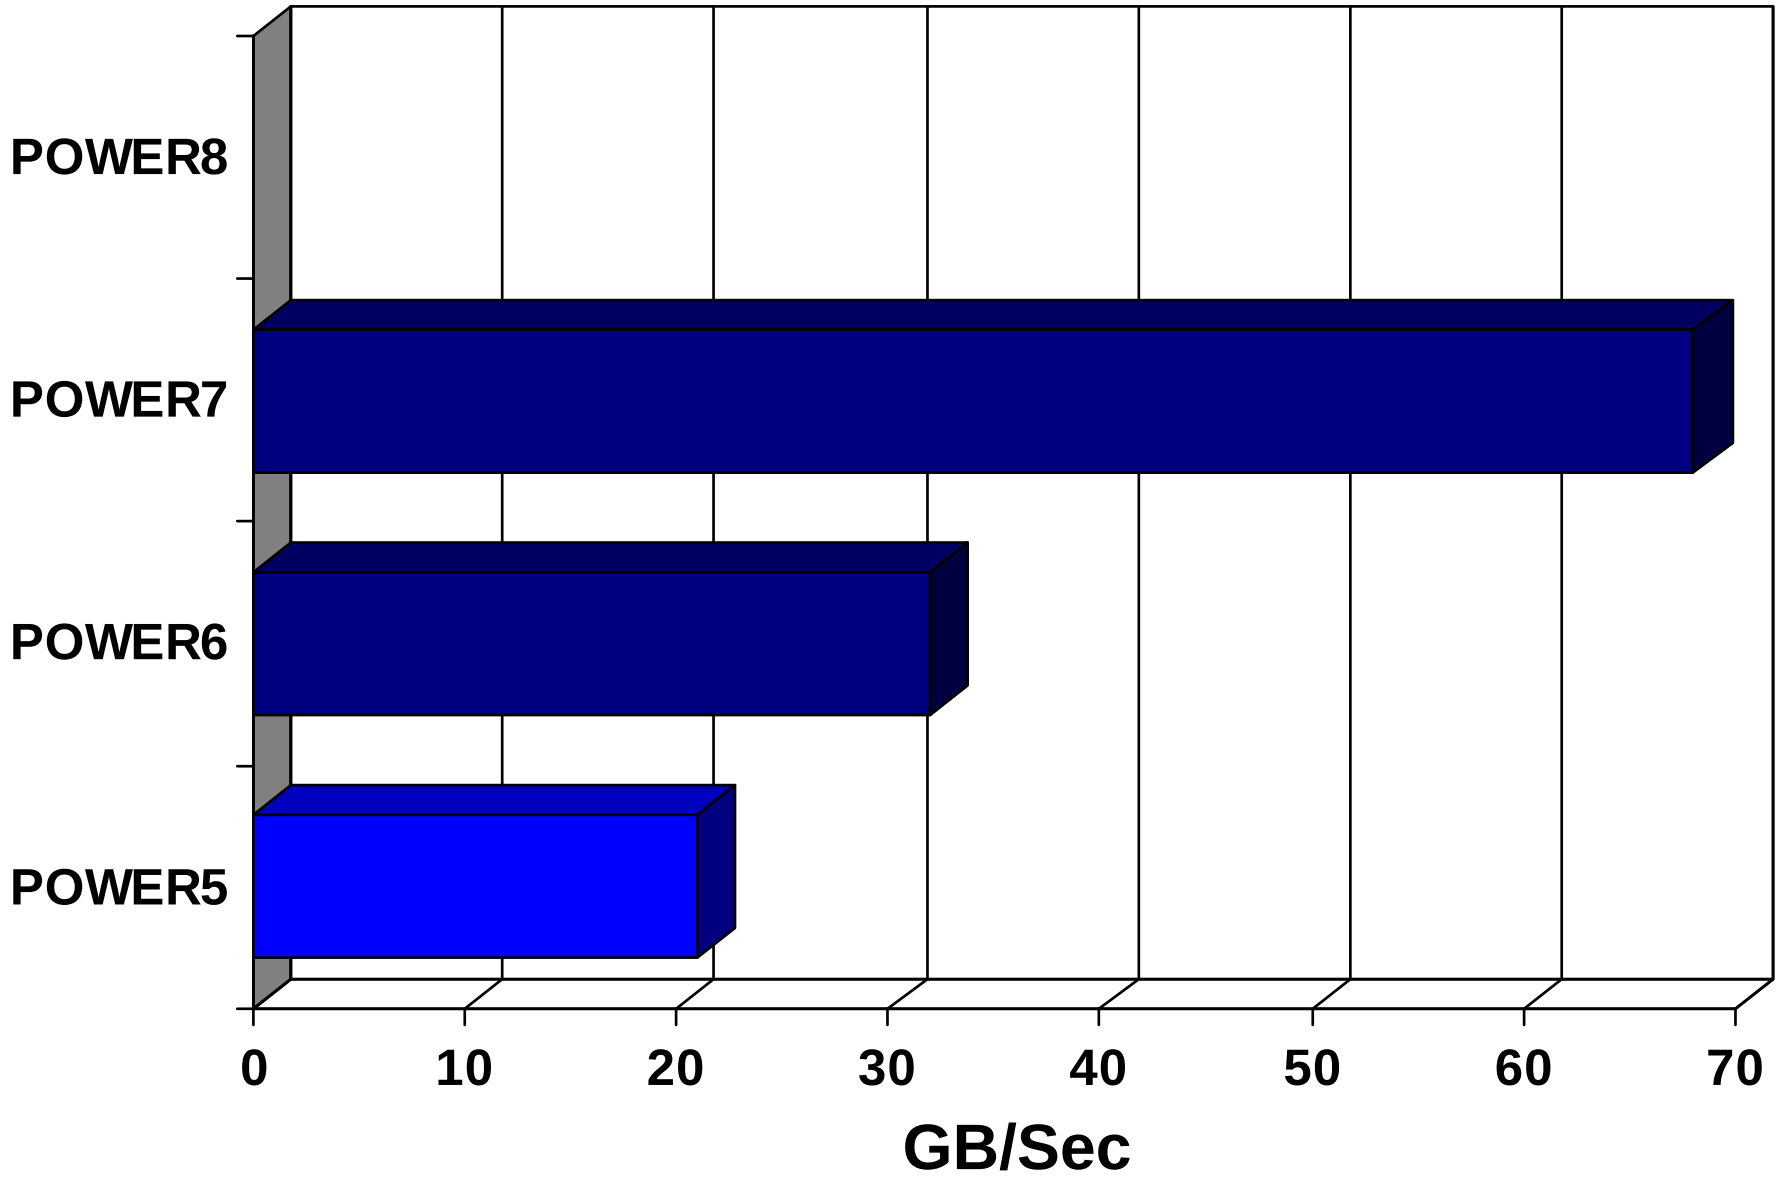

Memory Bandwidth per Socket

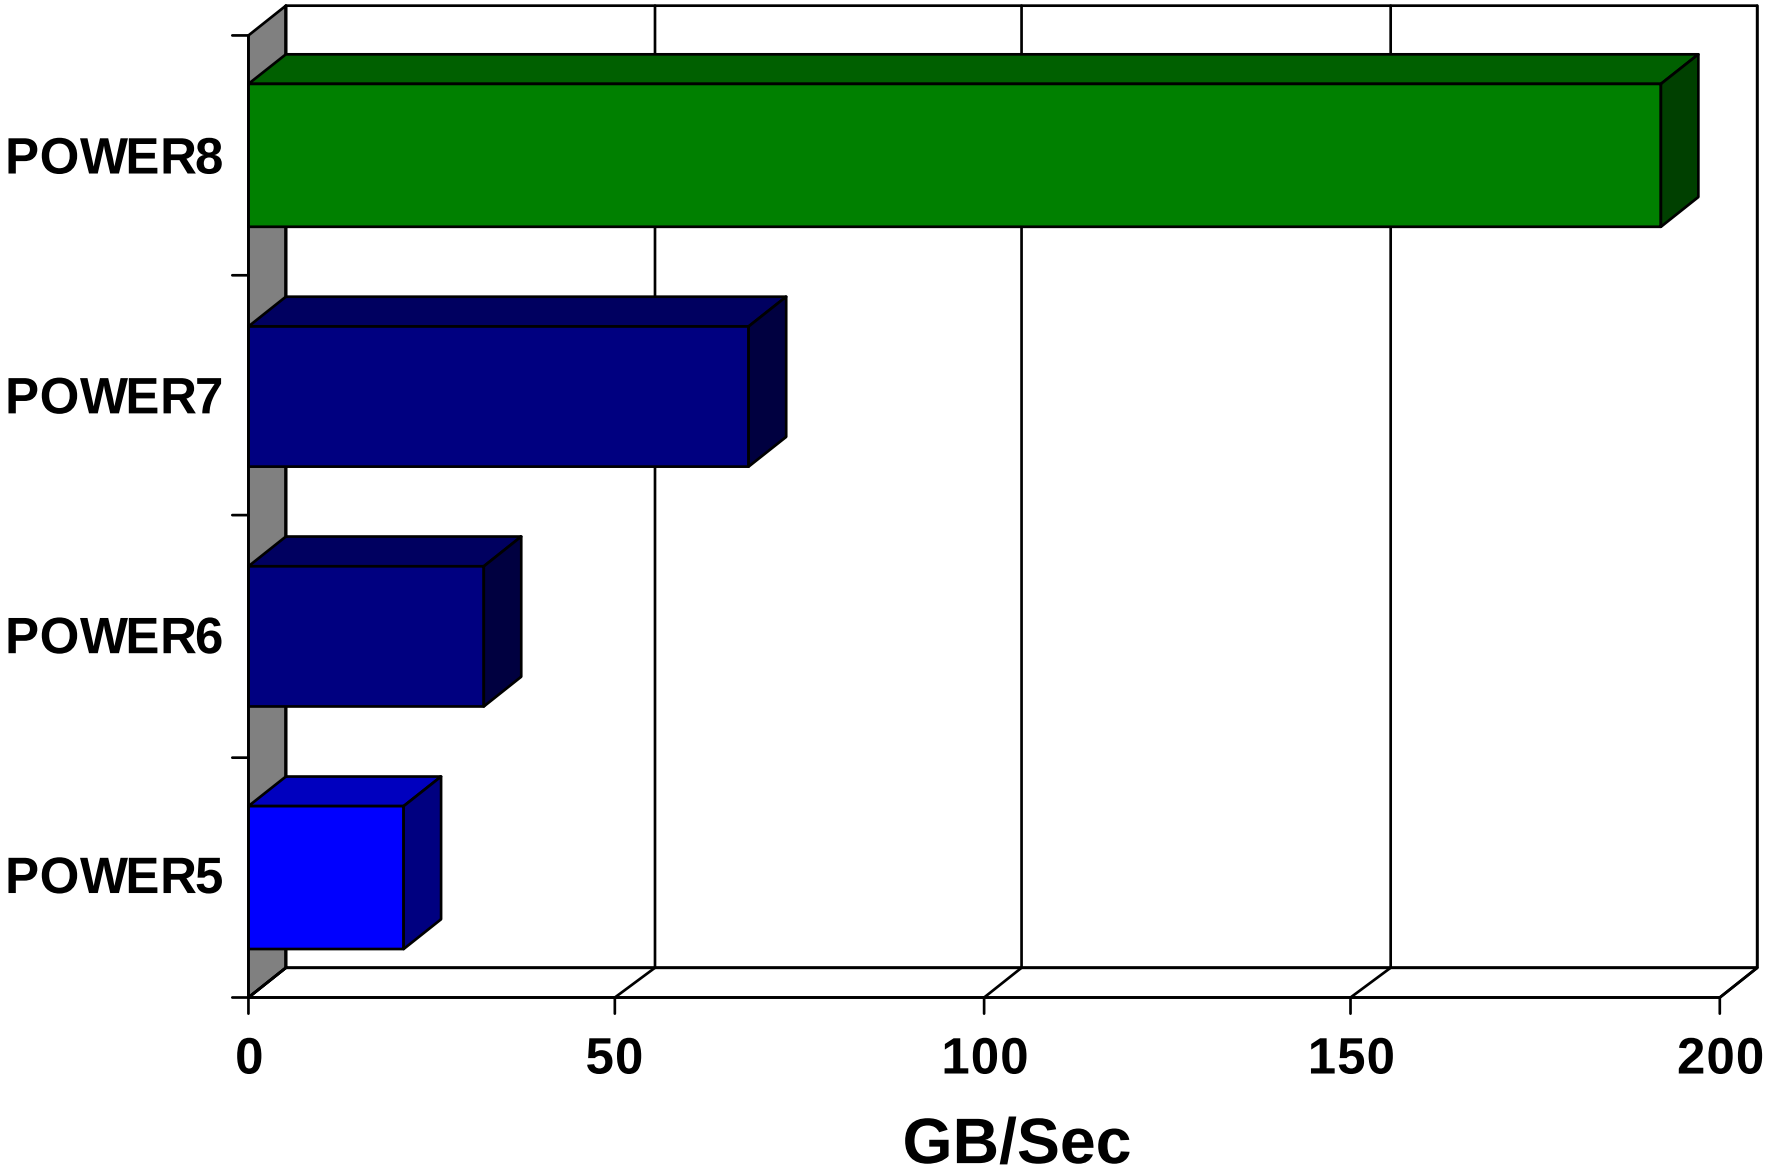

Servers

A Thing about Naming

Power 8 Naming Convention

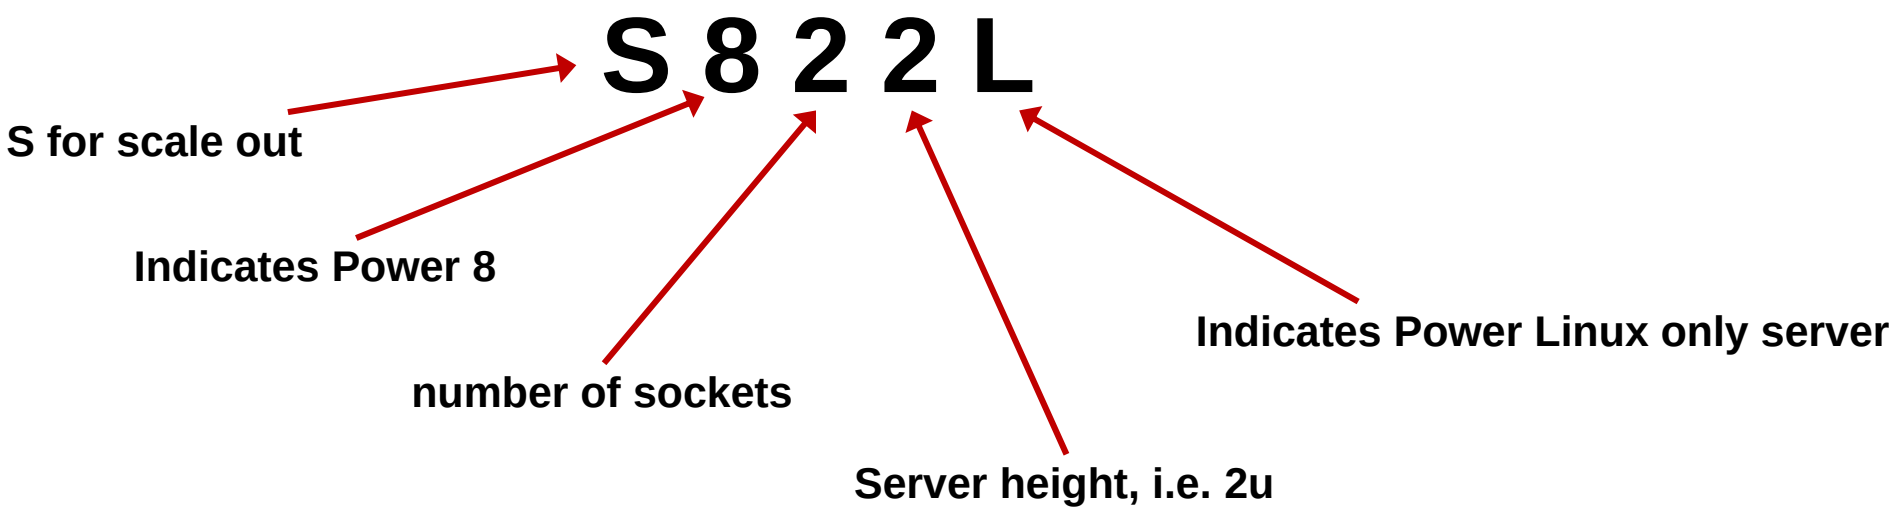

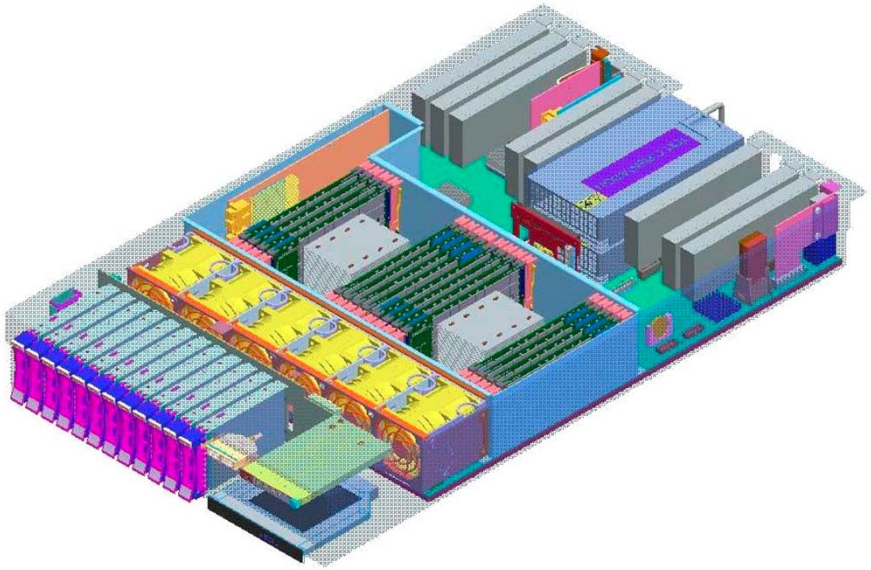

2U Server

POWER8 2S2U Layout

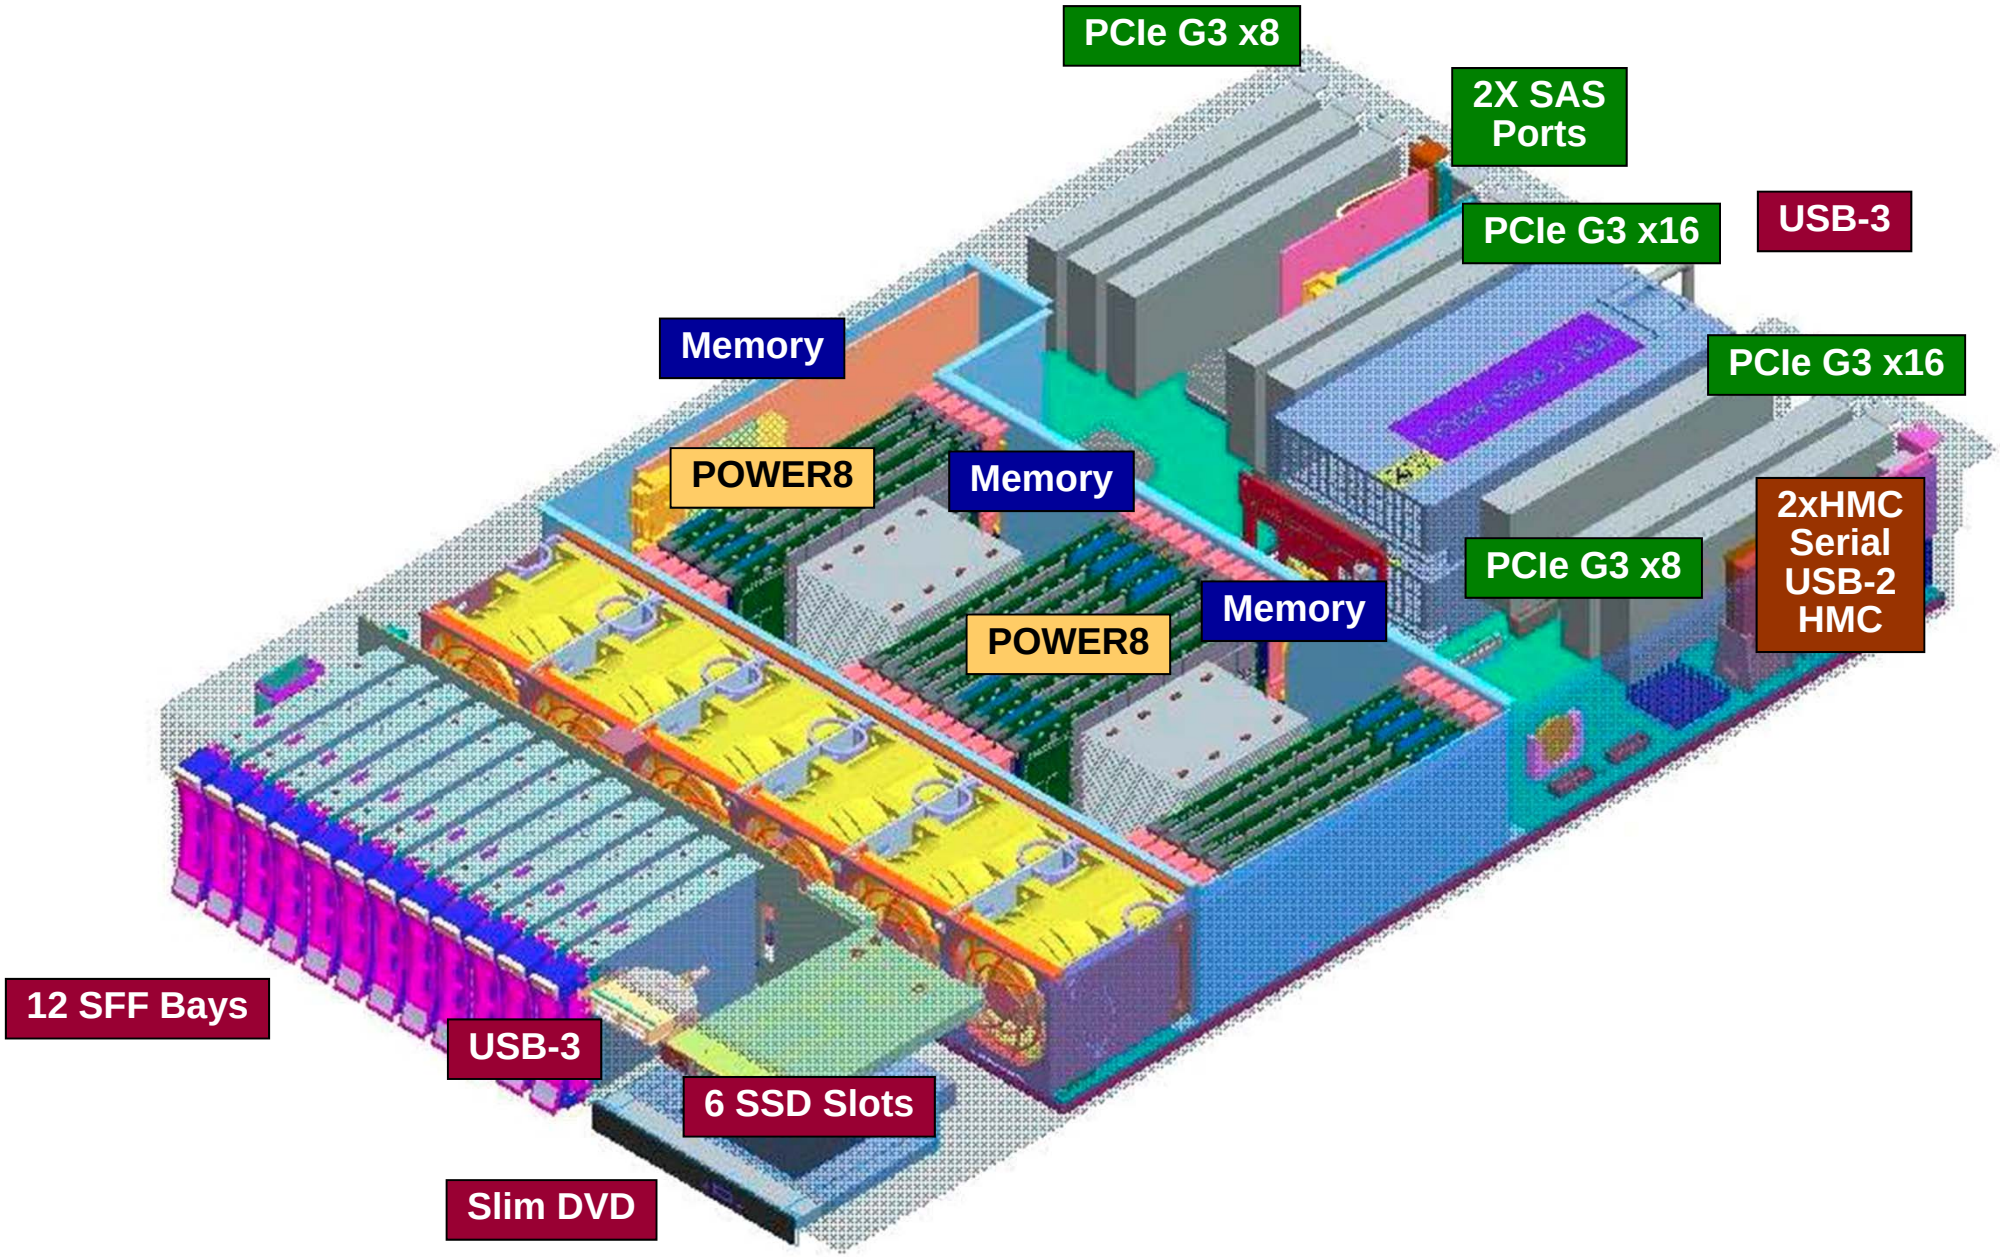

POWER8 2S2U Scale-Out System

Power S822

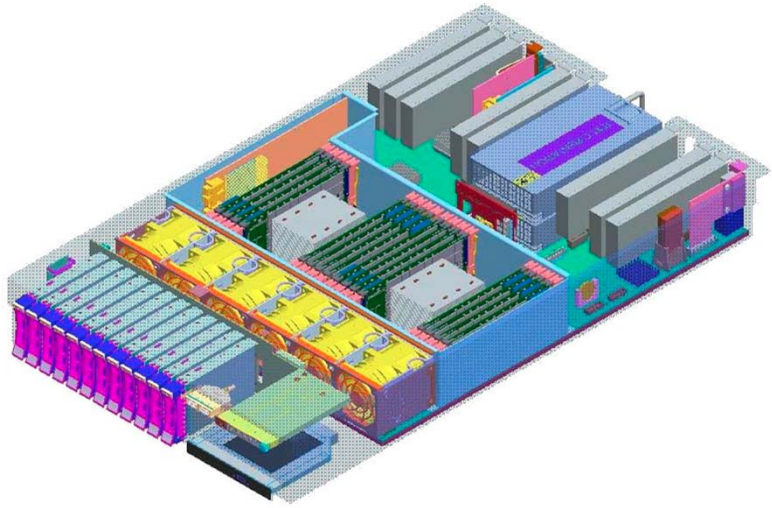

- Single Socket populated
  - Cores: 6 or 10
  - Memory: Up to 512 GB
  - PCIe Slots: 6 PCIe Gen3 LP (Hotplug)
- Both Sockets populated
  - Cores: 12 or 20
  - Memory: Up to 1 TB
  - PCIe Slots: 9 PCIe Gen3 LP (Hotplug)

- 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

3 Yr Warranty

- Hypervisor: PowerVM
- OS: AIX, Linux (not IBM i)

POWER8 Linux 2S2U Scale-Out System

Power S822L

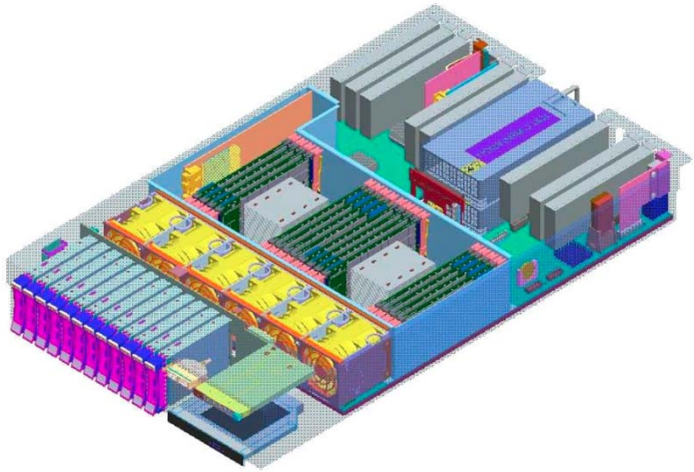

- Dual Sockets
  - Cores: 20 or 24
  - Memory: Up to 1 TB
  - Slots: 9 PCIe Gen3 LP (Hotplug)
- Ethernet: Quad 1 Gbt / (x8 slot)
- Integrated ports: USB (4), Serial (2), HMC (2)
- Internal Storage
  - DVD
  - 12 SFF Bays -- Split Backplane: 6 + 6
  - or 8 SFF Bays & 6 1.8” SSD Bays with Easy Tier with 7GB write cache
- Hypervisor: PowerVM or PowerKVM
- OS: Linux

3 Yr Warranty

POWER8 2S2U Scale-Out Comparison - S822

	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 MB @ 1066 MHz	**512 / 1024 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	Yes (3 + 3)	**Yes (6 + 6)**
NIC SRIOV Support	No	**SoD**
Service Processor	Generation 1	**Generation 2**

POWER8 2S2U Scale-Out Comparison - S822L

	Power 730	Power System S822L
Processor	POWER7+	**POWER8**
Sockets	2	**2**
Cores	16	**24**
Maximum Memory	512 MB @ 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	Yes (3 + 3)	**Yes (6 + 6)**
NIC SRIOV Support	No	**SoD**
Service Processor	Generation 1	**Generation 2**

POWER8 2S2U Layout

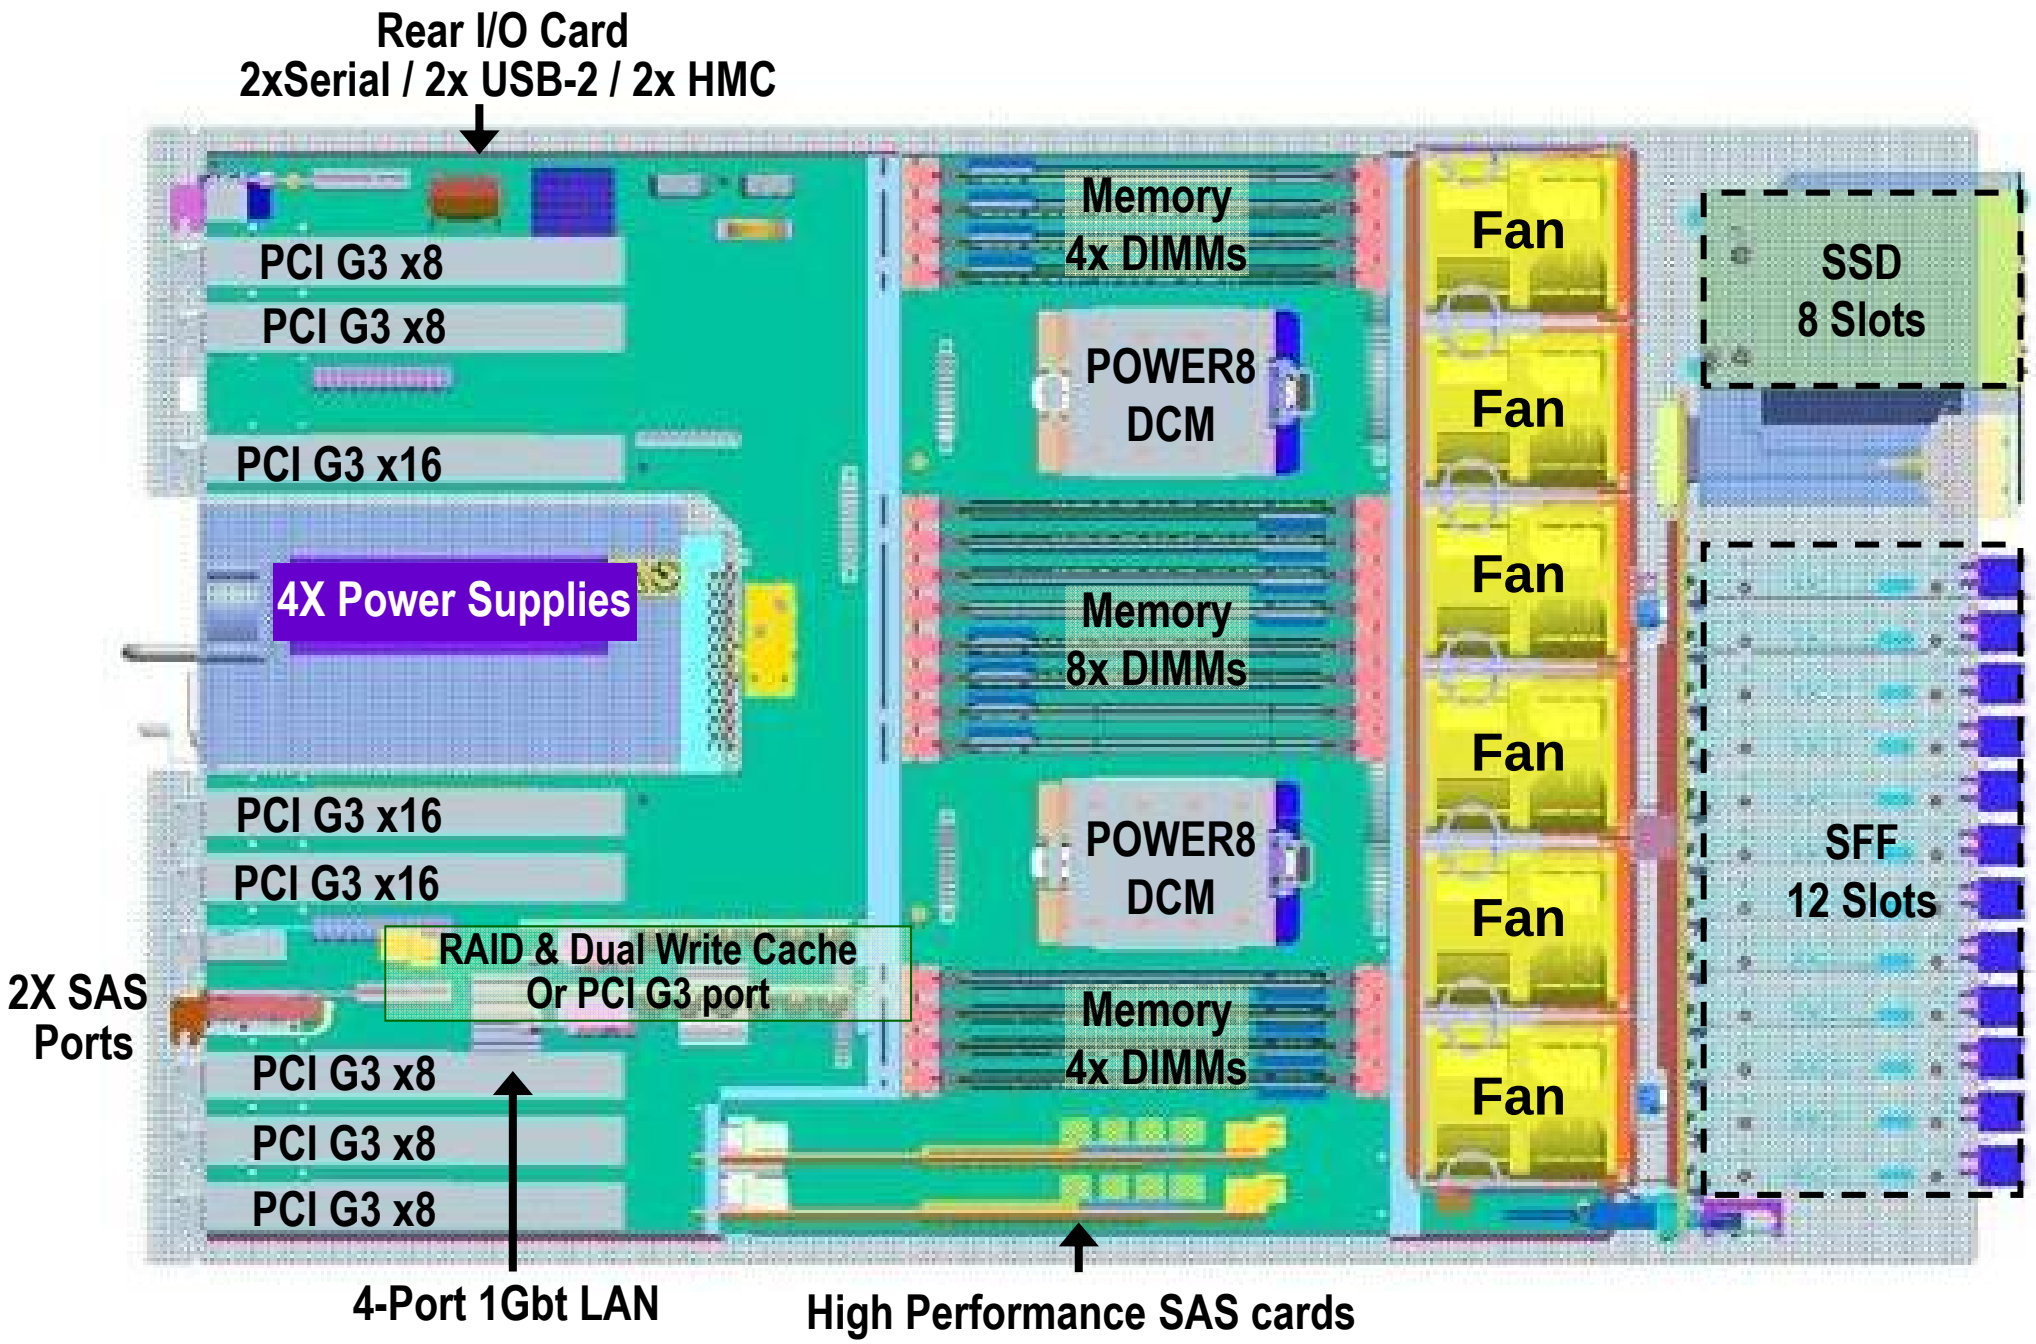

POWER8 2S2U Review View

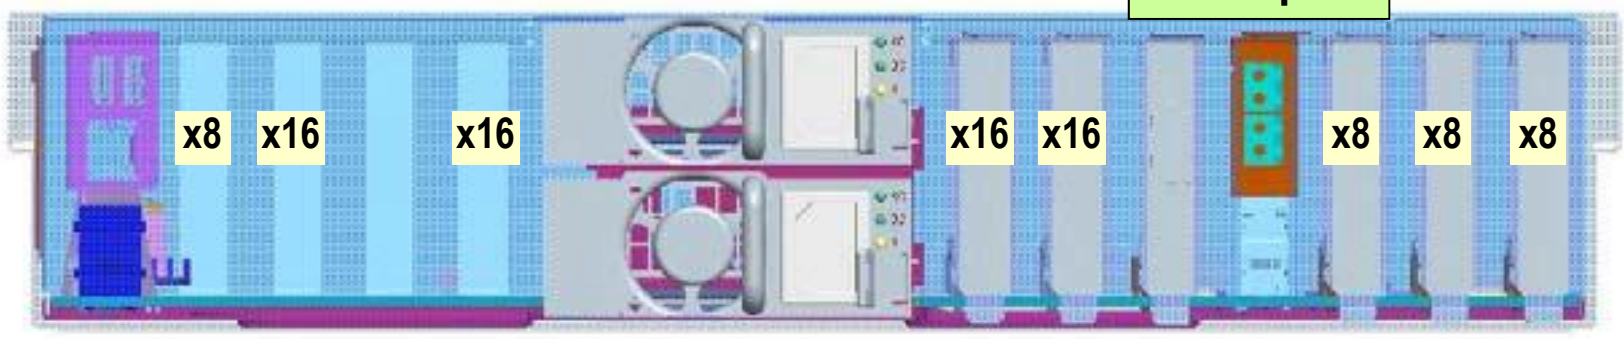

POWER8 2S2U Planar Layout

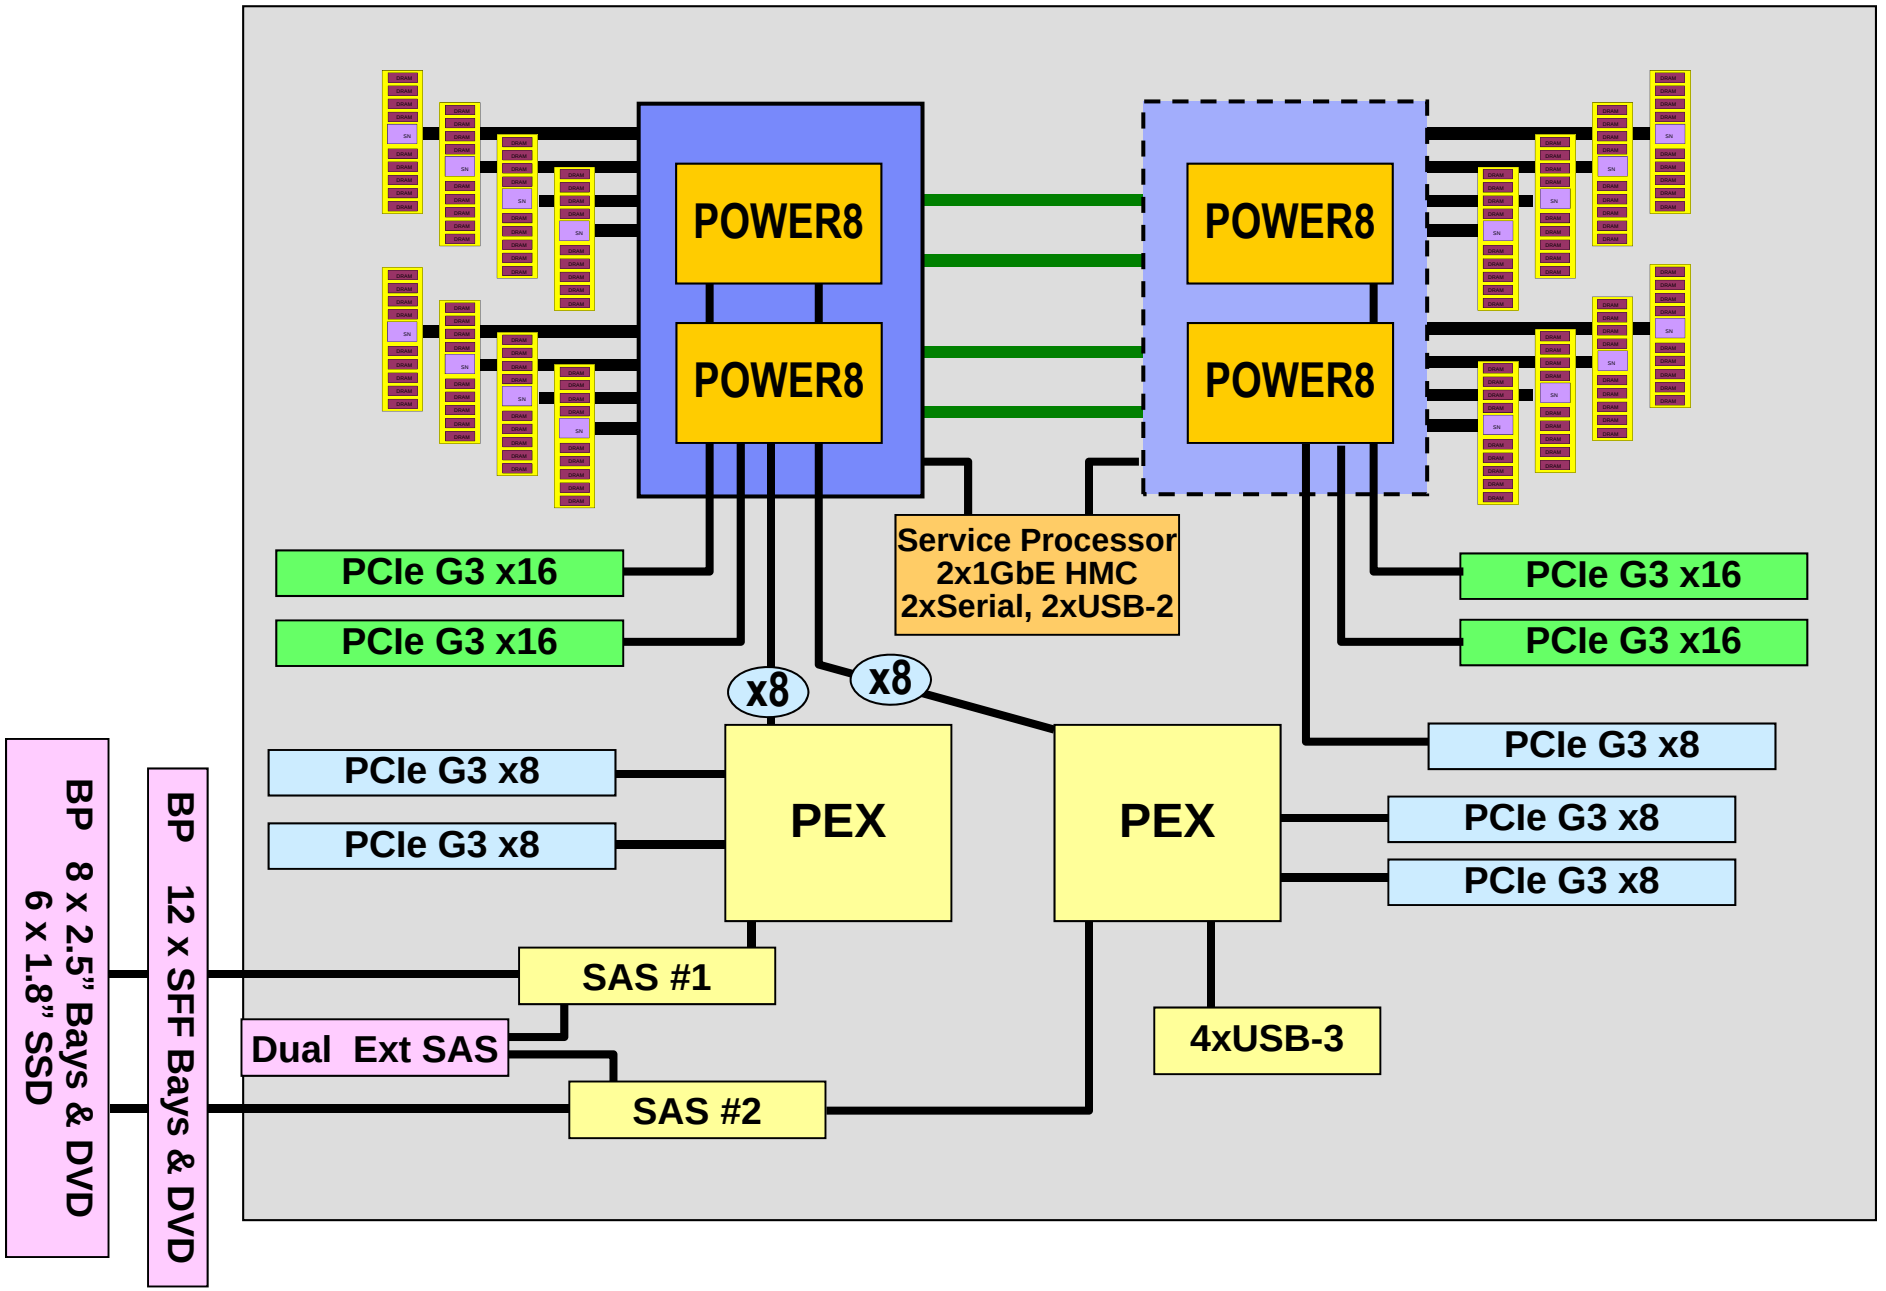

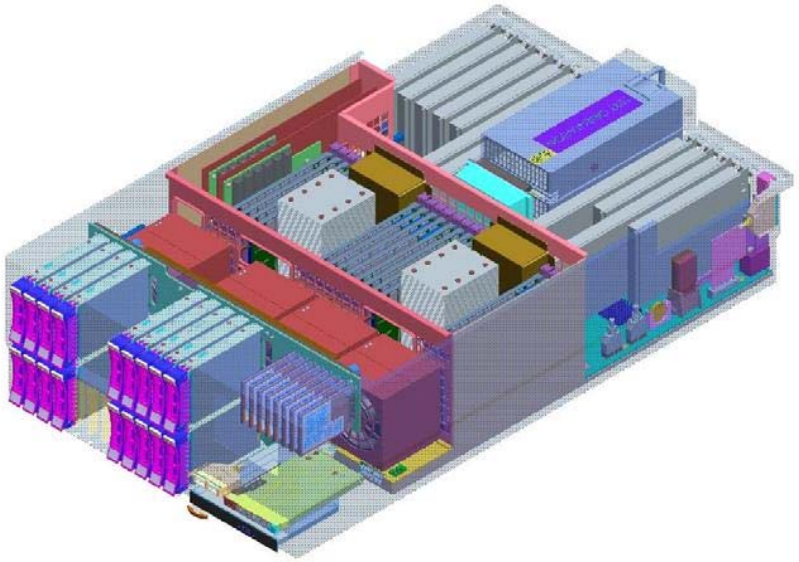

4U Servers

System 1S4U Layout

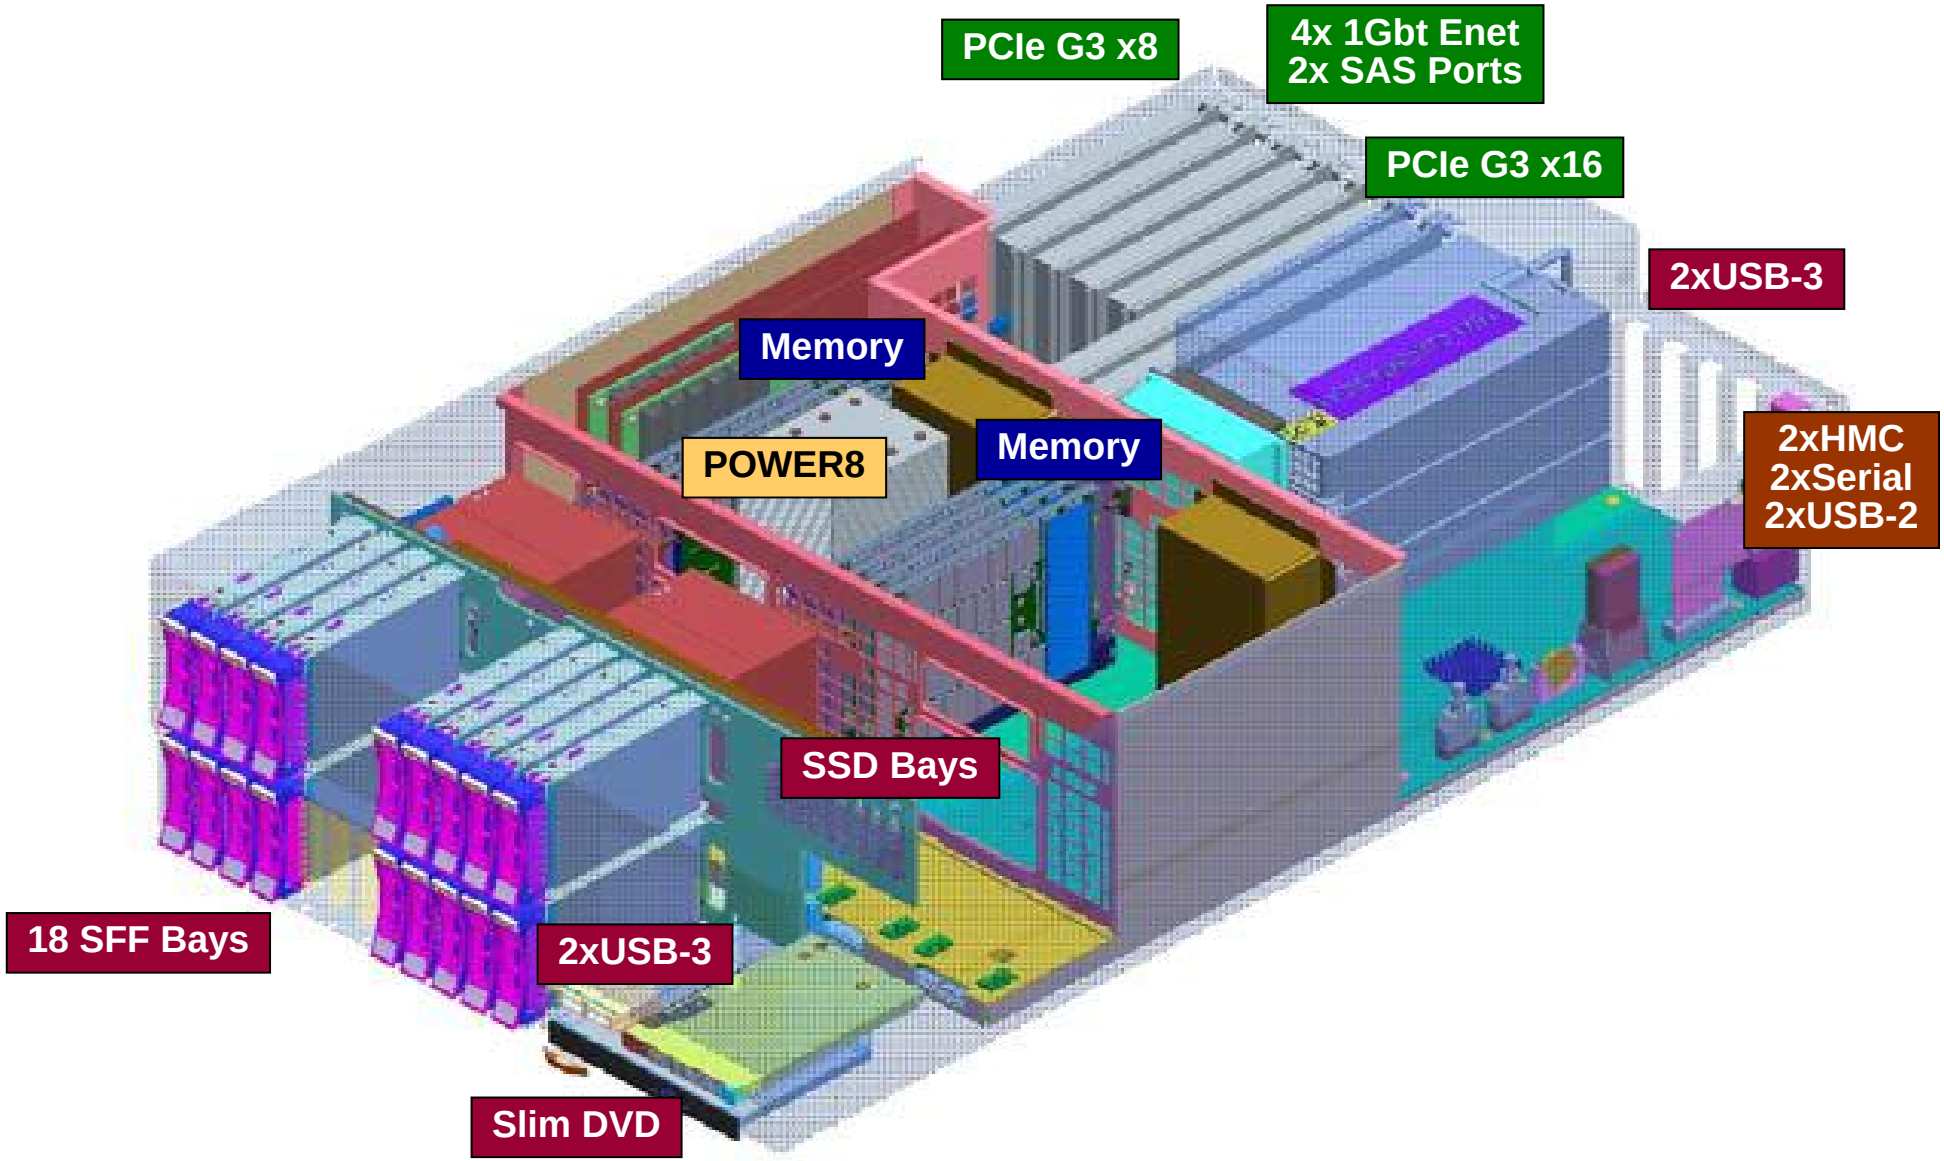

System 2S4U Layout

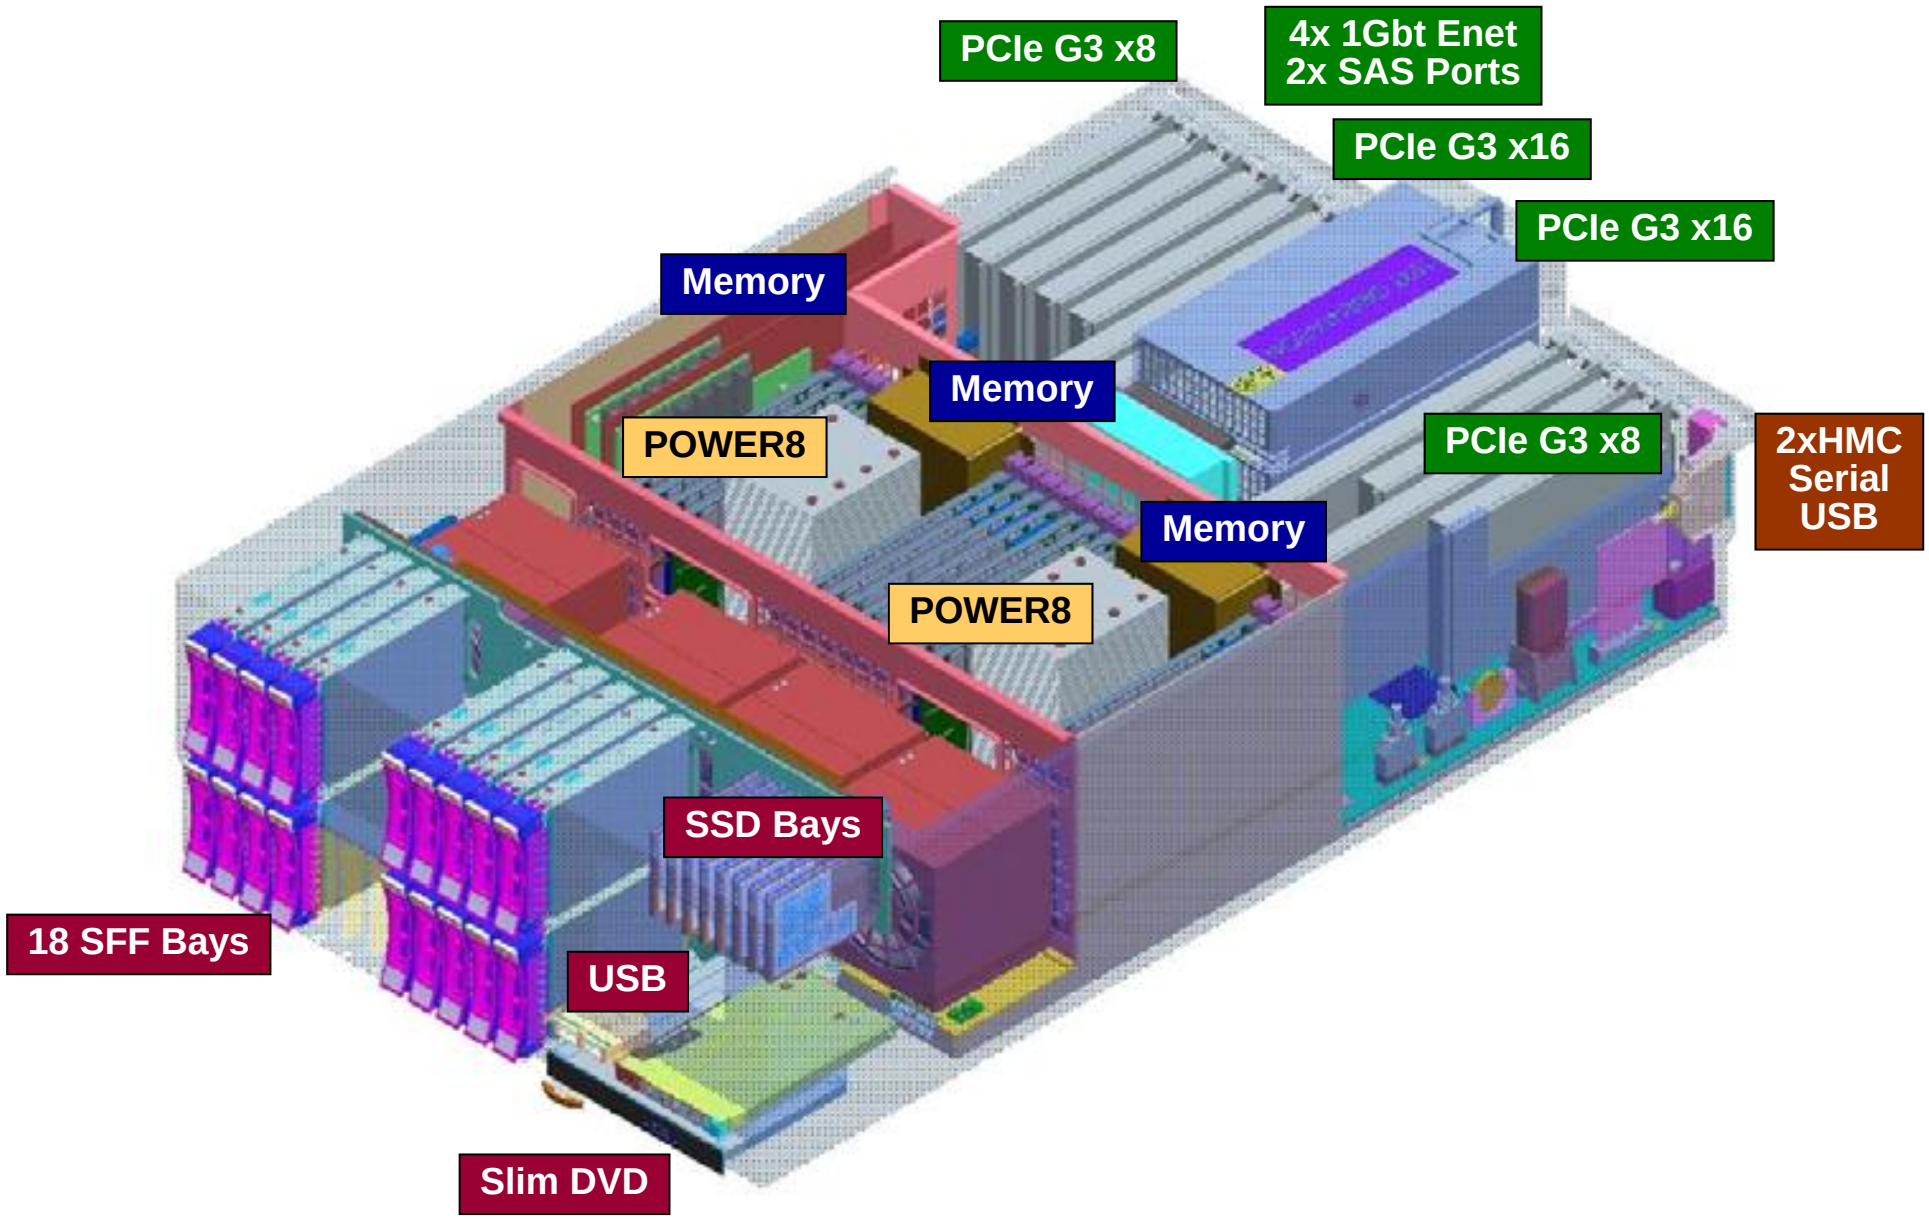

POWER8 1S4U Scale-Out System

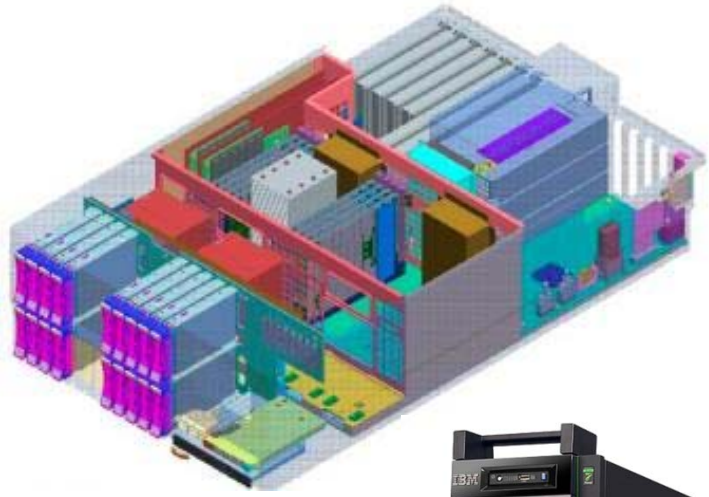

Power S814

- Form Factor: 4U or Tower
- Single Socket
  - Cores: 6 or 8
  - 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
  - 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 Yr Warranty

POWER8 2S4U Scale-Out System

Power S824

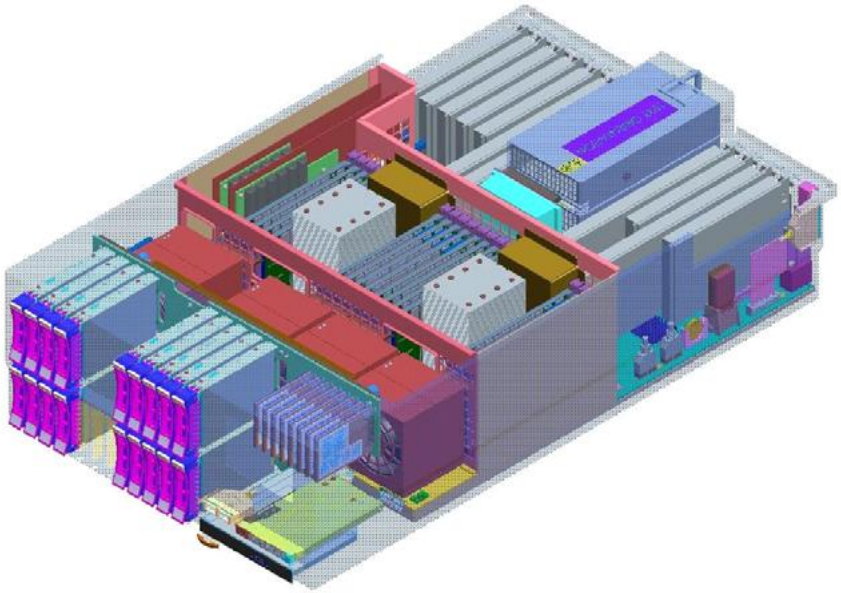

- Single Socket populated
  - Cores: 6 or 8
  - Memory: Up to 512 GB
  - Slots: 7 PCIe Gen3 full-high (Hotplug)
- Both Sockets populated
  - Cores: 12, 16, or 24
  - Memory: Up to 1 TB
  - Slots: 11 PCIe Gen3 full-high (Hotplug)

- Ethernet: Quad 1 Gbt / (x8 slot)
- Integrated ports: USB (4/5), Serial (2), HMC (2)
- Internal Storage
  - DVD
  - 12 SFF Bays -- Split Backplane: 6 + 6
  - or 18 SFF bays & 8 SSD bays with Easy Tier with 7GB write cache

3 Yr Warranty

- Hypervisor: PowerVM
- OS: AIX, IBM i, Linux

POWER8 4U Scale-Out Comparison – S814

	Power 720	Power System S814
Processor	POWER7+	**POWER8**
Sockets	1	**1**
Cores	4 / 6 / 8	**6 / 8**
Maximum Memory	512 @ 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)**
NIC SRIOV Support	No	**SOD**
Service Processor	Generation 1	**Generation 2**

POWER8 4U Scale-Out Comparison – S824

	Power 740	Power System S824
Processor	POWER7+	**POWER8**
Sockets	1 (upgradeable) / 2	**1 (upgradeable) / 2**
Max Cores	8 / 16	**8 / 24**
Maximum Memory	512GB / 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)**
SRIOV Support	No	**SOD**
Service Processor	Generation 1	**Generation 2**

POWER8 1S4U Planar Layout

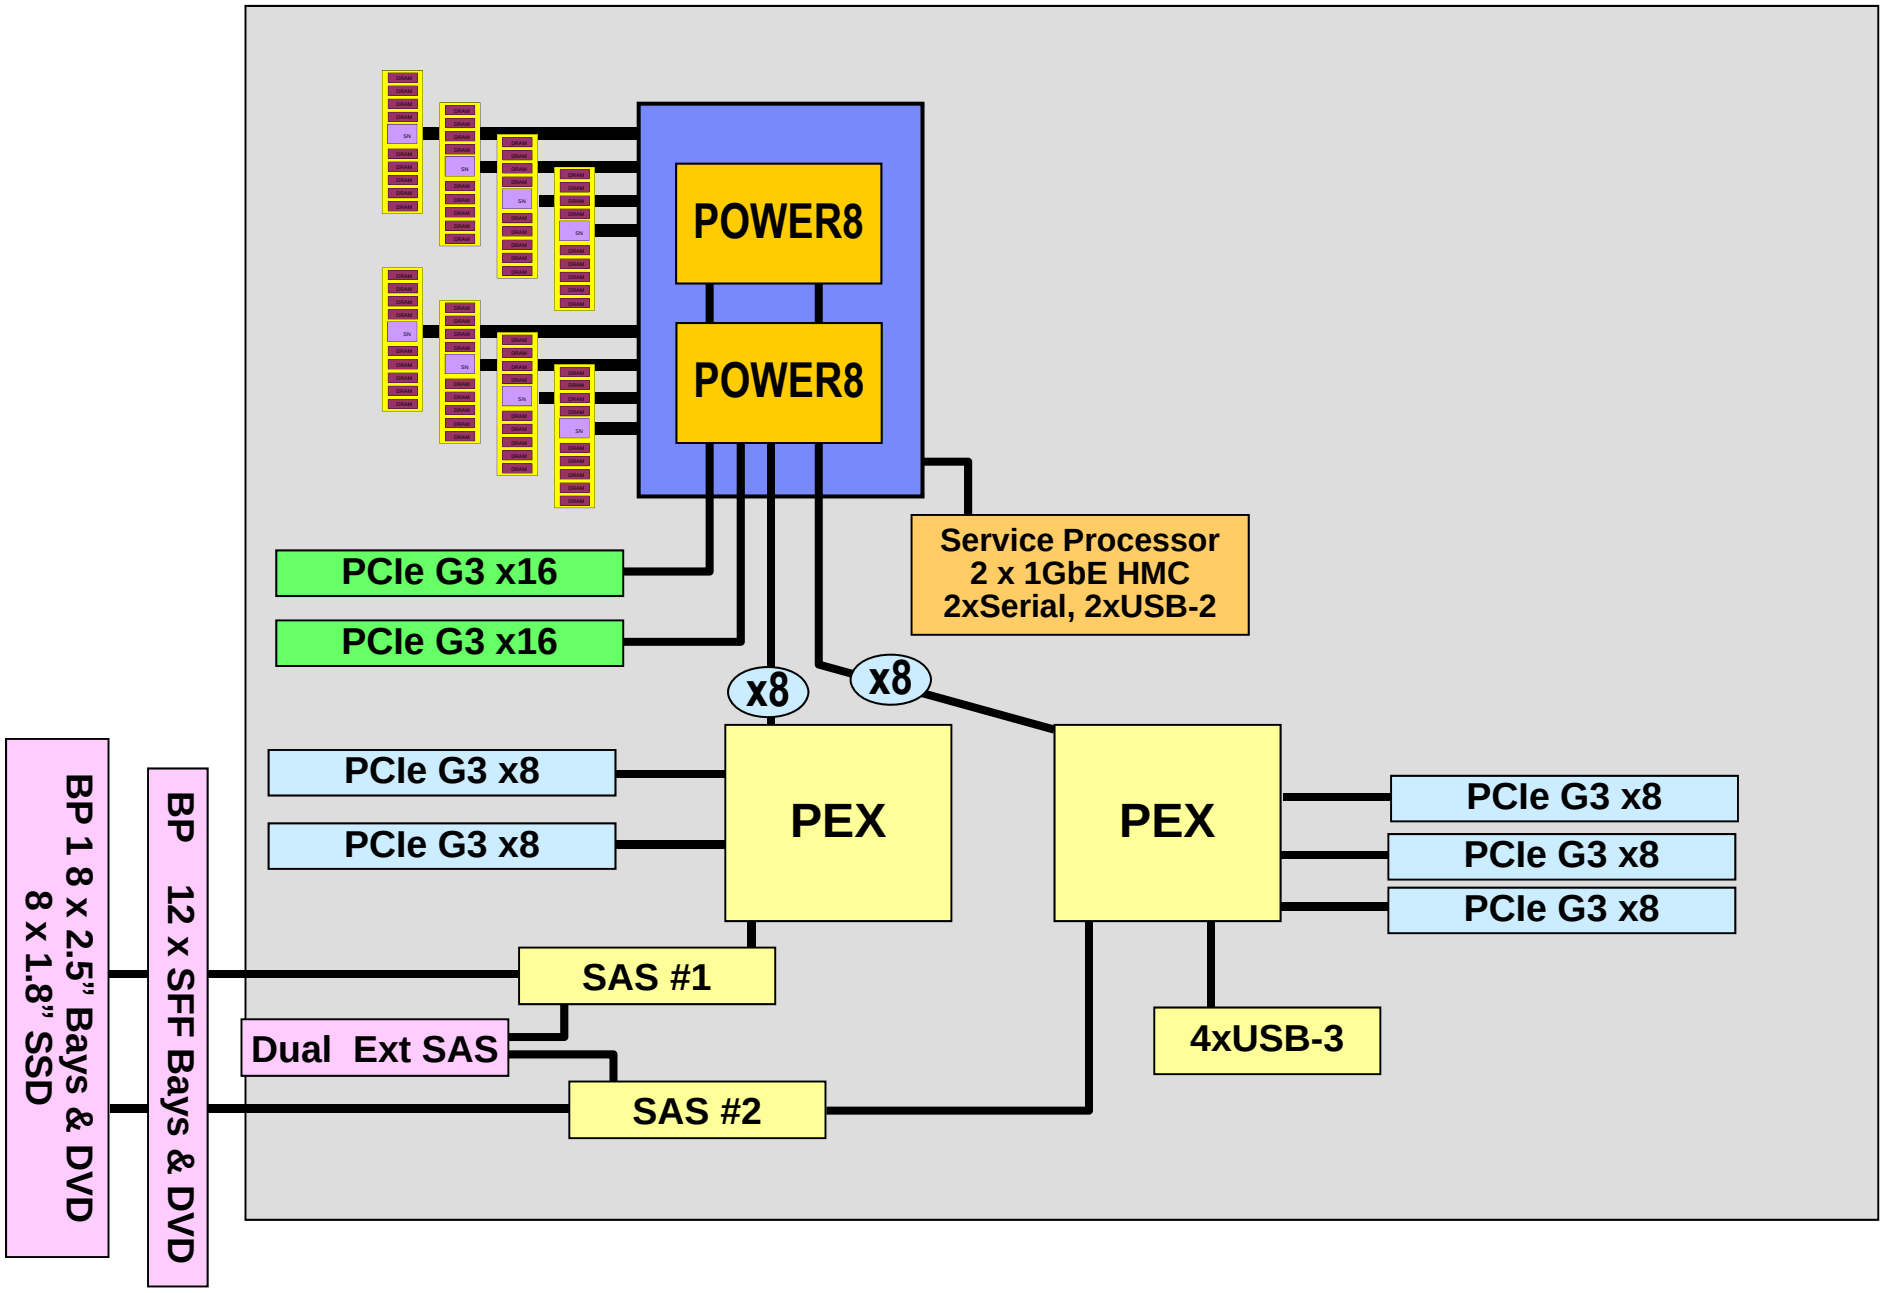

POWER8 2S4U Planar Layout

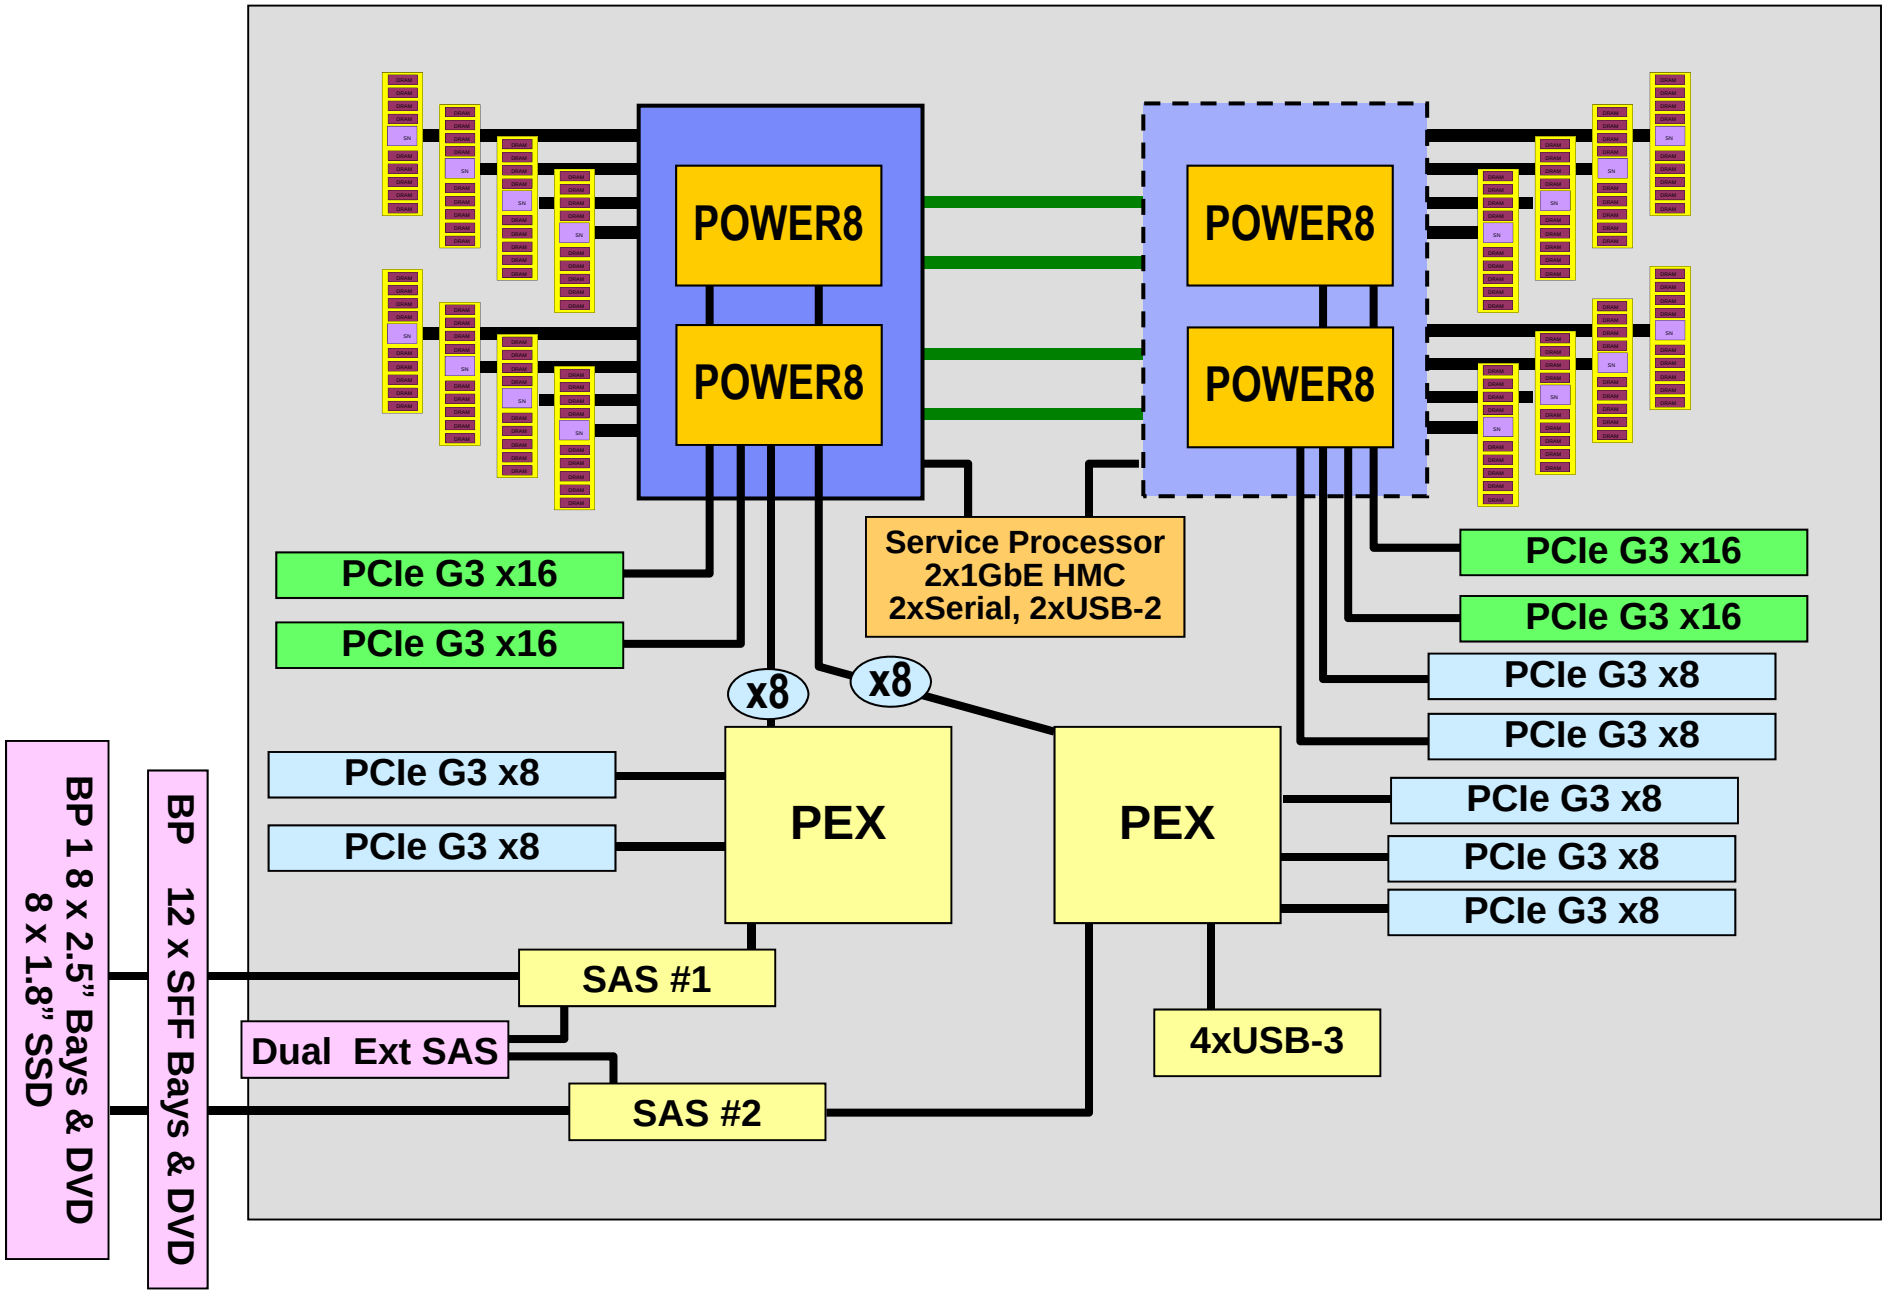

Top View POWER8 2S4U System

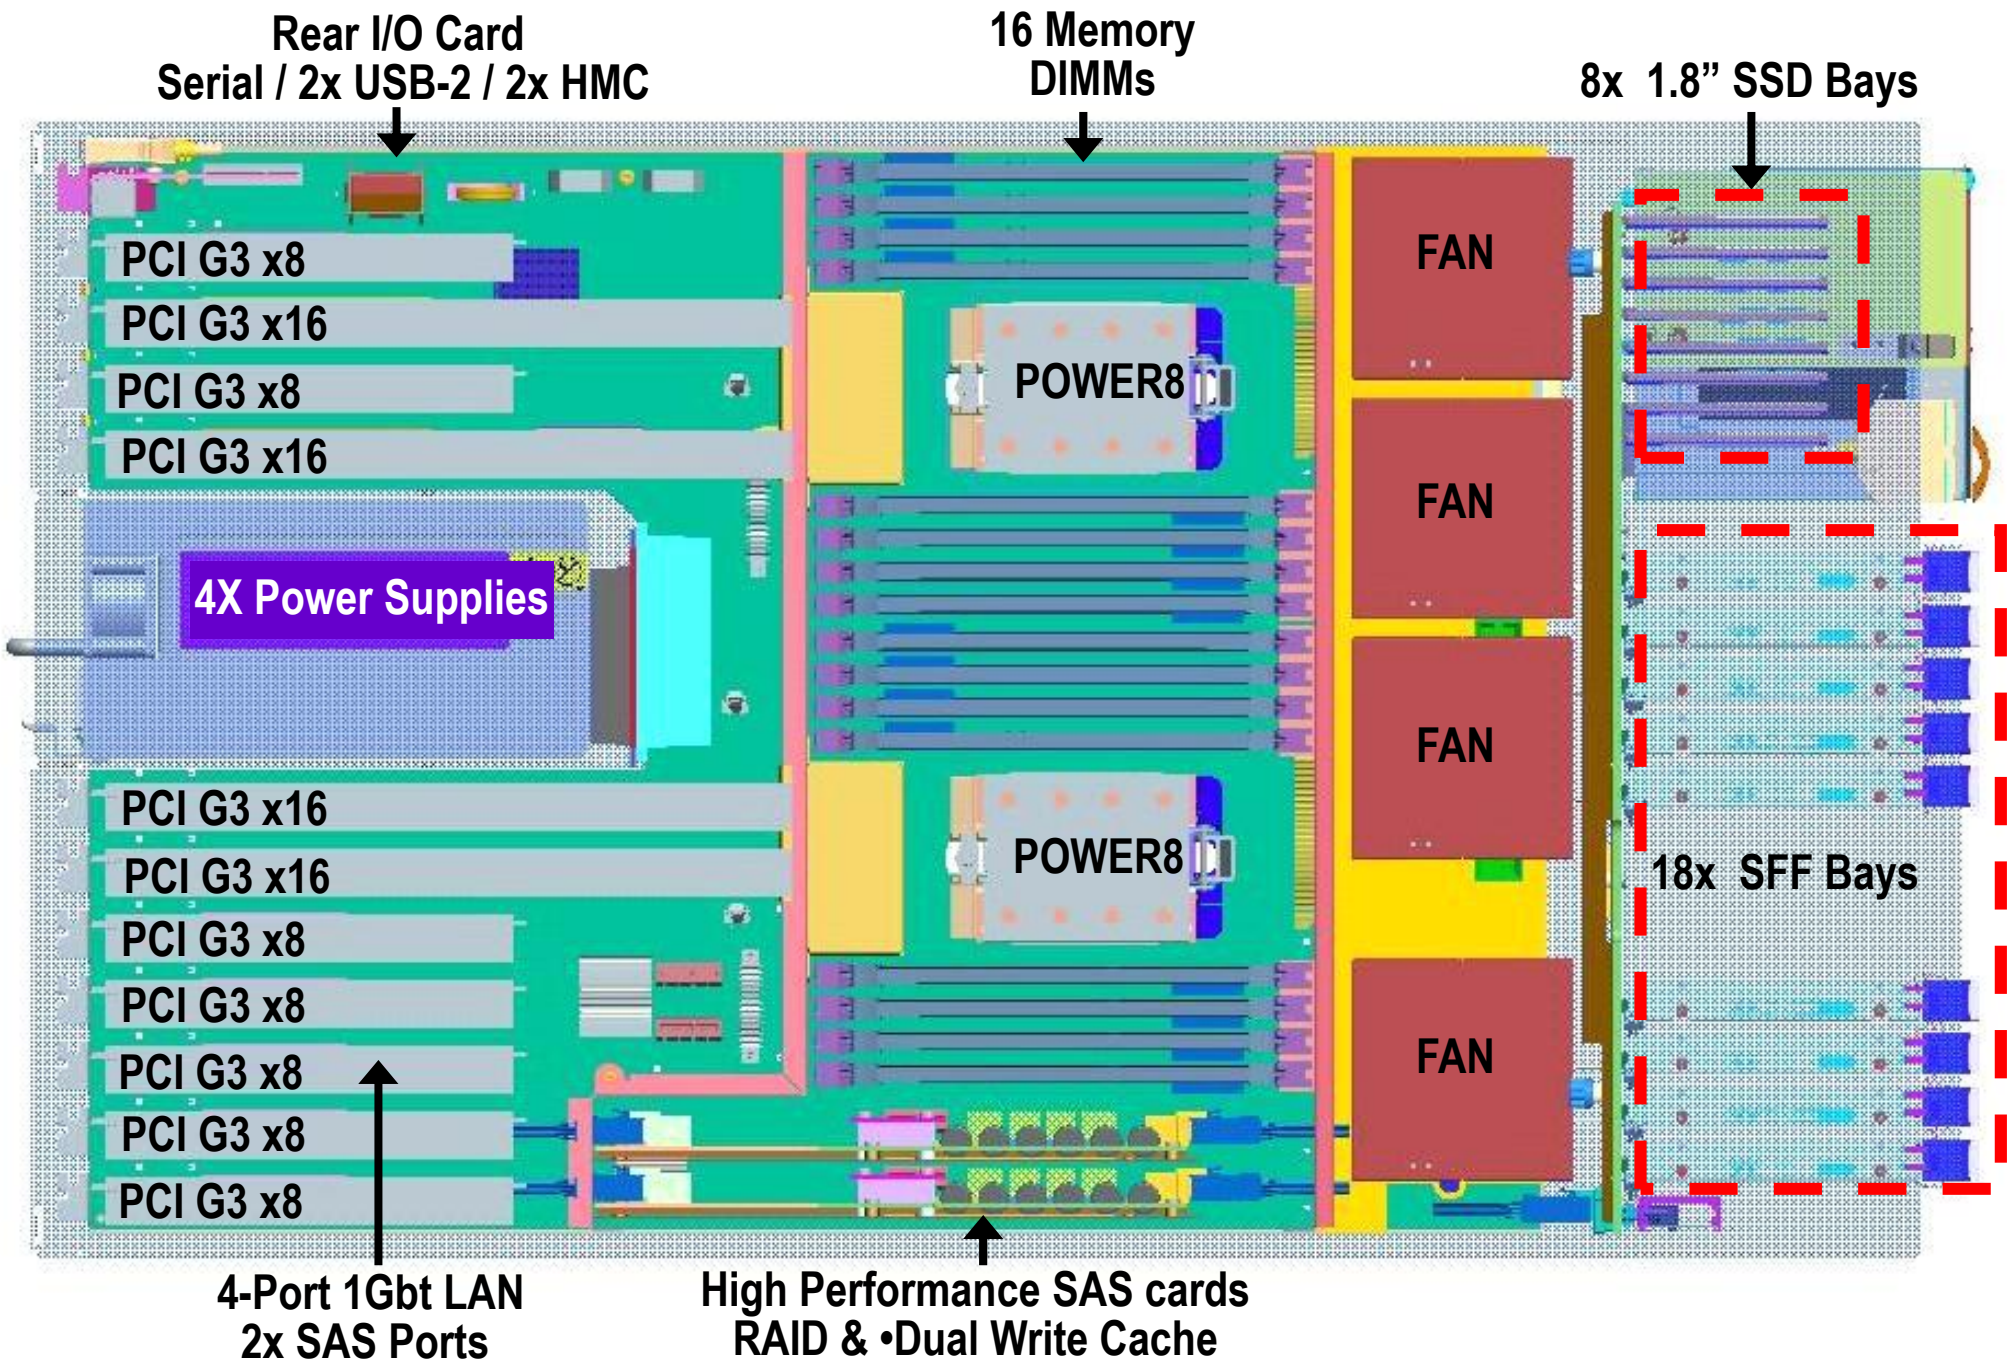

Power 4U Scale-Out Comparison

	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		

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
Add'l storage backplane option	28 Apr	17 July	Sept
SFF-3 146/139GB 15k & 300/283GB 10k HDD	28 Apr	27 May	25 July

Performance

Power 740+ / POWER8 S824 rPerf Comparisons

Performance per Core

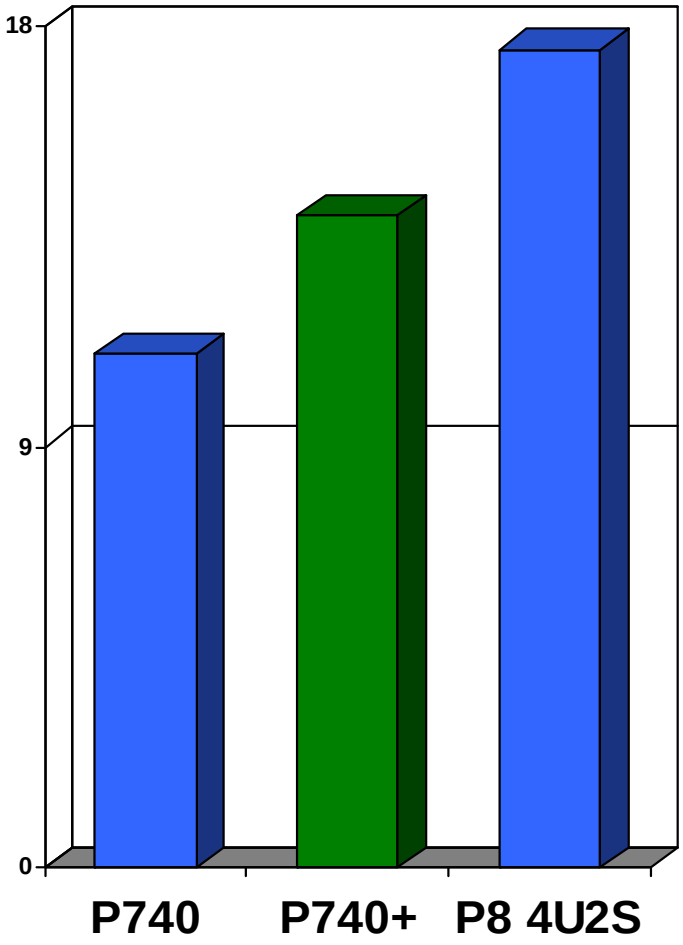

Performance per Socket

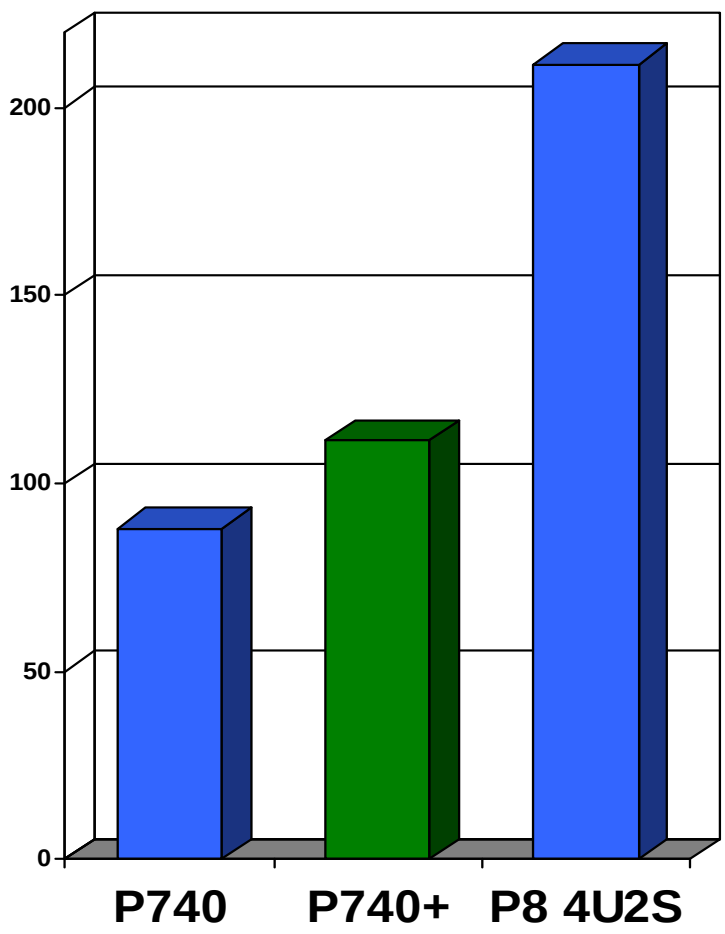

Performance per KW

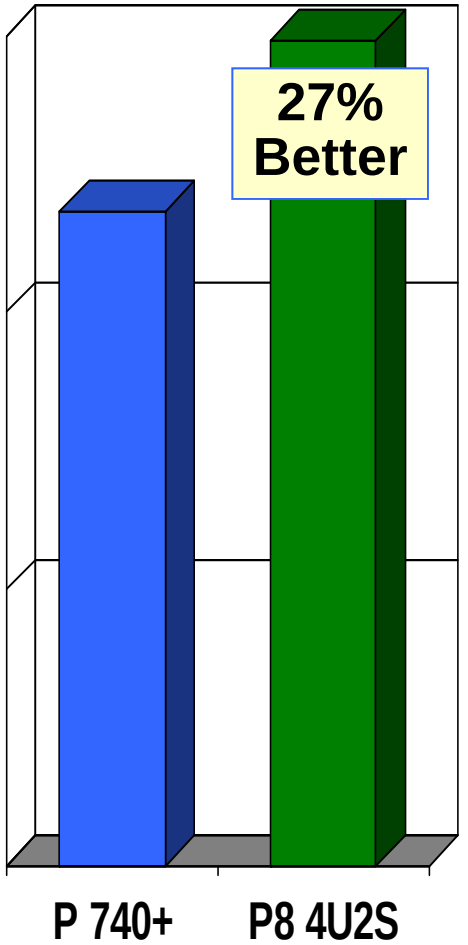

Power 740 / POWER8 S824 CPW Comparisons

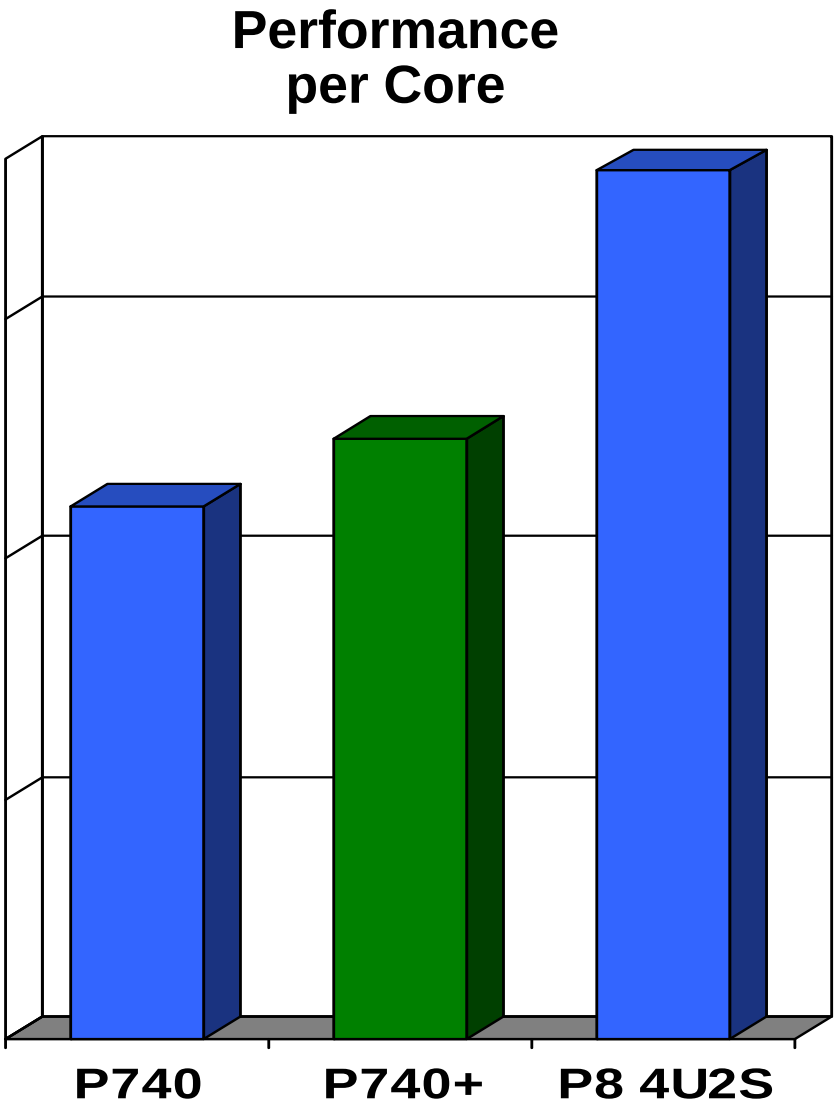

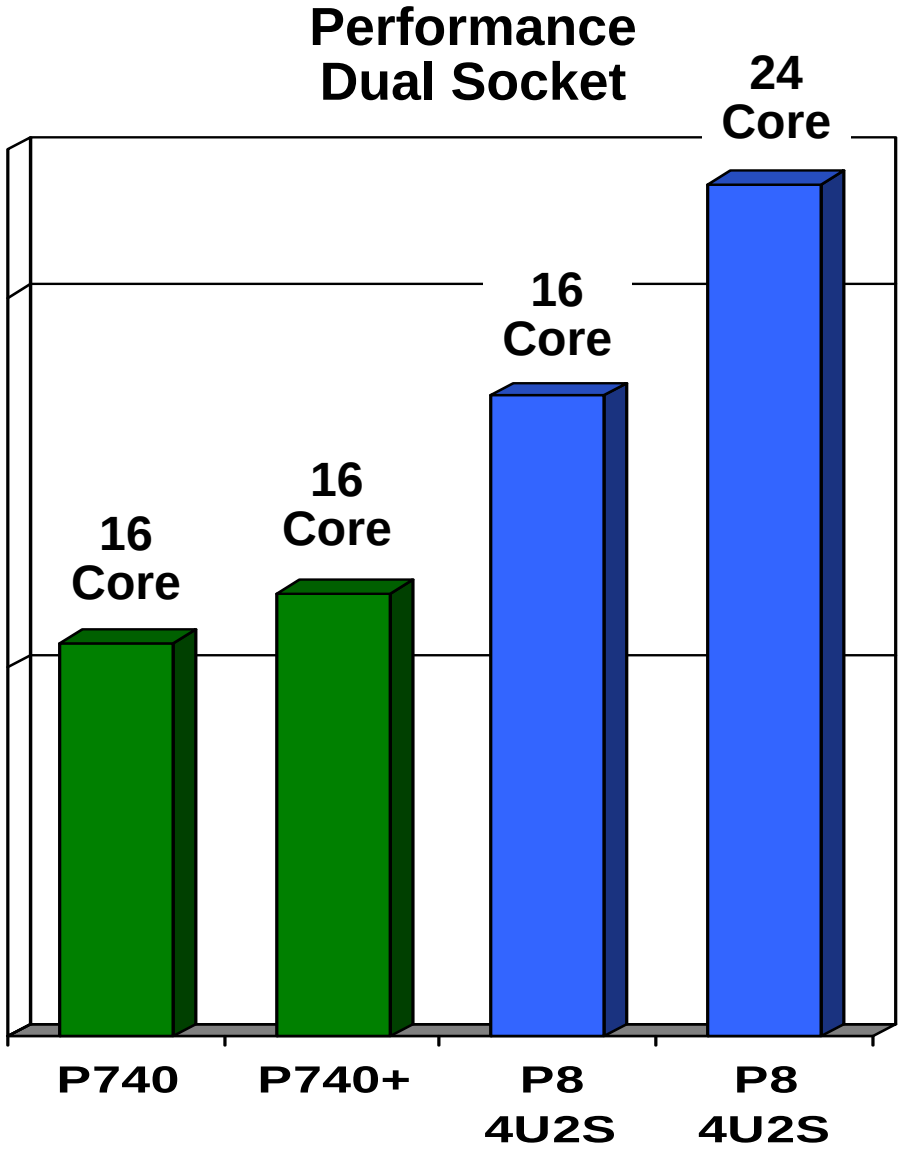

Power 750+ vs POWER8 4U2S

Performance

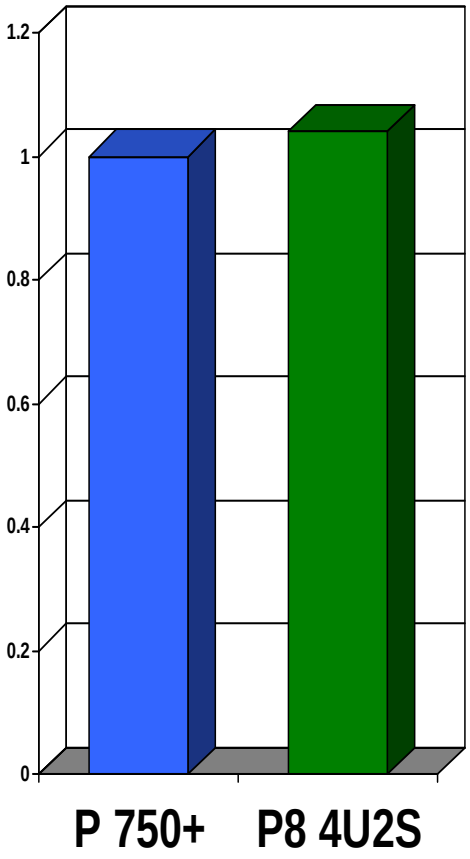

POWER8 ≥ Performance

Total # Cores

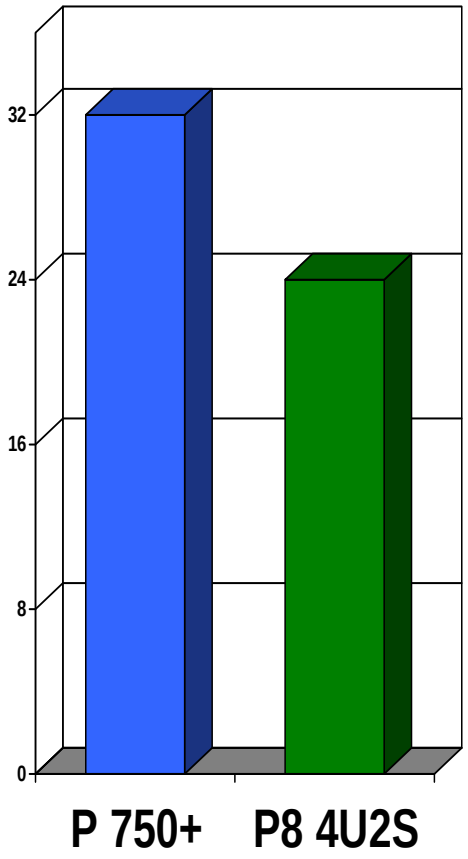

Fewer cores means Lower Software Costs

Does not reflect lower PVU savings

TCA

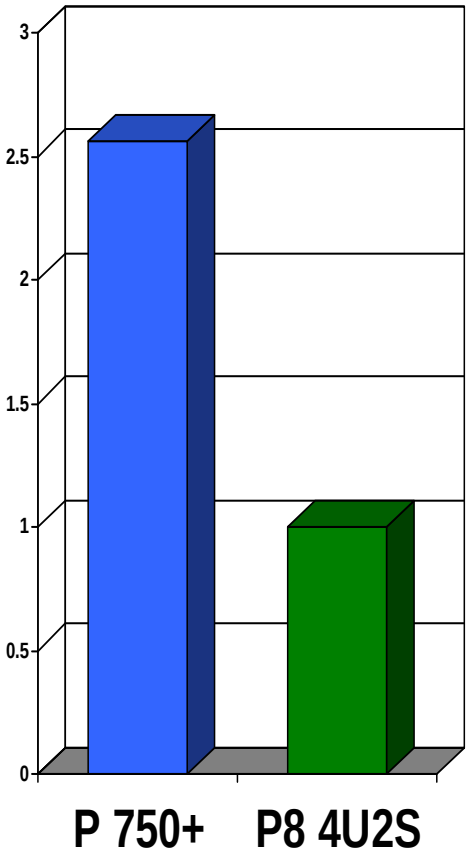

Lower investment costs

Performance per $

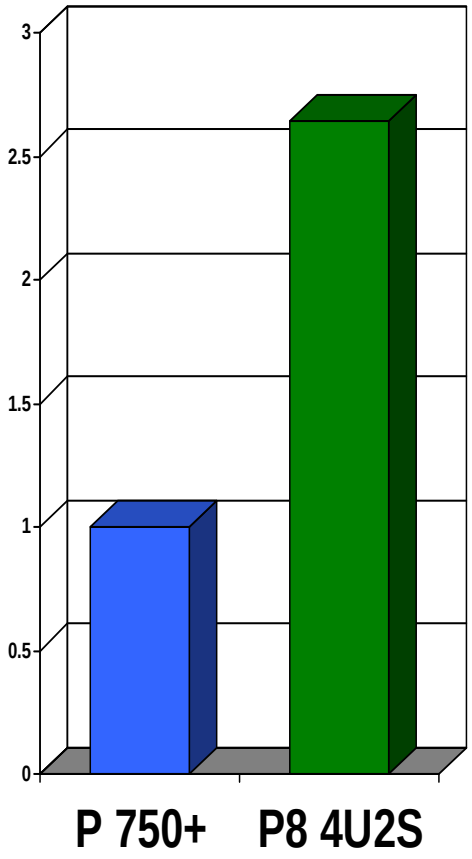

Better performance per dollar spent

POWER8 4U Scale-Out Comparison

	Power 750	Power System S824
Processor	POWER7+	**POWER8**
Sockets	4	**2**
Maximum Cores	32	**24**
Maximum Memory	1 TB @ 1066 MHz DDR2	**1 TB @ 1600 MHz DDR3**
Memory Cache	No	**Yes**
Total Memory Bandwidth	272 GB/sec Four Sockets	**384 GB/sec** **Two Sockets**
Memory DRAM Spare	No	**Yes**
PCIe Gen3 Support	No	**Yes**
CAPI Support	No	**Yes**
PCIe slots in system unit	6 PCIe Gen2 FH	**11 PCIe Gen3 FH**
PCIe Hot Plug Support	Yes	**Yes**
PCIe I/O drawers	Up to 4 drawers, Gen1 slots	**N/A - SOD PCIe Gen3**
Max IO bandwidth	78.8 GB/sec	**192 GB/sec**
Ethernet ports	Two 10GBt & Dual 1Gbt	**Quad 1 Gbt uses 1 PCIe slot**
SAS bays in system unit	6 SFF-1 bays (3+3)	**12 SFF-3 bays (6+6)** **Or 18 SFF-3 bays + 8 SSD bays**
Integrated SAS write cache	Opt 175MB	**Opt 3GB effectively 7GB**
Easy Tier function	No	**Yes**
SAS bays in system unit	6	**12**
Integrated split backplane	Yes (3 + 3)	**Yes (6 + 6)**
NIC SRIOV Support	No	**SOD**
Service Processor	Generation 1	**Generation 2**
Performance	~1	**~1.05**

POWER8 System Performance

P8 S824

P5+ 595

P4 690

POWER8 Turbo

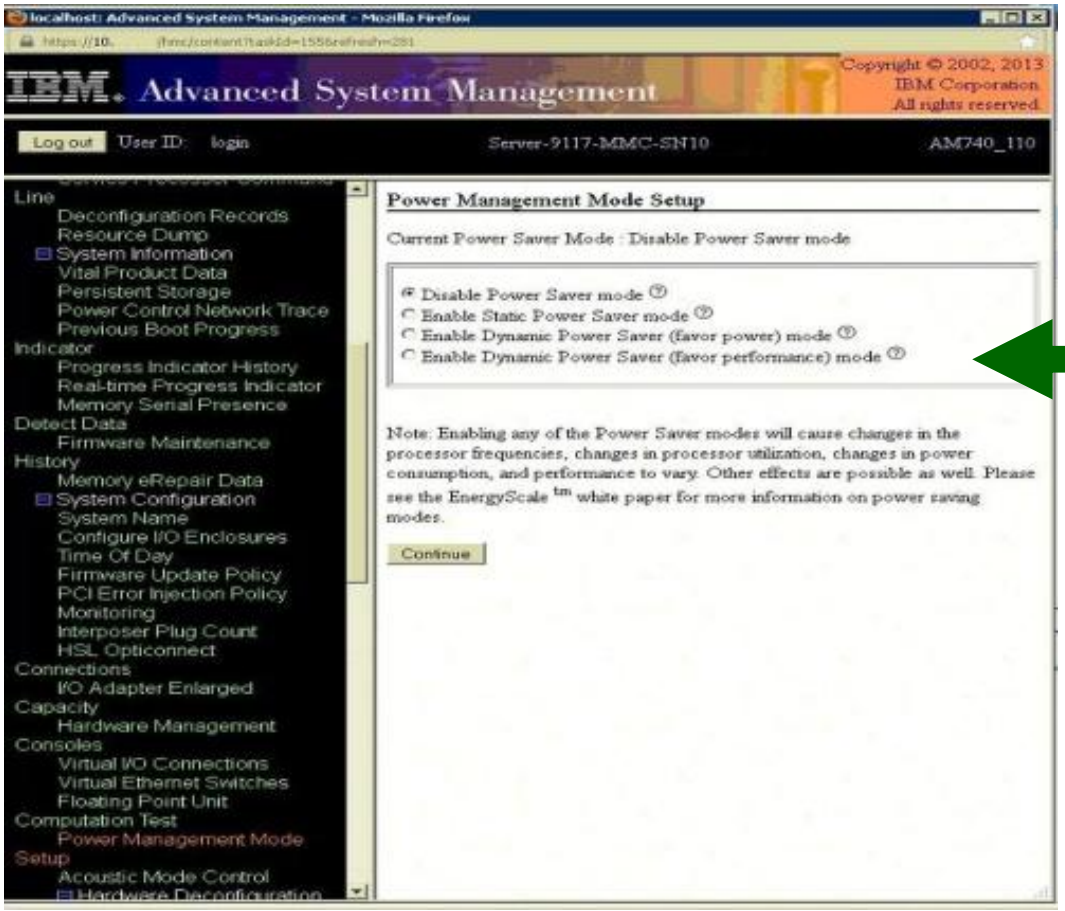

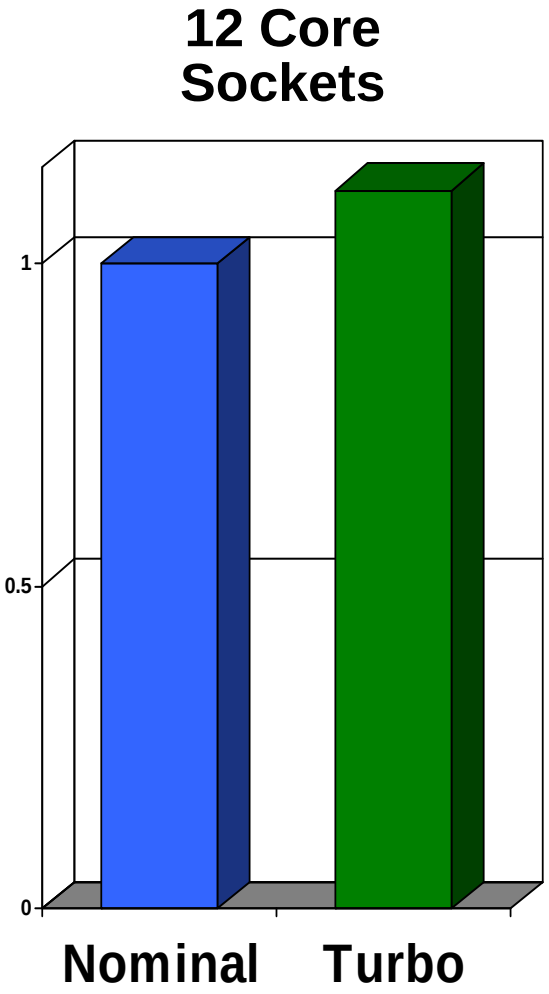

- **Configure via ASMI menu**
- **Potential increase in processor frequency: ~ 11+%**
- **Requires firmware 810 (POWER8 support)**
- **CPW & rPerf measured using nominal**

CPW

S814 (1 socket)

- 6-core 3.0 GHz 59,500
- 8-core 3.7 GHz 85,500

Measured using SMT8

SMT4 would be somewhat lower

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

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

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

+40%

+50%

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

+40%

- 8-core 4.1 GHz 94,500
- 16-core 4.1 GHz 173,500

+60%

- 12-core 1-socket not offered
- 24-core 3.5 GHz 230,500

+90%

rPerf (with SMT8)

S814 (1 socket)

- 6-core 3.0 GHz 97.5
- 8-core 3.7 GHz 143.9

S824 (1 or 2 socket)

- 6-core 3.8 GHz 120.8
- 12-core 3.8 GHz 235.6
- 8-core 4.1 GHz 166.0
- 16-core 4.1 GHz 323.6
- 12-core 1-socket not offered
- 24-core 3.5 GHz 421.8

S822 (1 or 2 socket)

- 6-core 3.8 GHz 120.8
- 12-core 3.8 GHz 235.6
- 10-core 3.4 GHz 177.8
- 20-core 3.4 GHz 346.7

rPerf – Multiple SMT Levels

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

IBM POWER8 versus Intel x86 “Ivy Bridge”

Published Industry Standard Benchmarks	IBM Power S824 (3.5GHz)	Intel x86 “Ivy Bridge”	POWER8 vs x86
	POWER8	Xeon E5-2697 v2	Ratio
	24 cores	24 cores	
SAP SD 2-Tier [1]	21,212	10,240	2.10x
SPECint_rate2006 [2]	1,750	1,020	1.70x
SPECfp_rate2006 [2]	1,370	734	1.90x
SPECjEnterprise2010 [3]	22,543	11,260	2.00x

1) IBM Power System S824 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 4 processors / 24 cores / 96 threads, POWER8; 3.52GHz, 512 GB memory, 21,212 SD benchmark users, running AIX® 7.1 and DB2® 10.5, dialog response: 0.98 seconds, line items/hour: 2,317,330, dialog steps/hour: 6,952,000 SAPS: 115,870 database response time (dialog/update): 0.011 sec / 0.019sec, CPU utilization: 99%, Certification #: * Results valid as of 3/24/14. * Certification # not available at press time. Source: http://www.sap.com/benchmark.
(1.1) Fujitsu RX300 S8 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 2 processors / 24 cores / 48 threads. Intel Xeon E5-2697 processor 2.70 GHz, 256 GB memory, 10,240 SD benchmark users, running Windows Server 2012 SE and SQL Server 2012, Certification #: 2013024

2) IBM Power S824 results submitted to SPEC, waiting for approval. Supermicro SuperServer 6027AX-TRF (X9DAX-iF, Intel Xeon E5-2697 v2, 2.70 GHz) results. Source: http://www.spec.org

3) IBM WebSphere Application Server V8.5.5.2 and DB2 10.5 on IBM Power S824 result of 22,543.34 published on Apr 22, 2014. Oracle Weblogic Server Standard Edition Release 12.1.2 and Oracle Database 12c on Oracle Sun Server X4-2 result of 11,259.88 published on Sep 23, 2013. Source:http://www.spec.org

PCIe Slots

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	S822	S822L S822
	8286-41A	8286-42A	8286-42A	8284-22A	8247-22L 8284-22A
	1S	Only 1S in 2S box	2S	Only 1S in 2S box	2S
Total PCIe slots x16 x8	7 2 x16 5 x8	7 2 x16 5 x8	11 4 x16 7 x8	6 2 x16 4 x8	9 4 x16 5 x8
Required LAN adapter (available for client use)	1 x8	1 x8	1 x8	1 x8	1 x8
PCIe slots after required LAN adapter	6 2 x16 4 x8	6 2 x16 4 x8	10 4 x16 6 x8	5 2 x16 3 x8	8 4 x16 4 x8
However if use high performance, expanded function backplane	-1 x8	-1 x8	-1 x8	-1 x8	-1 x8
PCIe slots after required LAN and if using high performance backplane	5 2 x16 3 x8	5 2 x16 3 x8	9 4 x16 5 x8	4 2 x16 2 x8	7 4 x16 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 Gen3

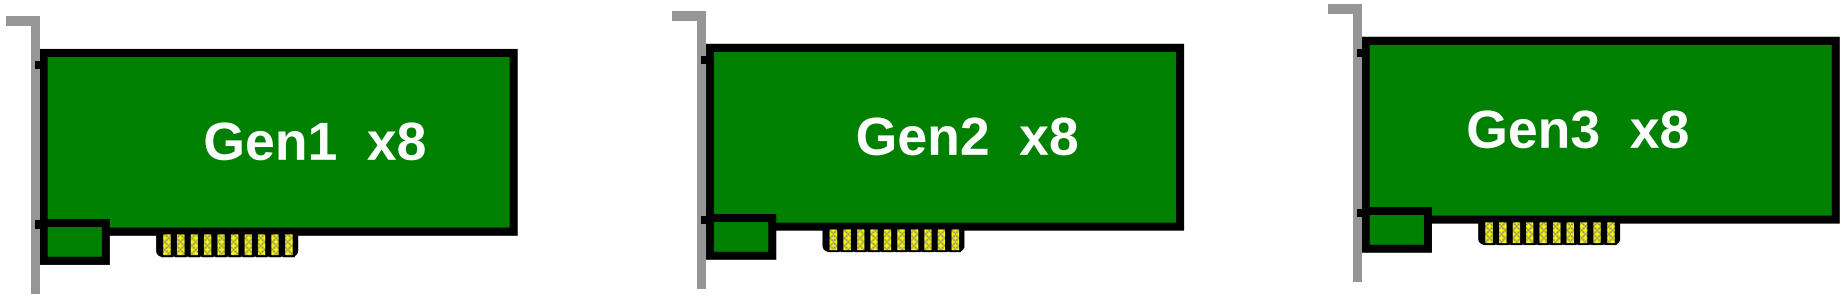

Though these cards physically look the same … and fit in the same slots
Gen3 cards/slots have up to 2X more **bandwidth** than Gen2 cards/slots
Gen3 cards/slots have up to 4X more **bandwidth** than Gen1 cards/slots

- More virtualization
- More consolidation
- More ports per adapter

} saving PCI slots and I/O drawers

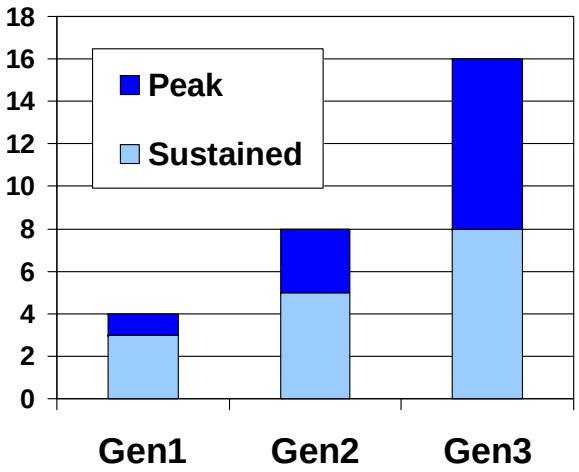

A Gen1 x8 PCIe adapter has a theoretical max (peak) bandwidth of 4 GB/sec.
A Gen2 x8 adapter has a peak bandwidth of 8 GB/sec.
A Gen3 x8 adapter has a peak bandwidth of 16 GB/sec.

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)

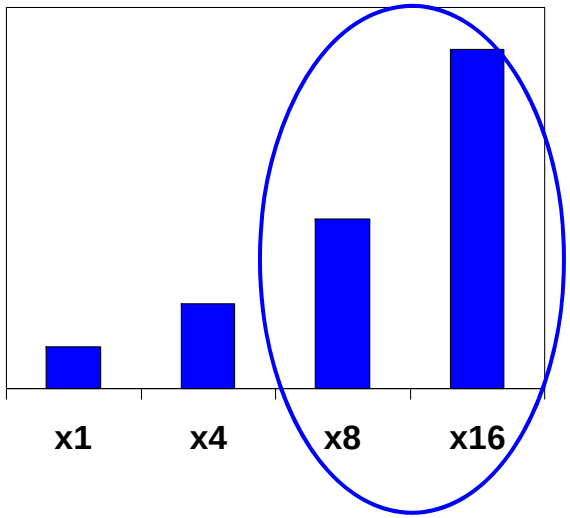

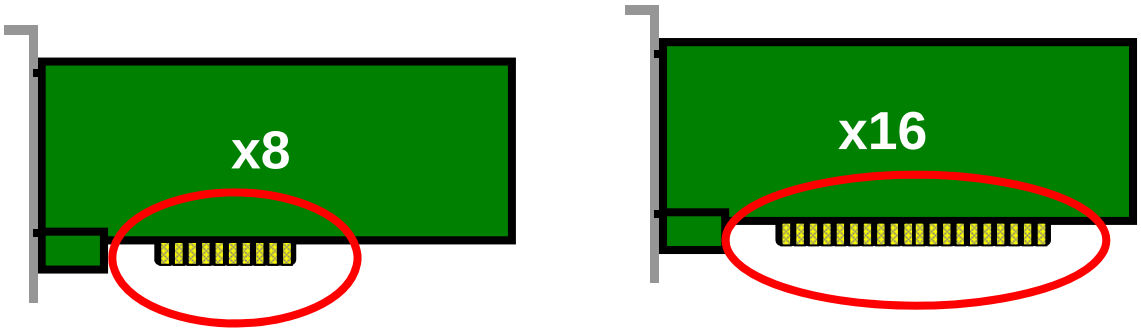

x16 slot/card has more connections than a x8 slot/card

"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

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/revised" 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: 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 (page 1)

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 (page 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**

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

PCIe Adapters NOT Supported

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	

New POWER8 PCIe Adapters

- LAN
  - 2-port 40GbE NIC & RoCE QSFP+ cables
  - 2-port 10GbE RJ45 (10GBaseT) only $950
  - 4-port 10GbE+1GbE Copper twinax & RJ45 only $1200
  - 4-port 10GbE+1GbE SR optical & RJ45 only $1600

  No NIM
  - SOD for NIC SR-IOV two specific 10GbE adapters

- USB-3 4-port (old USB-2 not supported)

- SAS PCIe3
  - 4-port Tape/DVD only $1000
  - Reprice 4-port zero cache HDD/SSD to $1000

- Graphics adapter for future RHEL 7 (Planned June announce)

Prices are USA PRELIMINARY list prices for a Power S822 and are subject to change. Reseller prices may vary.

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

Prices are USA list prices for a Power S822 and are subject to change. Reseller prices may vary.

New I/O Announcing 8 April for POWER7/POWER+

PCIe Adapters

- SAS expandeded usage
  - PCIe3 4-port zero cache on more models
  - PCIe3 4-port huge cache on more models
  - PCIe2 3-port large cache on more POWER7 models

- Easy Tier function for POWER7+ models with #EJ0L huge cache SAS adapter for AIX/Linux/VIOS

- IBM Flash Adapter 90 support on more models

Prices are USA list prices for a Power 720 and are subject to change. Reseller prices may vary.

Storage Backplanes

POWER8 2U Storage Backplane Options

Base: 12 SFF HDD/SSD Bays
- **RAID 0,5,6,10 / JBoD**

Optional: Split Disk Backplane = 6 + 6

Optional: High Performance RAID 0,5,6,10,5T2,6T2,10T2
- **8 SFF Bays (HDD / SSD) and 6 SSD**

POWER8 4U Storage Backplane Options

Must select one →	#EJ0N		#EJ0N+EJ0S		#EJ0P *	
	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
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 #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 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)

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 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**

AIX / IBM i / Linux	Yes		Yes		Yes
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 with #EJ0U**

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**

Controller Write Cache Value

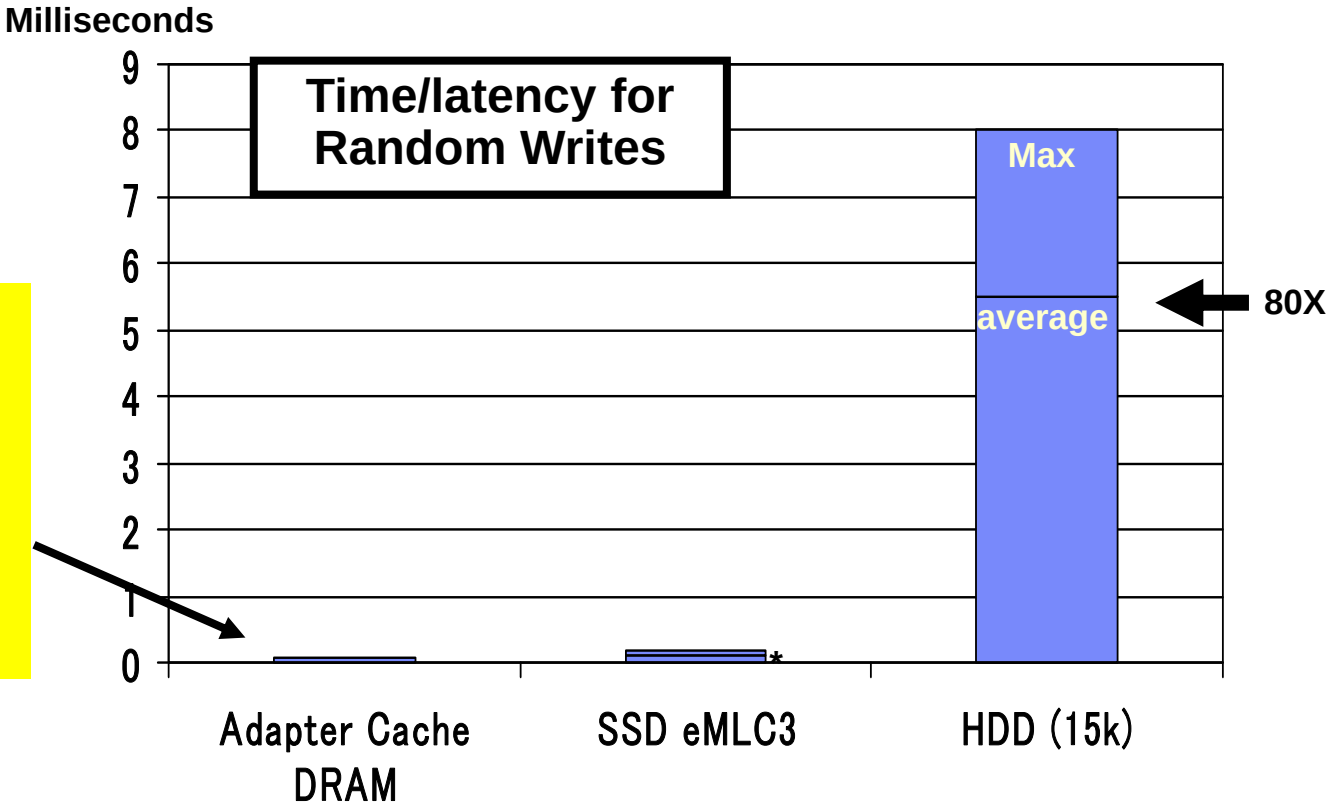

Controller cache is

- **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

Controller cache is

- **Up to 1.5 to 2.5X faster than SSD**

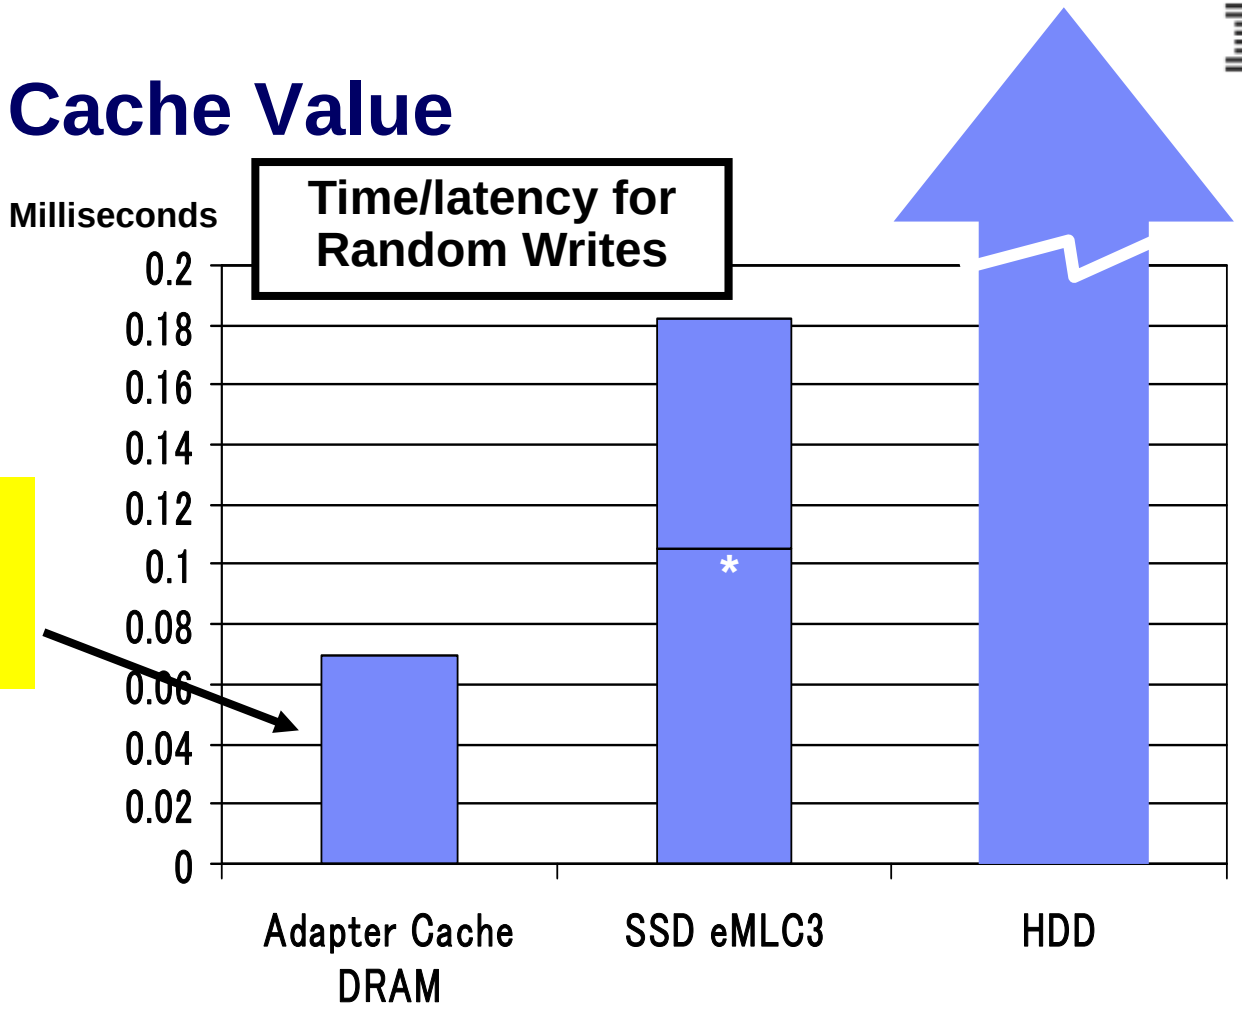

Value depends on the amount/percentage of Writes

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 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.

SAS Bays for POWER8 Servers

New SSF-3

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

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

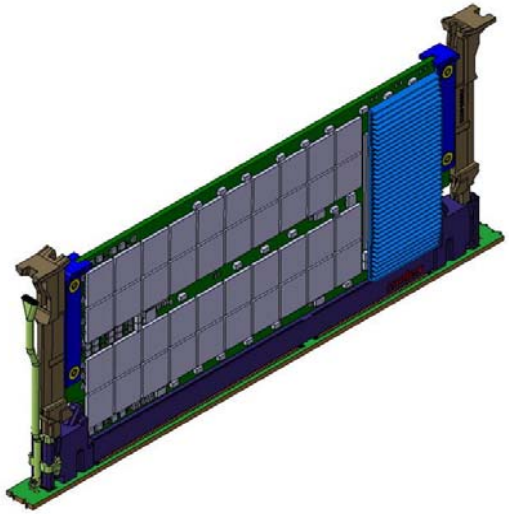

POWER8 Memory

POWER8 Memory Card

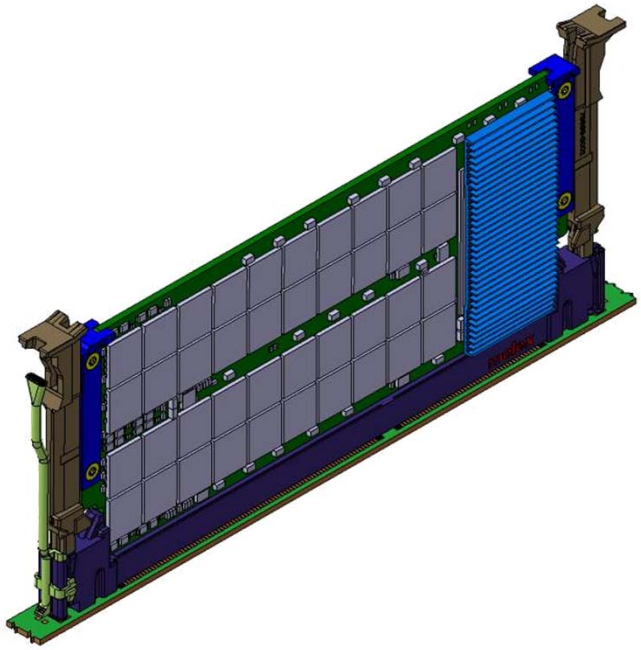

POWER8 Memory Card

- Capacity: 16GB / 32 GB / 64 GB
- 1600 MHz
- Memory Sparing - RAS improvement
- 8 Cards per socket (Scale-Out Systems)
- Low Profile

Memory slots & Memory DIMMs

	4U		2U	
	1S 4U	2S 4U	1S 2U	2S 2U
Max number DIMMs	8	16	8	16
Min number DIMMs	1	2	1	2
Max GB memory	512GB	1TB	512GB	1TB

- **Memory DIMMs are all DDR3 1.6GHz**
- **Memory is not hot plug**
- **1 memory feature = 1 DIMM**
- **Memory plugged in pairs – 1S exception**
- **2U & 4U server memory physically packaged differently**

Memory DIMMs
16 GB
32 GB
64 GB

Activations: all memory 100% activated. CoD/CUoD not implemented.

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 Mandatory #EJ0Z feature. Ports take a PCIe slot	2 rear No feat code. Ports do not take a PCIe slot

* 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.

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

POWER8 OS Support

Compatible Mode Architecture

P6 MODE (*and P6+ Mode*)*	P7 MODE *(No P7+ Mode)*	P8 MODE
2-Thread SMT	4-Thread SMT, *IntelliThreads*	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

IBM i Cores/Partition – SMT – Threads

	IBM i 7.1 TR8
POWER7	Max published cores/partition = 32 cores Max cores/partition = 64 cores **with Lab Services approval RPQ** SMT = ST, SMT2, STM4 (not SMT8) Threads = up to 128 threads / single partition (256 threads with Lab Services RPQ)
POWER8	Max published cores/partition = 24 cores (32 core theoretical announce) Max cores/partition = 64 cores(when announced) **with Lab Services approval RPQ** SMT = ST, SMT2, SMT4, **SMT8** Threads = up to **192** (with 24-core server – theoretical max of **256**) threads / single partition

	IBM i 7.2
POWER7	Max published cores/partition = 32 cores Max cores/partition = **96 cores** **with Lab Services approval RPQ** SMT = ST, SMT2, STM4 (not SMT8) Threads = up to 128 threads / single partition **(384 threads** with Lab Services RPQ)
POWER8	Max published cores/partition = 24 cores (**48 cores theoretical announce)** Max cores/partition = **96** cores(when announced) **with Lab Services approval RPQ** SMT = ST, SMT2, SMT4, **SMT8** Threads = **up to 192** (with 24-core server – theoretical max of **768**) threads / single partition

Linux Plans

	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

AIX Levels

	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

Power Virtualization Options

PowerKVM

Initial Offering: Q2 2014

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

PowerVM

Initial Offering: 2004

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

PowerVM vs PowerKVM Comparison

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

POWER6 / POWER7 / POWER8 Partition Mobility

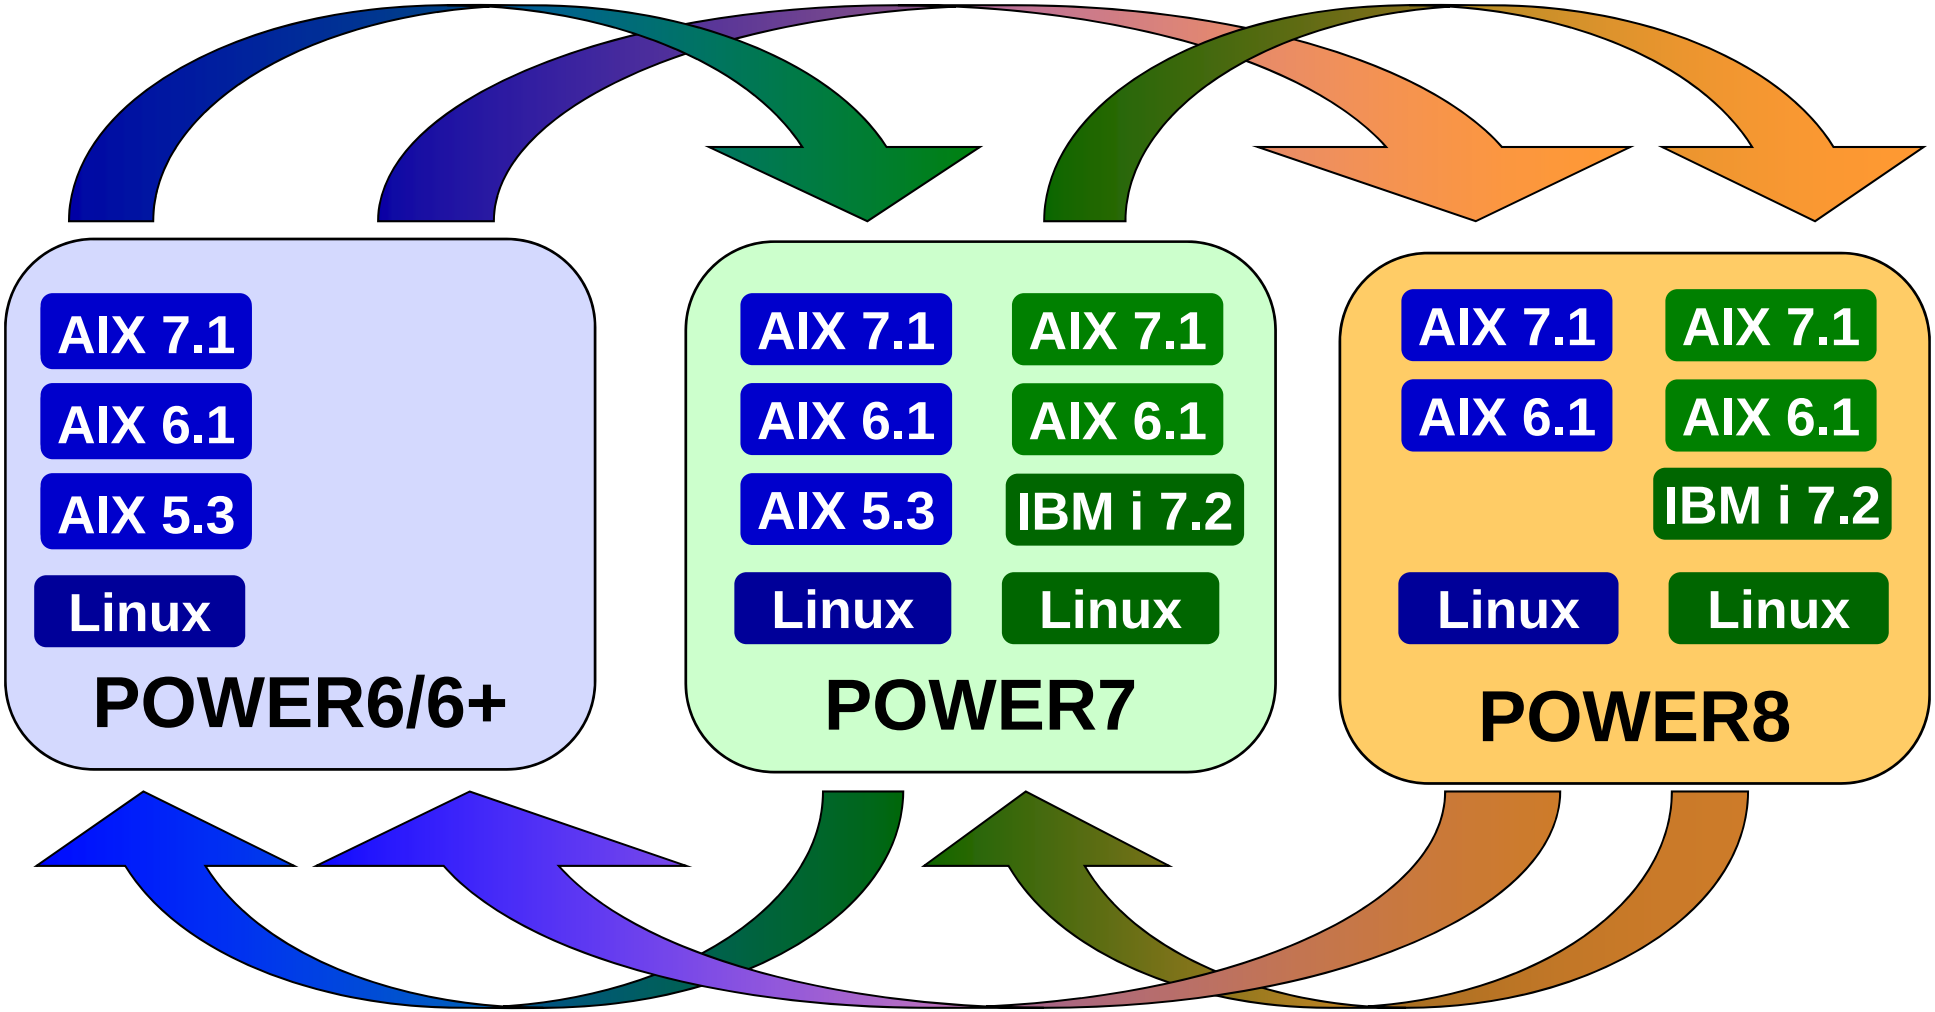

Leverage POWER6 / POWER7 Compatibility Modes

LPAR Migrate between POWER6 / POWER7 / POWER8 Servers

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

HMC Model 7042-CR8

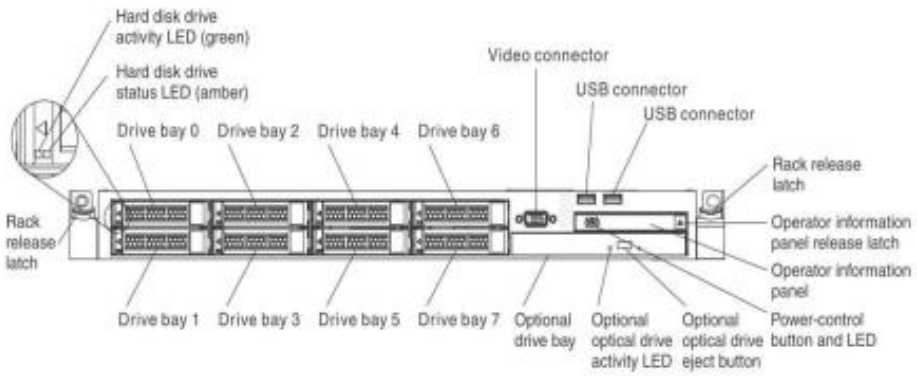

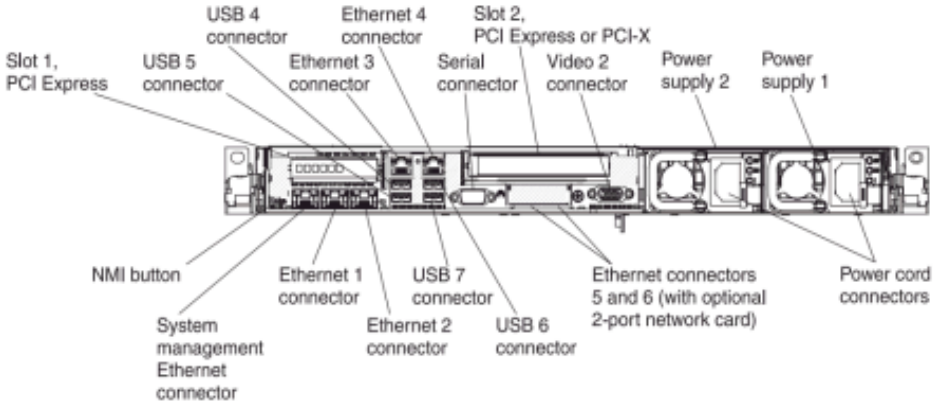

7042-CR8 is System x model 7914 Ref3

- 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

New I/O
28 April

New I/O
8 April

The End…