Power Systems Hardware:

Today and Tomorrow

November 2015

Mark Olson

olsonm@us.ibm.com

POWER8 Chip

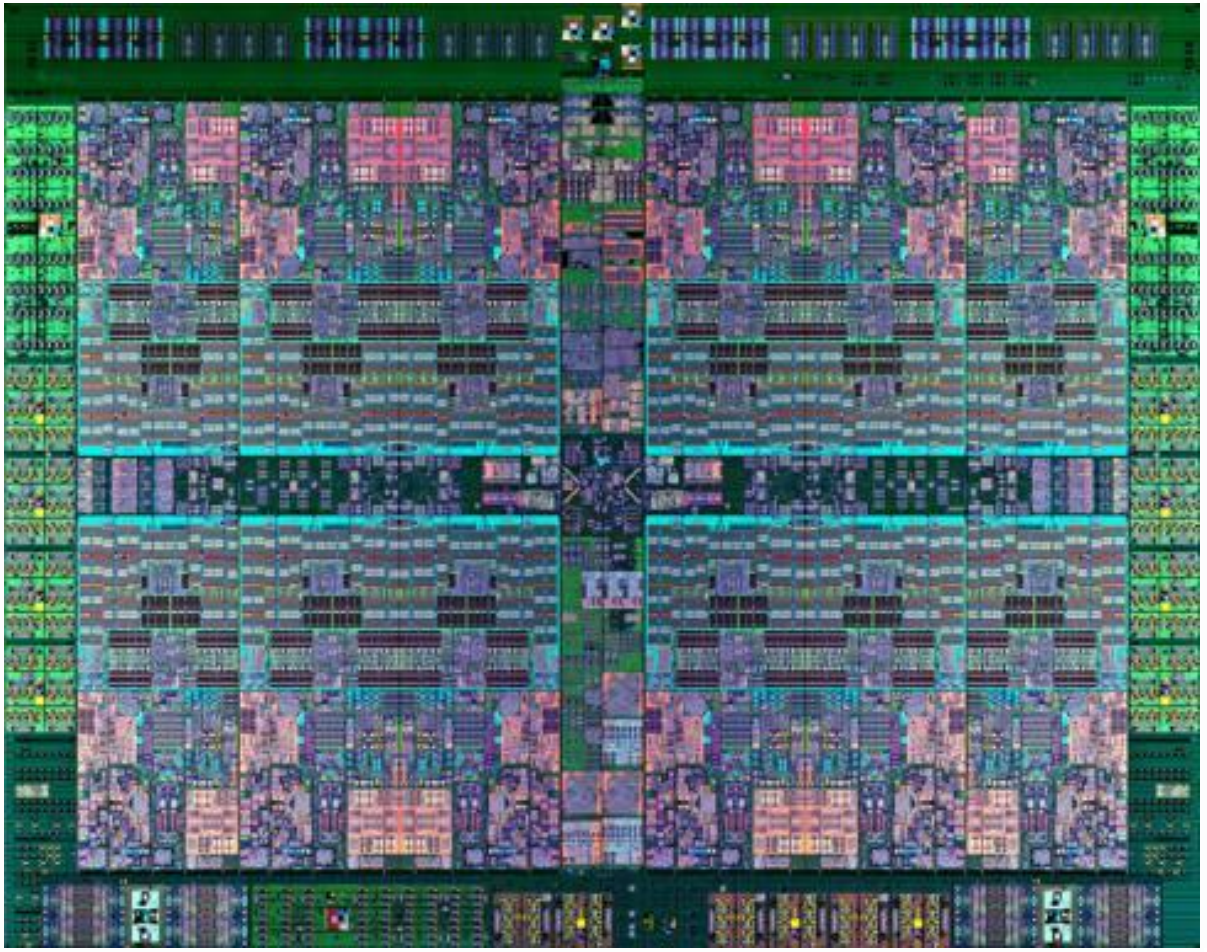

Processor Technology Roadmap

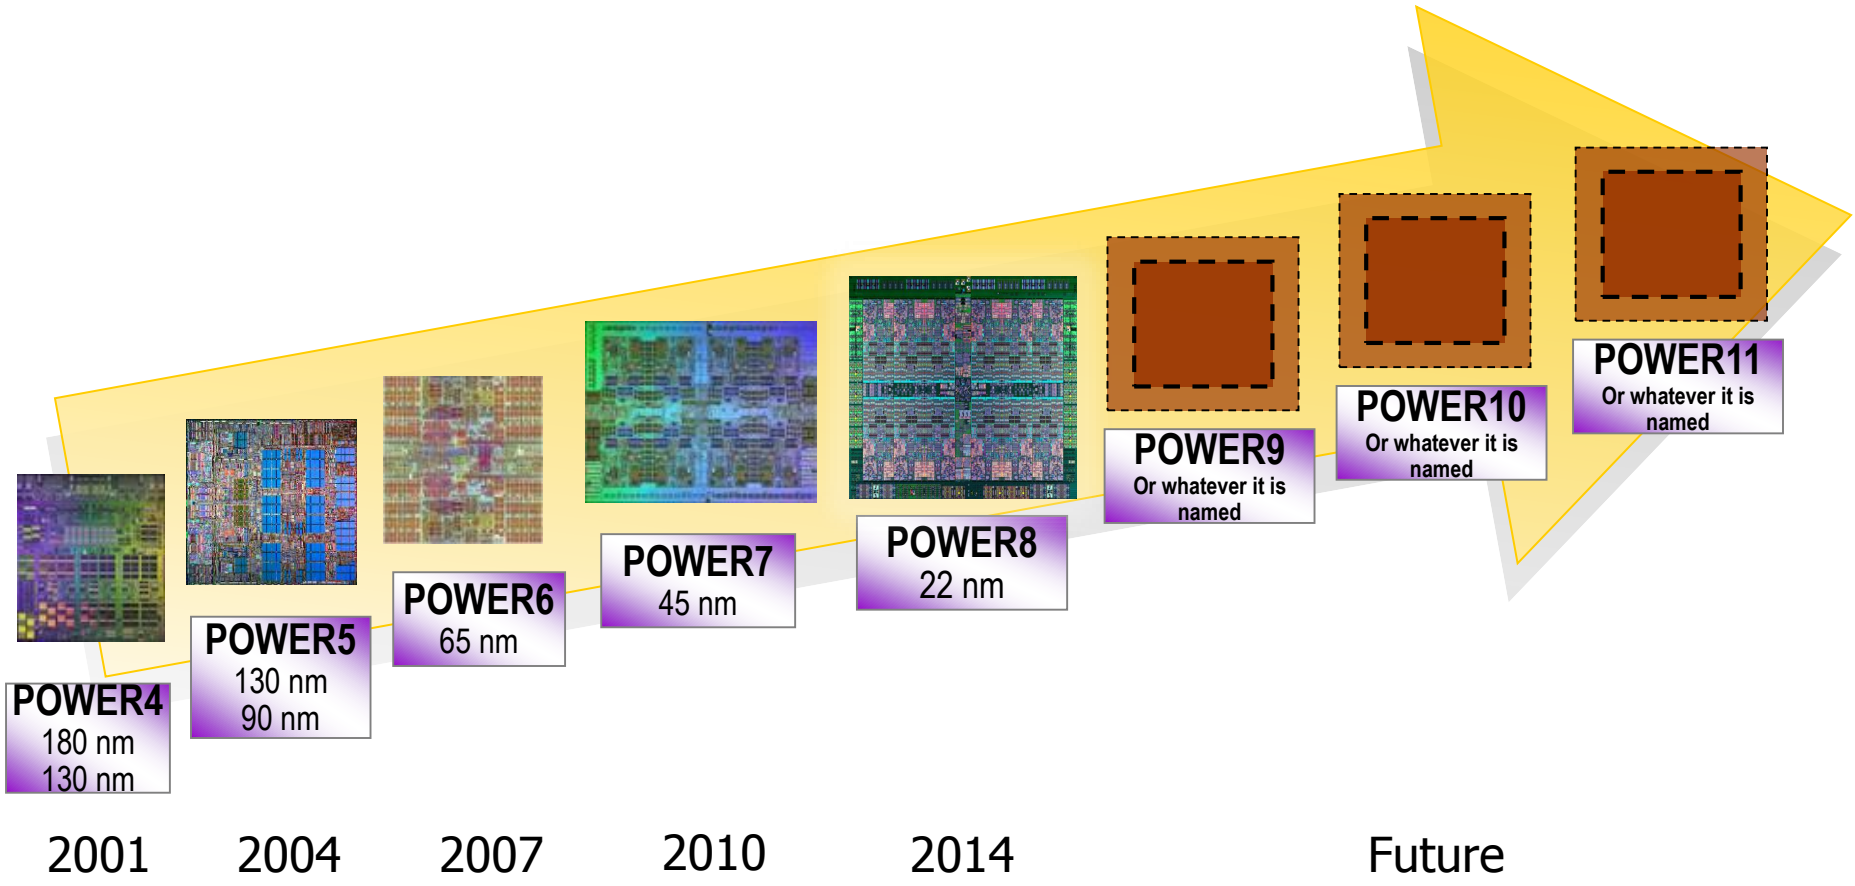

Processor Chip Comparisons

	POWER5 *2004*	**POWER6** *2007*	**POWER7** *2010*	**POWER7+** *2012*	**POWER8**
	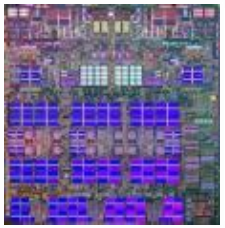		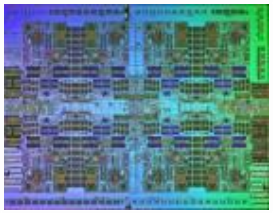	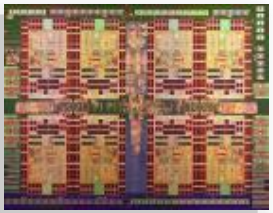	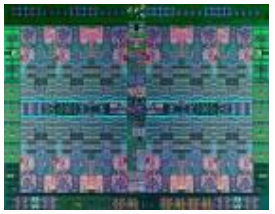
Technology	130nm SOI	65nm SOI	45nm SOI eDRAM	32nm SOI eDRAM	22nm SOI eDRAM
Compute					
Cores	2	2	8	8	12
Threads	SMT2	SMT2	SMT4	SMT4	SMT8
Caching					
On-chip	1.9MB (L2)	8MB (L2)	2 + 32MB (L2+3)	2 + 80MB (L2+3)	6 + 96MB (L2+3)
Off-chip	36MB (L3)	32MB (L3)	None	None	128MB (L4)
Bandwidth					
Sust. Mem.	15GB/s	30GB/s	100GB/s	100GB/s	230GB/s
Peak I/O	6GB/s	20GB/s	40GB/s	40GB/s	96GB/s

Processor Designs

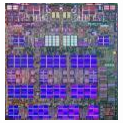

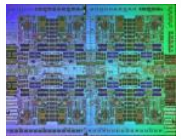
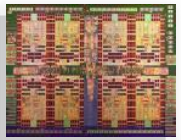
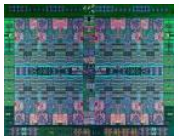

	POWER5+	POWER6	POWER7	POWER7+	POWER8
Max cores	4	2	8	8	12
Technology	90nm	65nm	45nm	32nm	22nm
Size	245 mm^2	341 mm^2	567 mm^2	567 mm^2	650 mm^2 *
Transistors	276 M	790 M	1.2 B	2.1 B	4.2 B *
Frequencies	1.9 GHz	4 - 5 GHz	3 – 4 GHz	Up to 4.4 GHz	Up to 4.1 GHz **
SMT (threads)	2	2	4	4	8
L2 Cache	1.9MB Shared	4MB / Core	256KB / core	256KB / core	512KB/core
L3 Cache	36MB	32MB	4MB / Core On chip	10MB / Core On chip	8MB / Core On chip
L4 Cache	--	--	--	--	Up to 128MB
Bandwidth **Sust memory** **Peak I/O**	15GB/s 6GB/s	30GB/s 20GB/s	100GB/s 40GB/s	100GB/s 40GB/s	230GB/s 96GB/s

*** with 12-core chip**
**** announced so far**

Processor Designs

	POWER4	POWER4+	POWER5	POWER5+	POWER6	POWER7	POWER7+	POWER8
Max cores	2	2	2	4	2	8	8	12
Technology	180nm	130nm	130nm	90nm	65nm	45nm	32nm	22nm
Size	412mm^2	mm^2	389 mm^2	245 mm^2	341 mm^2	567 mm^2	567 mm^2	650 mm^2 *
Transistors	170 M	180 M	276 M	276 M	790 M	1.2 B	2.1 B	4.2 B *
Frequencies	1.3 GHz	1.9 GHz	1.65-1.9 GHz	1.9-2.2 GHz	3 - 5 GHz	3 – 4.25 GHz	Up to 4.4 GHz	3 -- 4.35 GHz
SMT (threads)	1	1	2	2	2	4	4	8
L2 Cache	1.4 MB Shared	1.5 MB Shared	1.9MB Shared	1.9MB Shared	4MB / Core	256KB / core	256KB / core	512KB/core
L3 Cache	32MB	32MB	36MB	36MB	32MB	4MB / Core On chip	10MB / Core On chip	8MB / Core On chip
L4 Cache	--	--	--	--	--	--	--	Up to 128MB
Bandwidth **Sust memory** **Peak I/O**	6GB/s 2GB/s	6GB/s 2GB/s	15GB/s 6GB/s	15GB/s 6GB/s	30GB/s 20GB/s	100GB/s 40GB/s	100GB/s 40GB/s	230GB/s 96GB/s

*** with 12-core chip**

Innovation Drives Performance

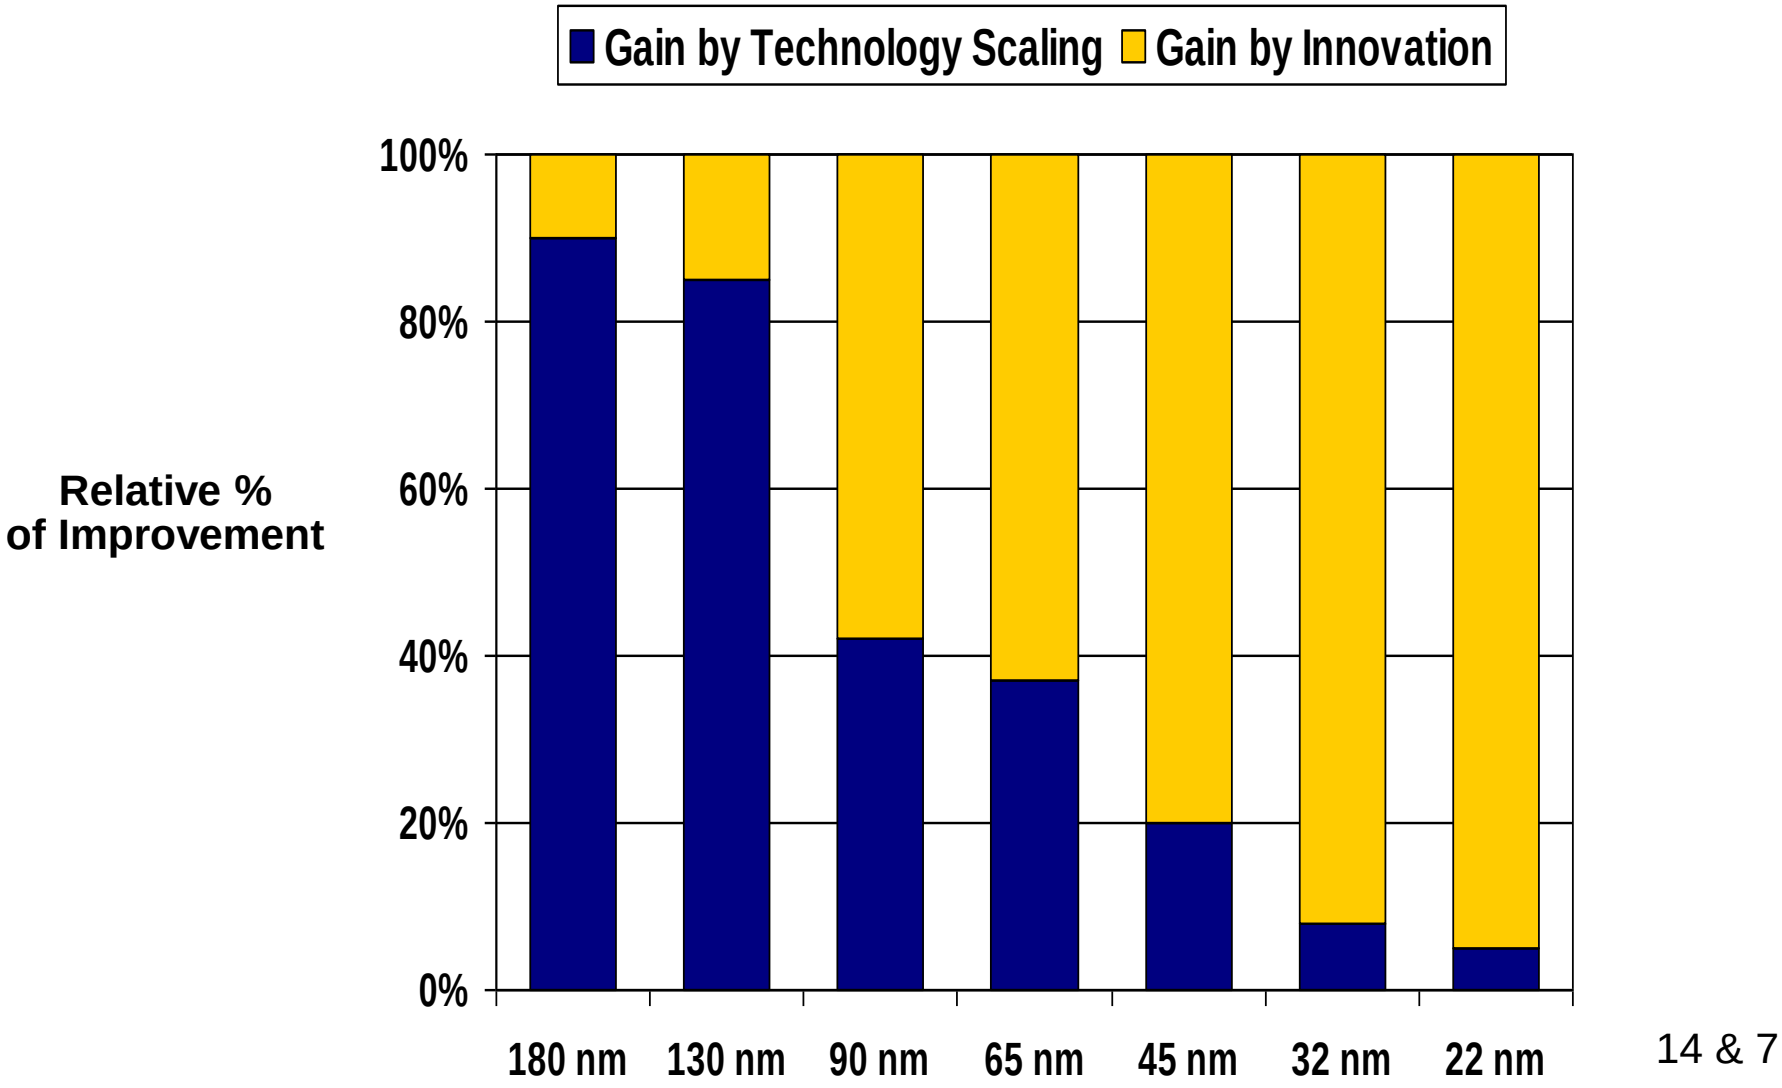

14 & 7

POWER8 Chip Packaging

Technology
- 22nm SOI, eDRAM, 650mm2
- 15-layers = great bandwidth

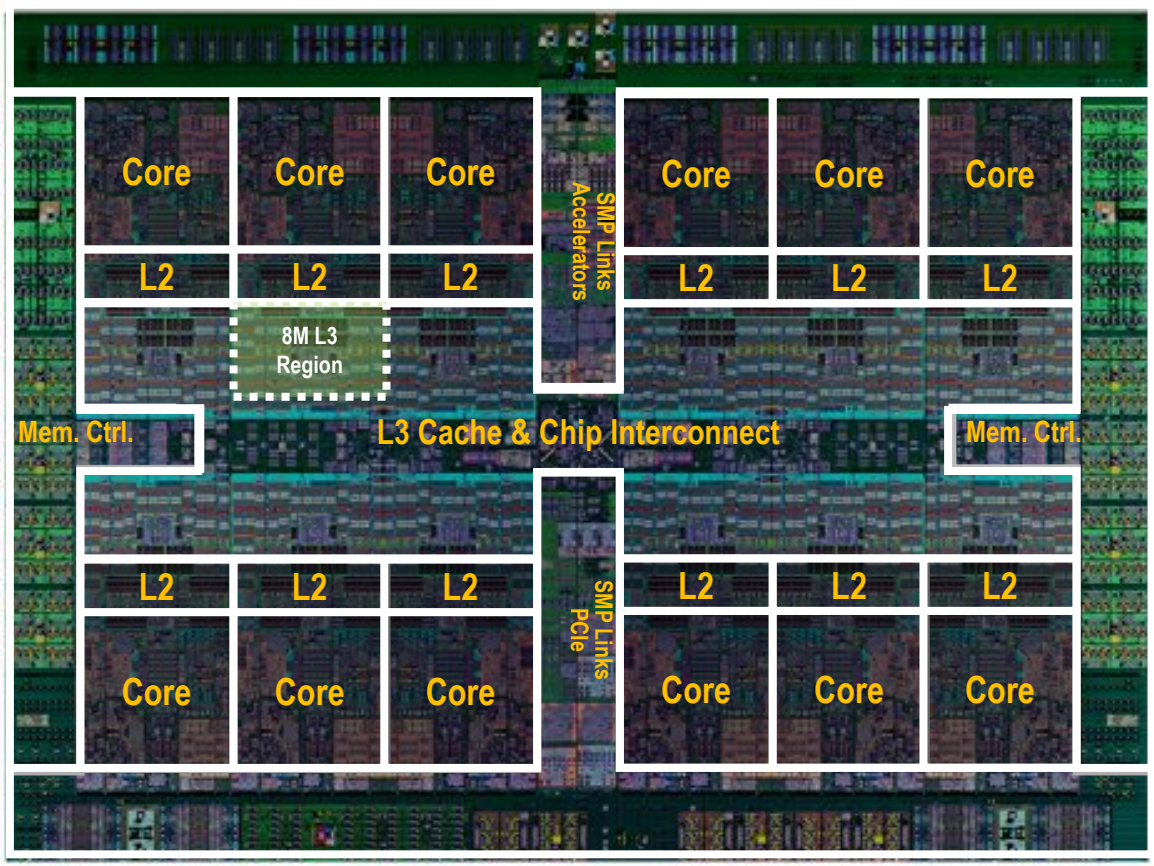

World Class 22nm Semiconductor Technology

Silicon On Insulator
-Faster Transistor, Less Noise

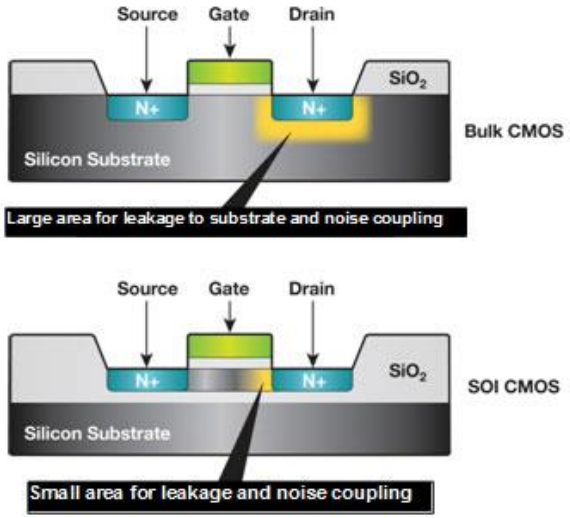

On-chip eDRAM
- 6x latency improvement
- No off-chip signaling rqmt
- 8x bandwidth improvement
- 3x less area than SRAM
- 5x less energy than SRAM

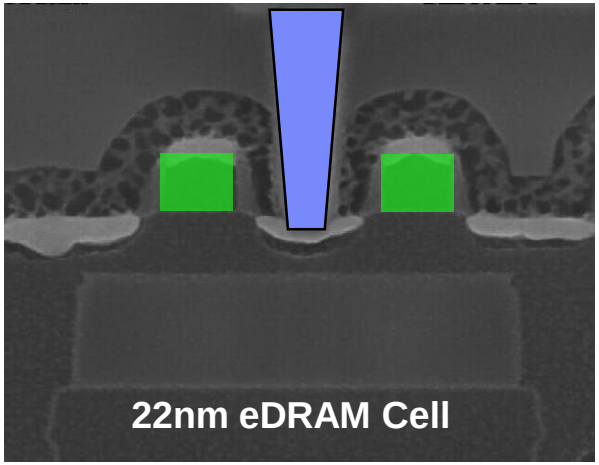

22nm 15-layer copper wire

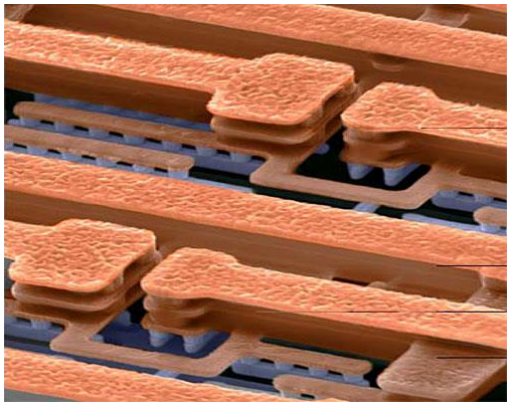

Dense interconnect
- Faster connections
- Low latency distance paths
- High density complex circuits
- 2X wire per transistor

IBM 22nm Technology

Low level wires used for dense local circuit interconnect.

Top level wires used for for power distribution, clocks, and off-chip signaling.

15 Layer IBM metal stack

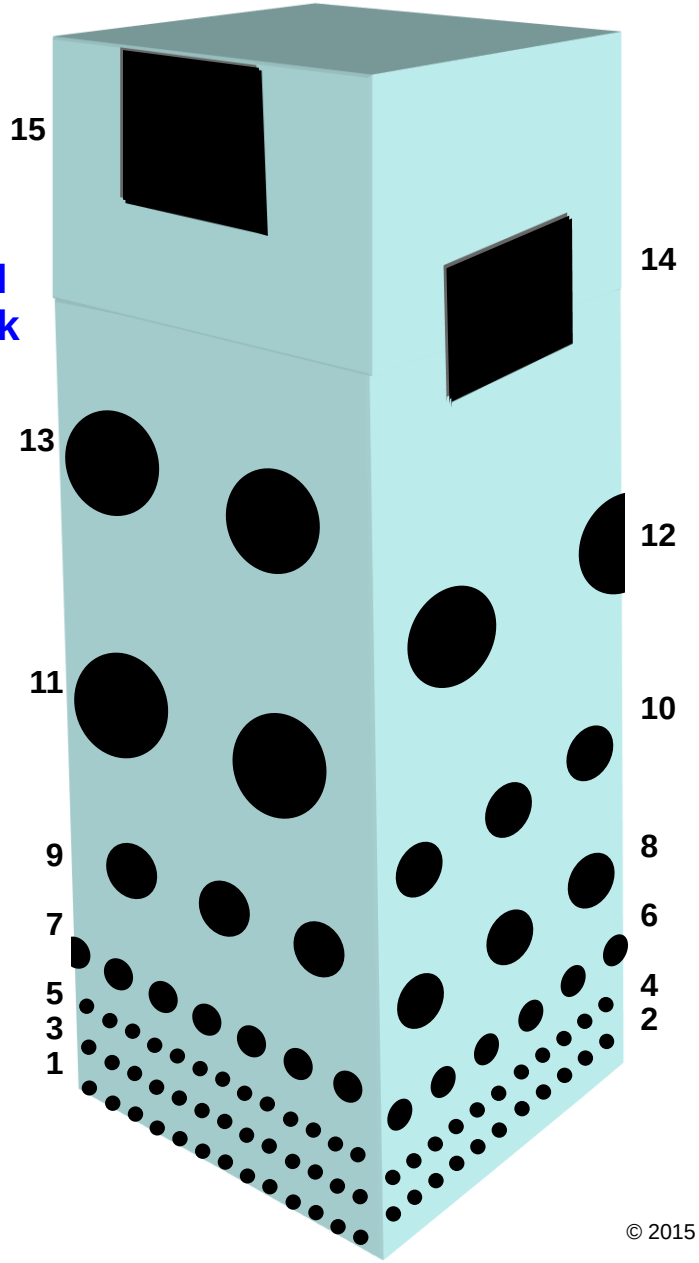

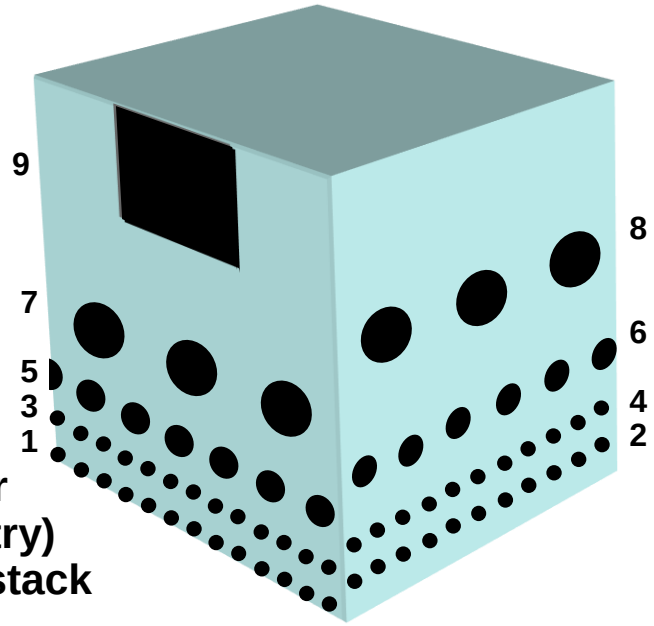

9 Layer (industry) metal stack

POWER8 Chip Packaging

Cores
- 12 cores (SMT8)
- 8 dispatch, 10 issue, 16 exec pipe
- 2X internal data flows/queues
- Enhanced prefetching

Caches
- 64K Data cache (L1)
- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)

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

Bus Interfaces
- Integrated PCIe Gen3
- SMP Interconnect
- CAPI

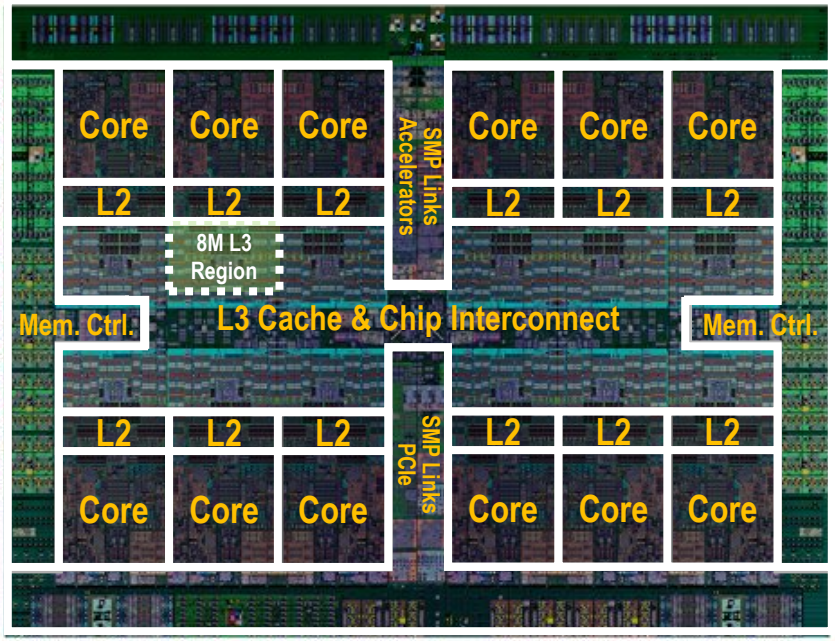

Energy Management
- On-chip Power Management Micro-controller
- Integrated Per-core VRM
- Critical Path Monitors

Memory
- Dual memory Controllers
- 230 GB/sec Sustained bandwidth

POWER8 Memory Buffer Chip

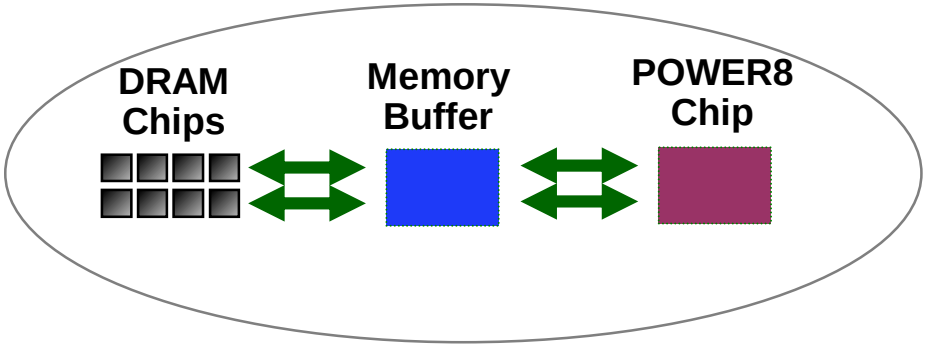

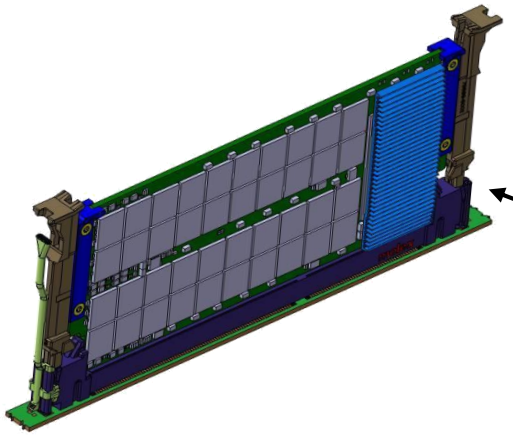

“L4 cache”

Intelligence Moved into Memory

- Previously on POWER7+ chip onto buffer

Processor Interface

- High speed interface

Performance Value

Memory Bandwidth per Socket

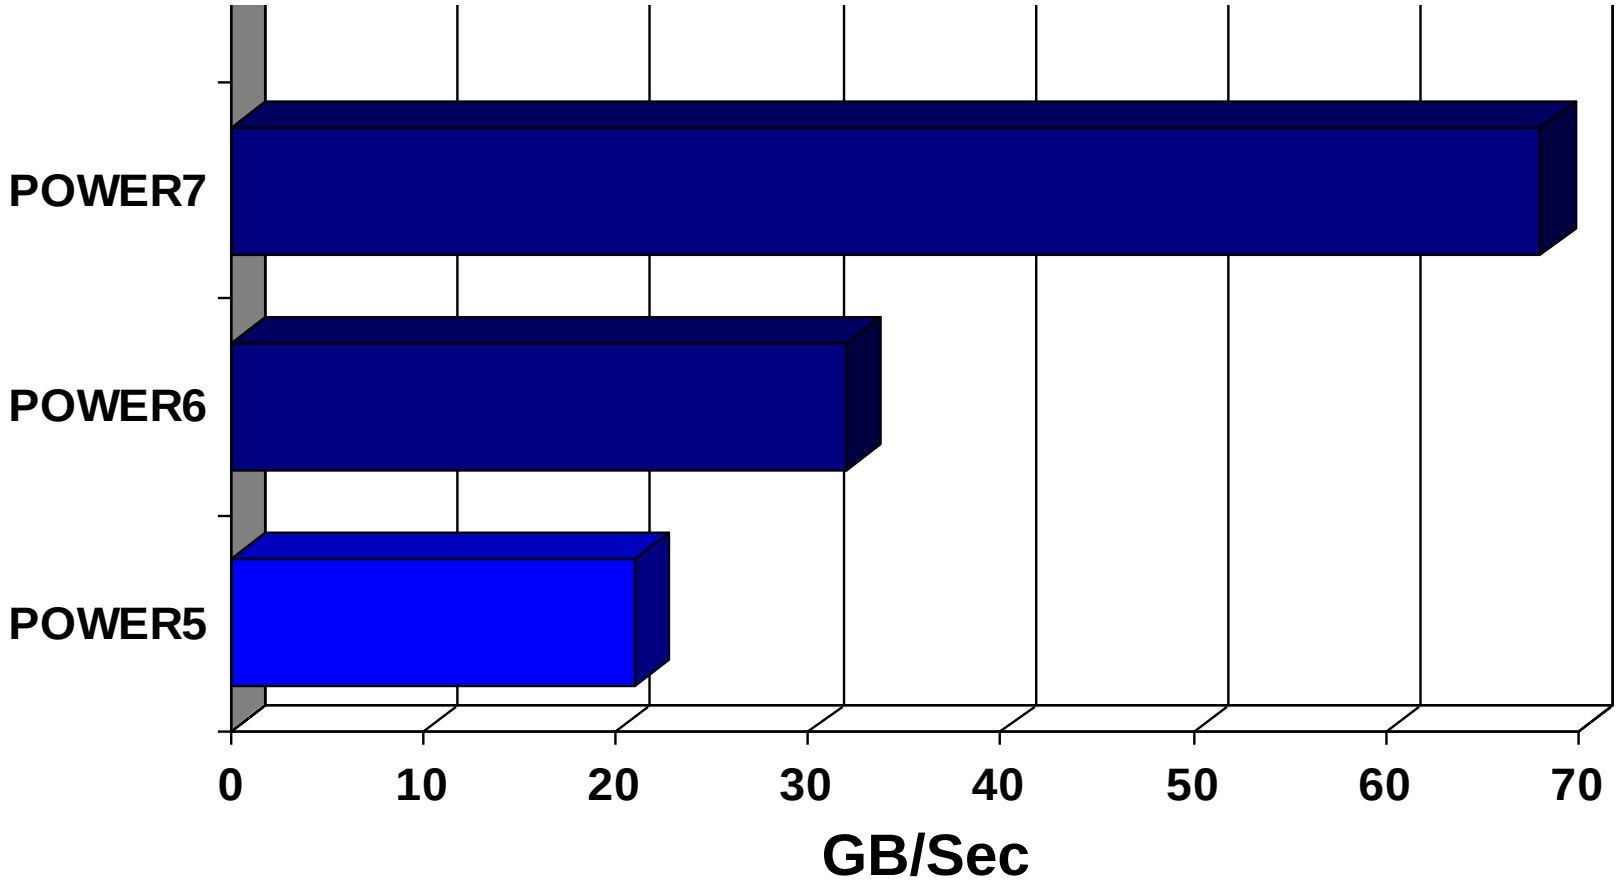

Memory Bandwidth per Socket

Reset the scale

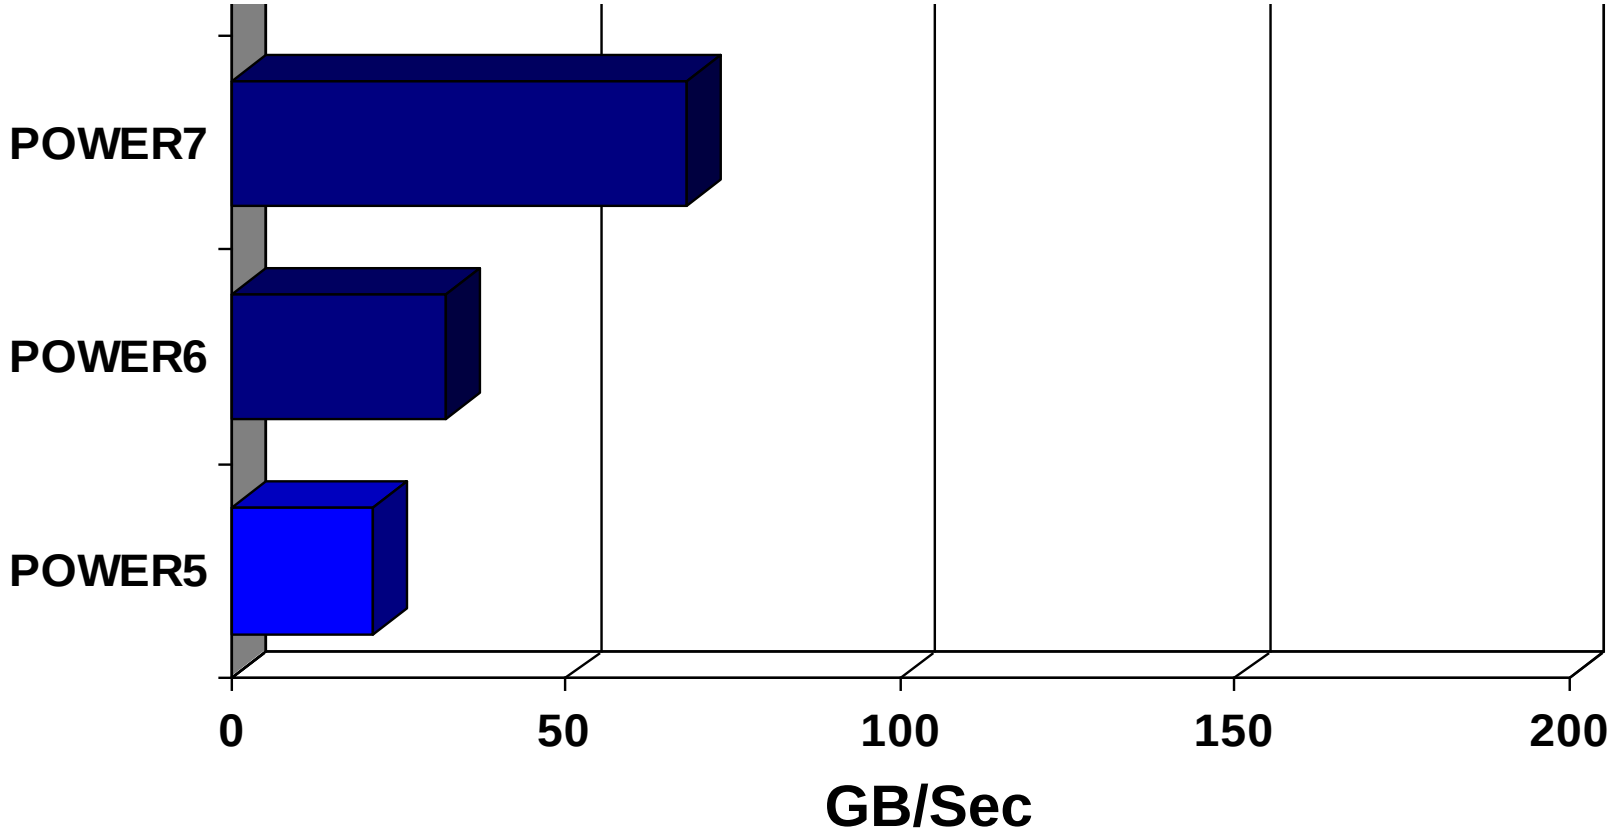

POWER8 Memory Bandwidth per Socket

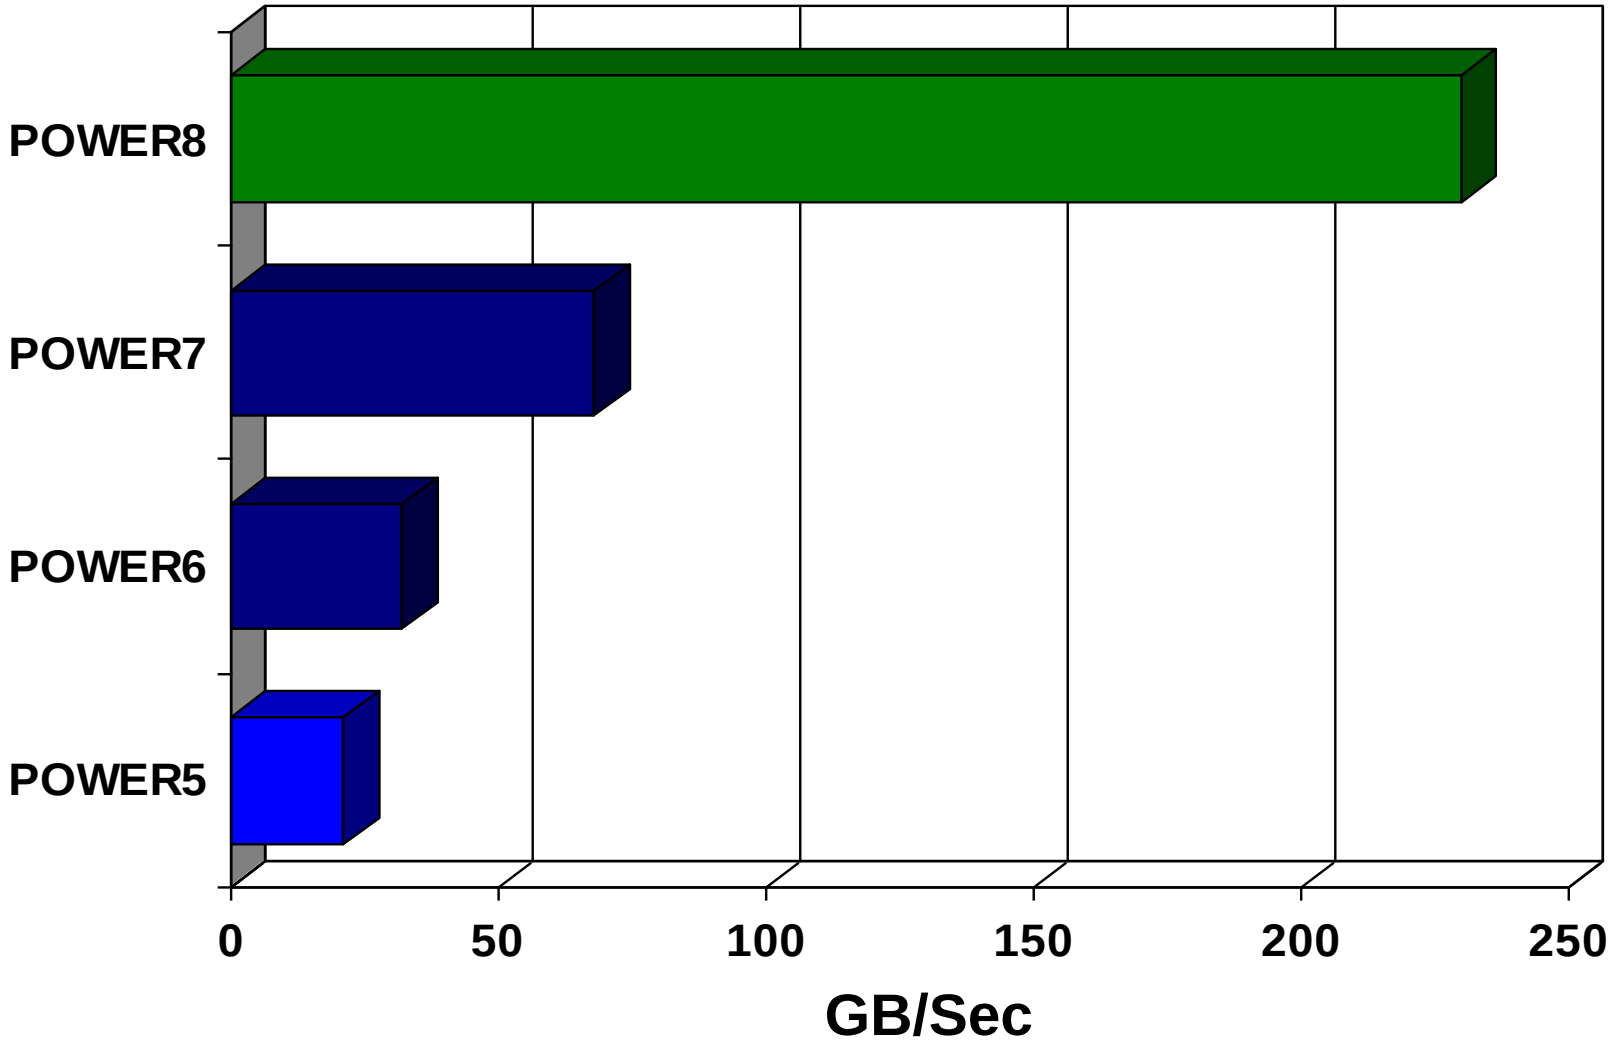

POWER8 Integrated PCI Gen 3

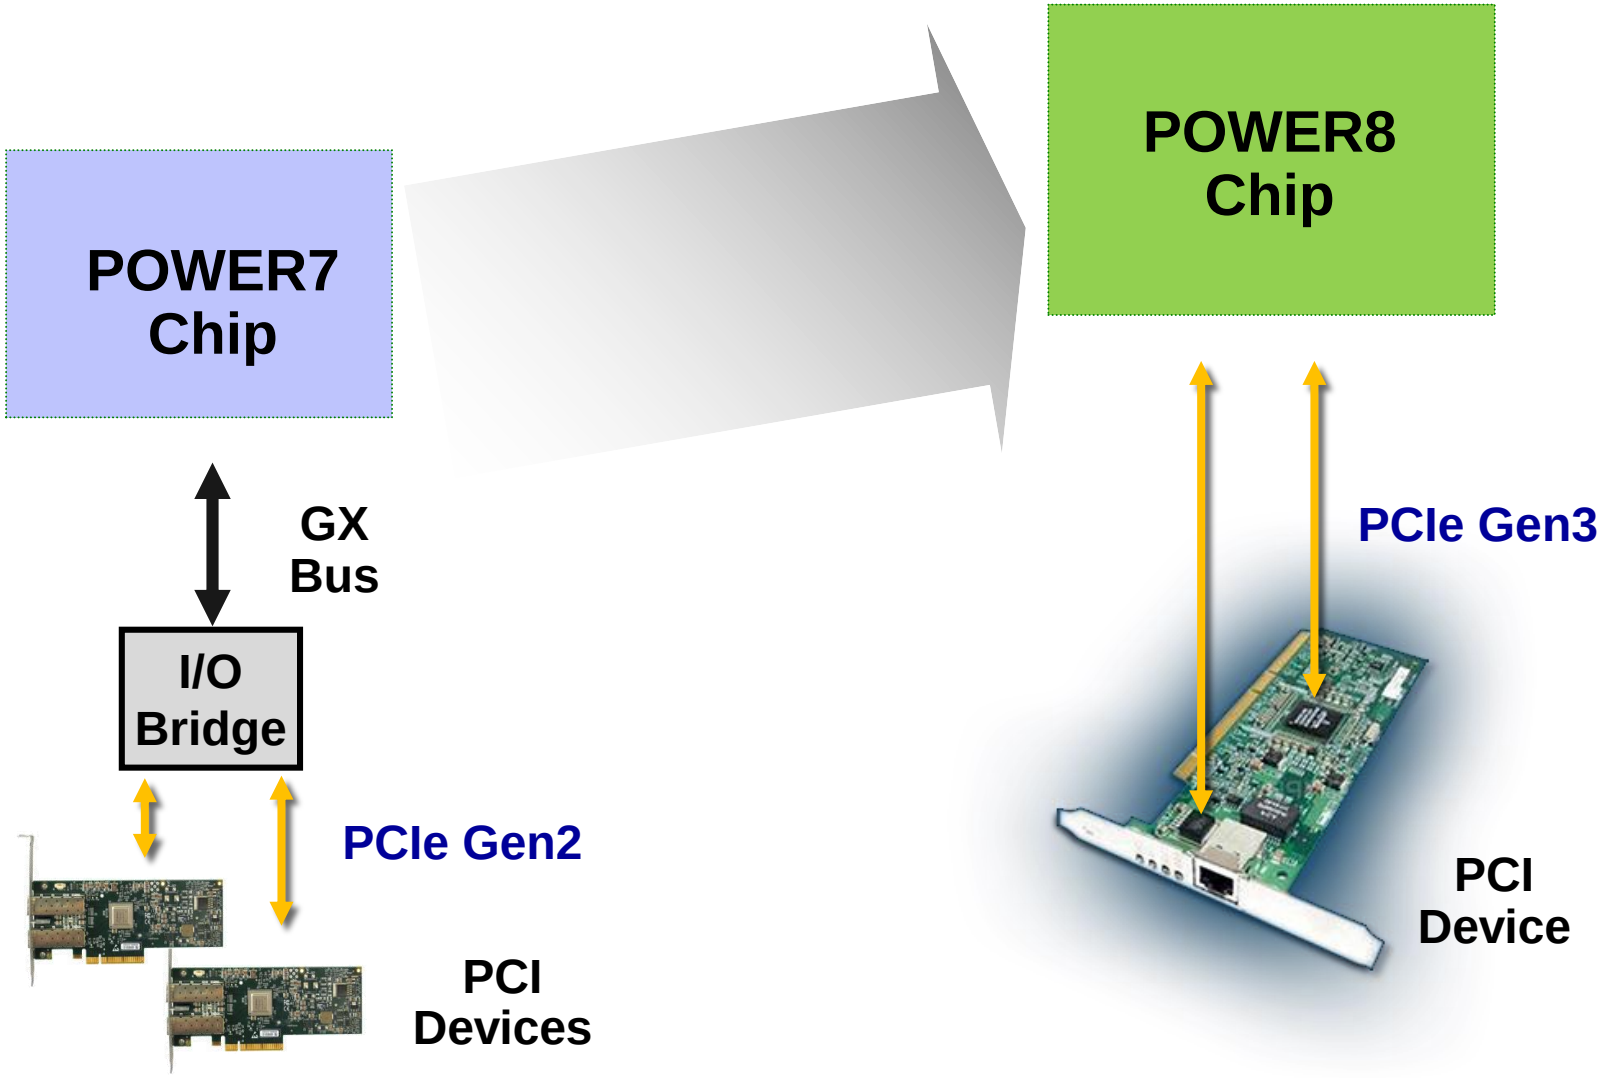

Power 770/780 Node I/O Bandwidth

(System node or processor enclosure or CEC drawer)

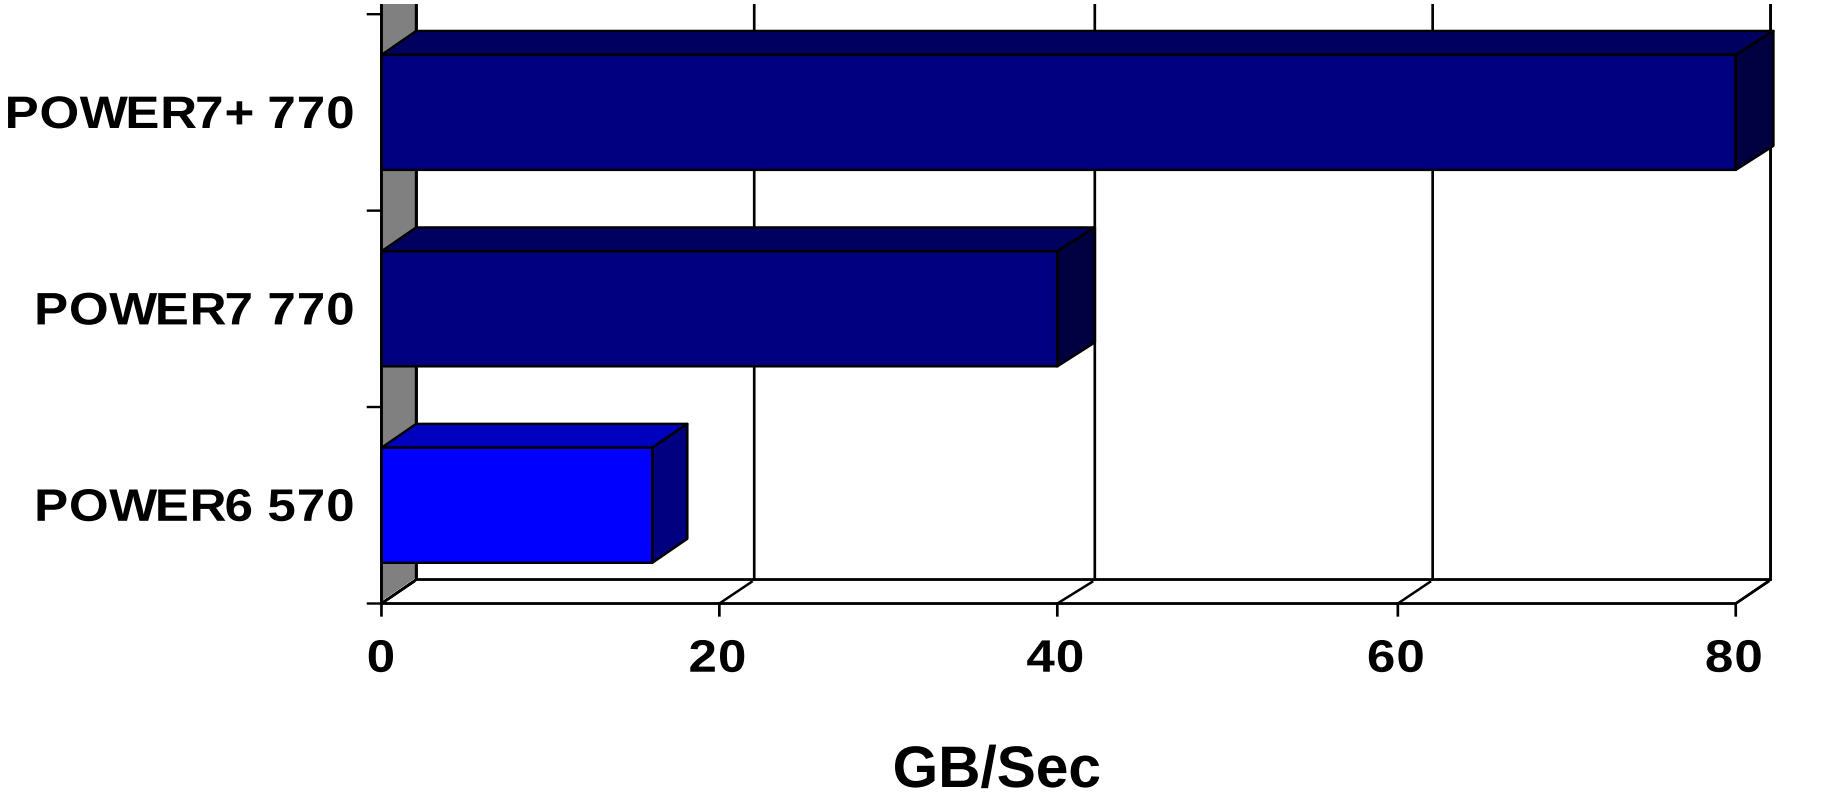

Power 770/780 Node I/O Bandwidth

(System node or processor enclosure or CEC drawer)

Reset the scale

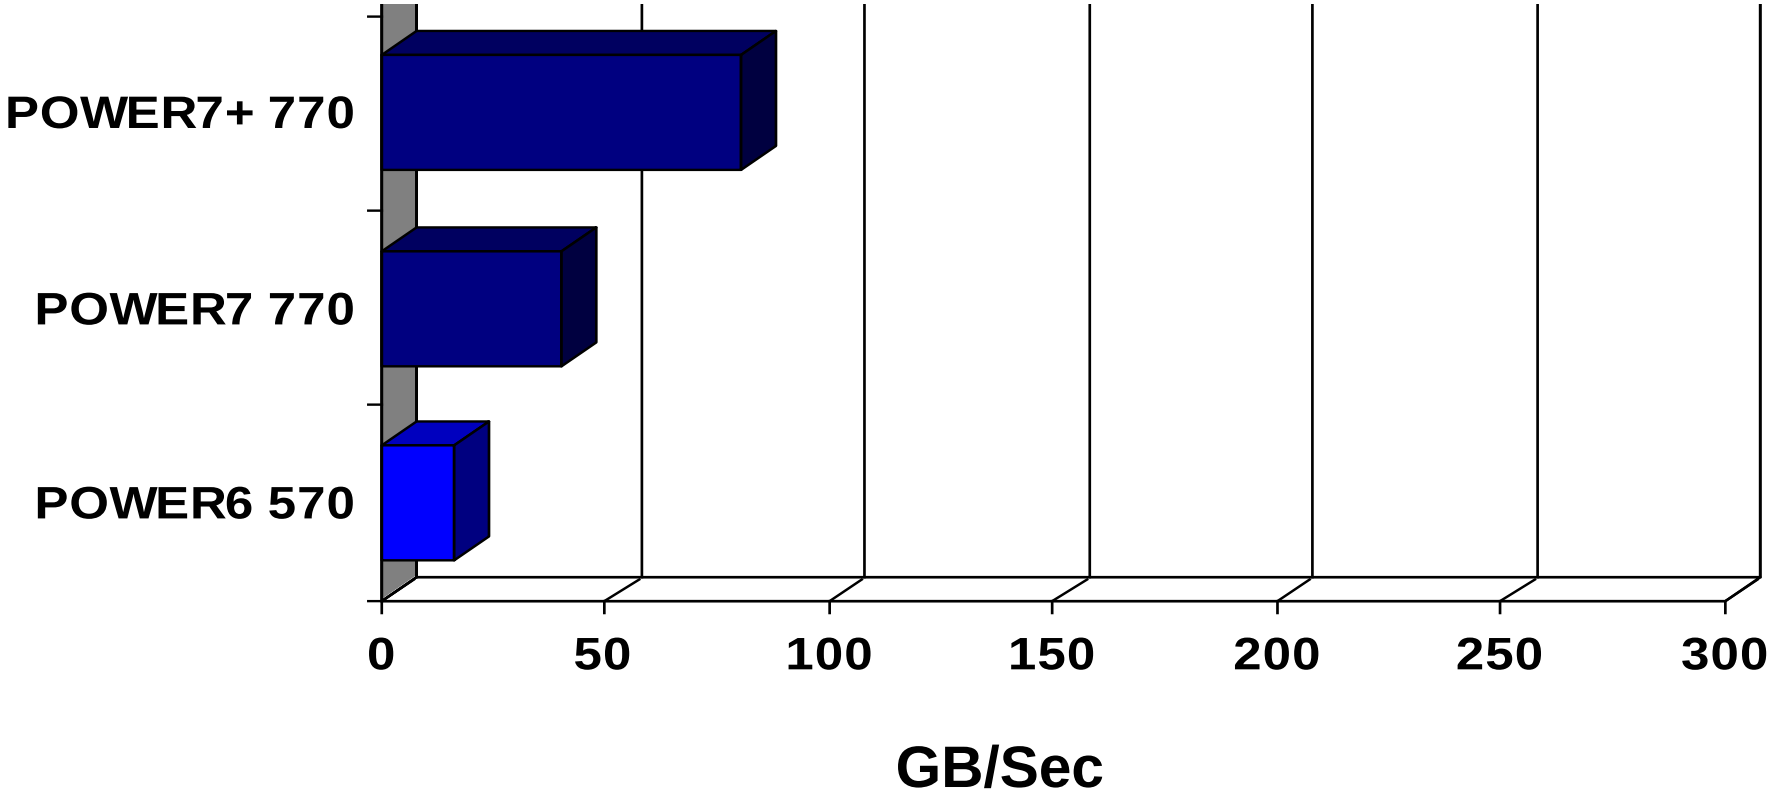

E870/E880 Node I/O Bandwidth

(System node or processor enclosure or CEC drawer)

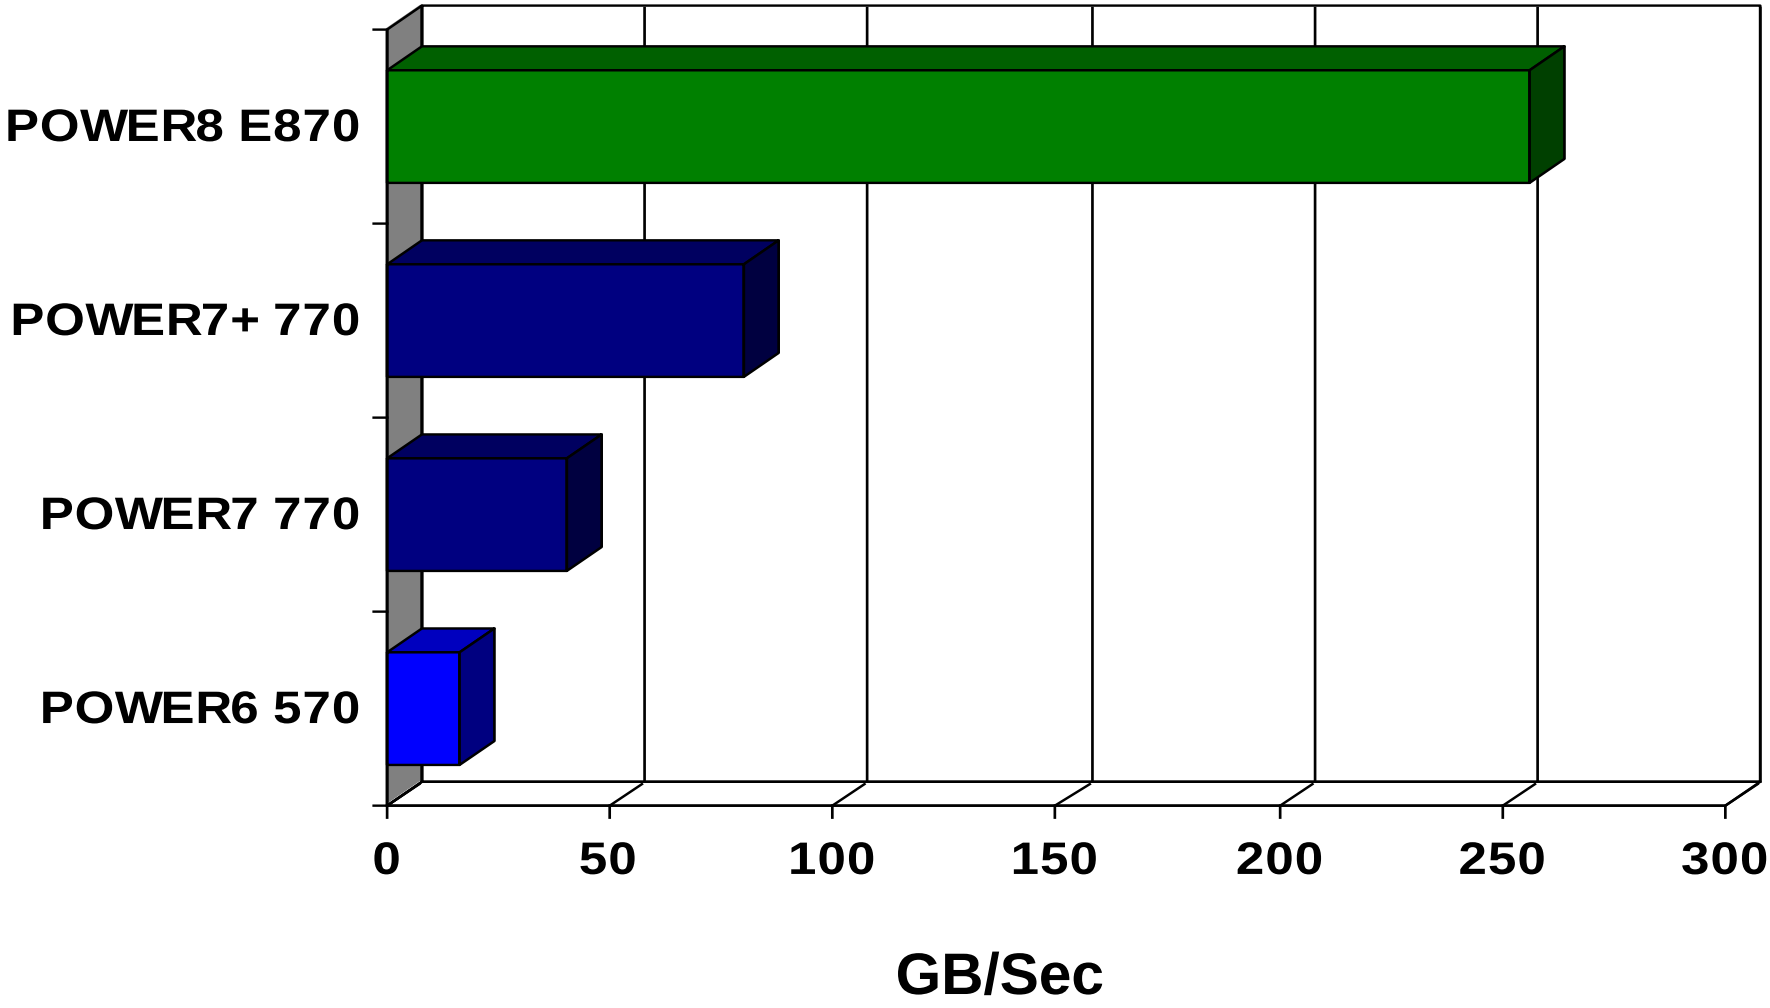

IO Bandwidth Comparing 2-Socket Servers

Reset the scale

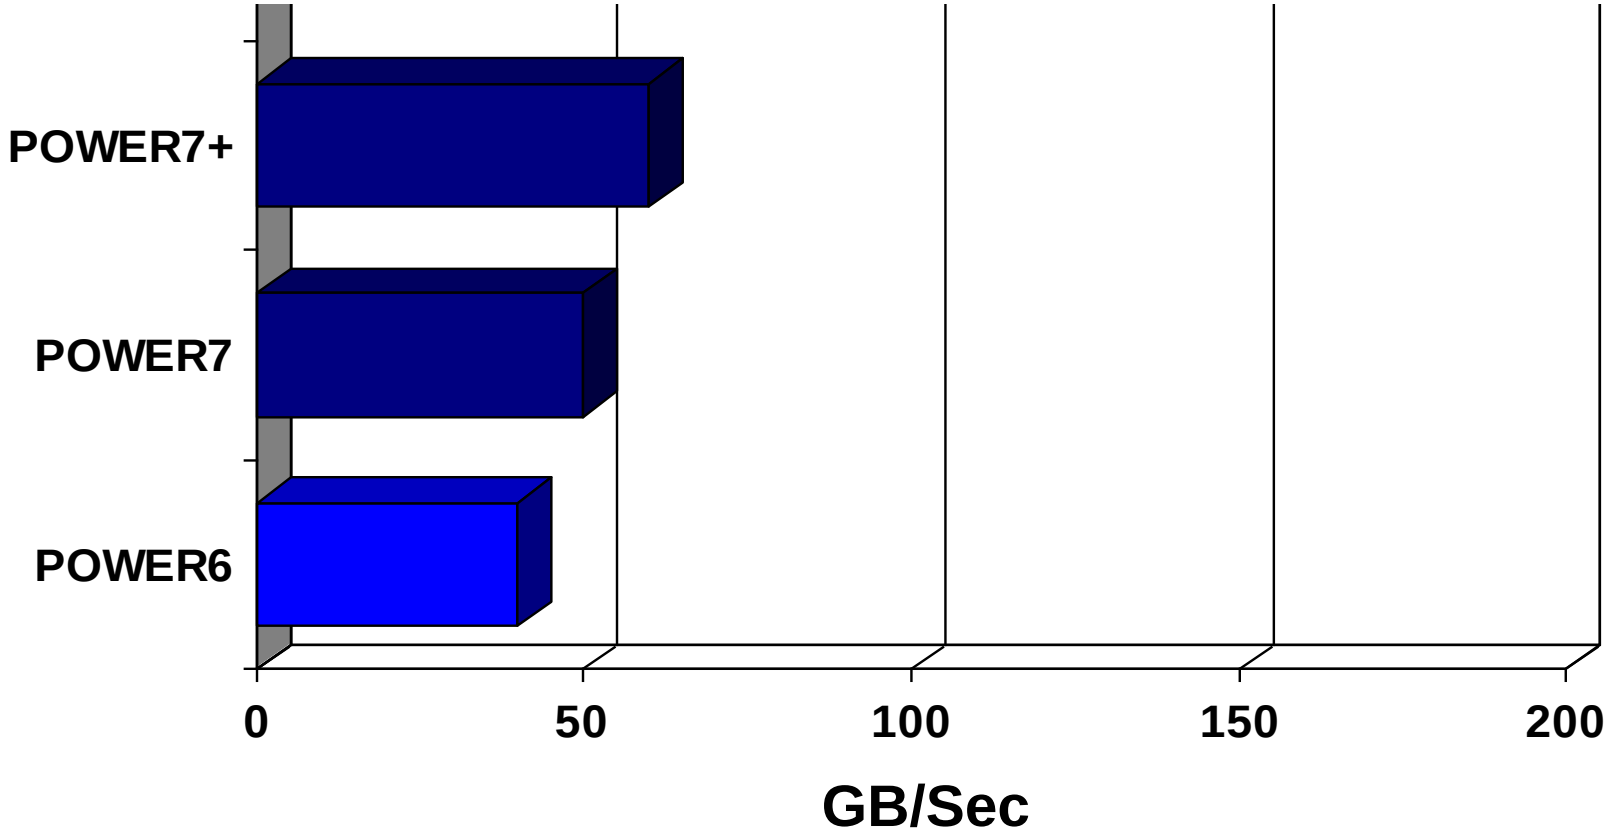

POWER8 IO Bandwidth Comparing 2-Socket Servers

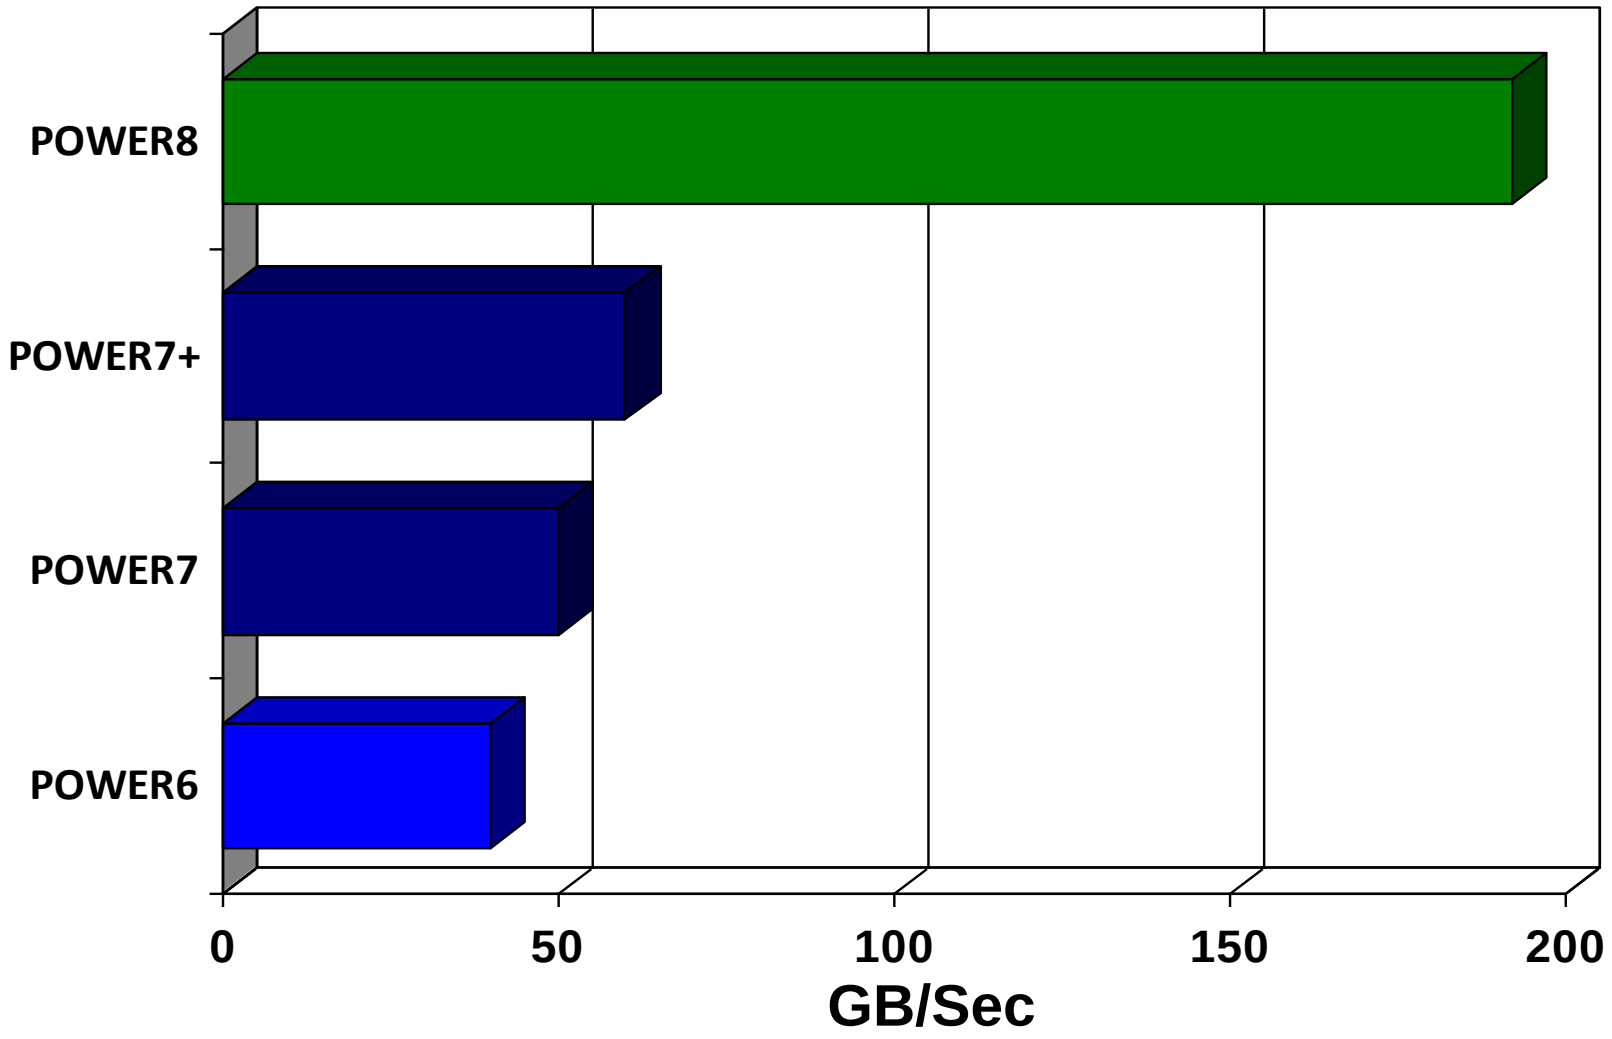

POWER8 CAPI (Coherent Accelerator Processor Interface)

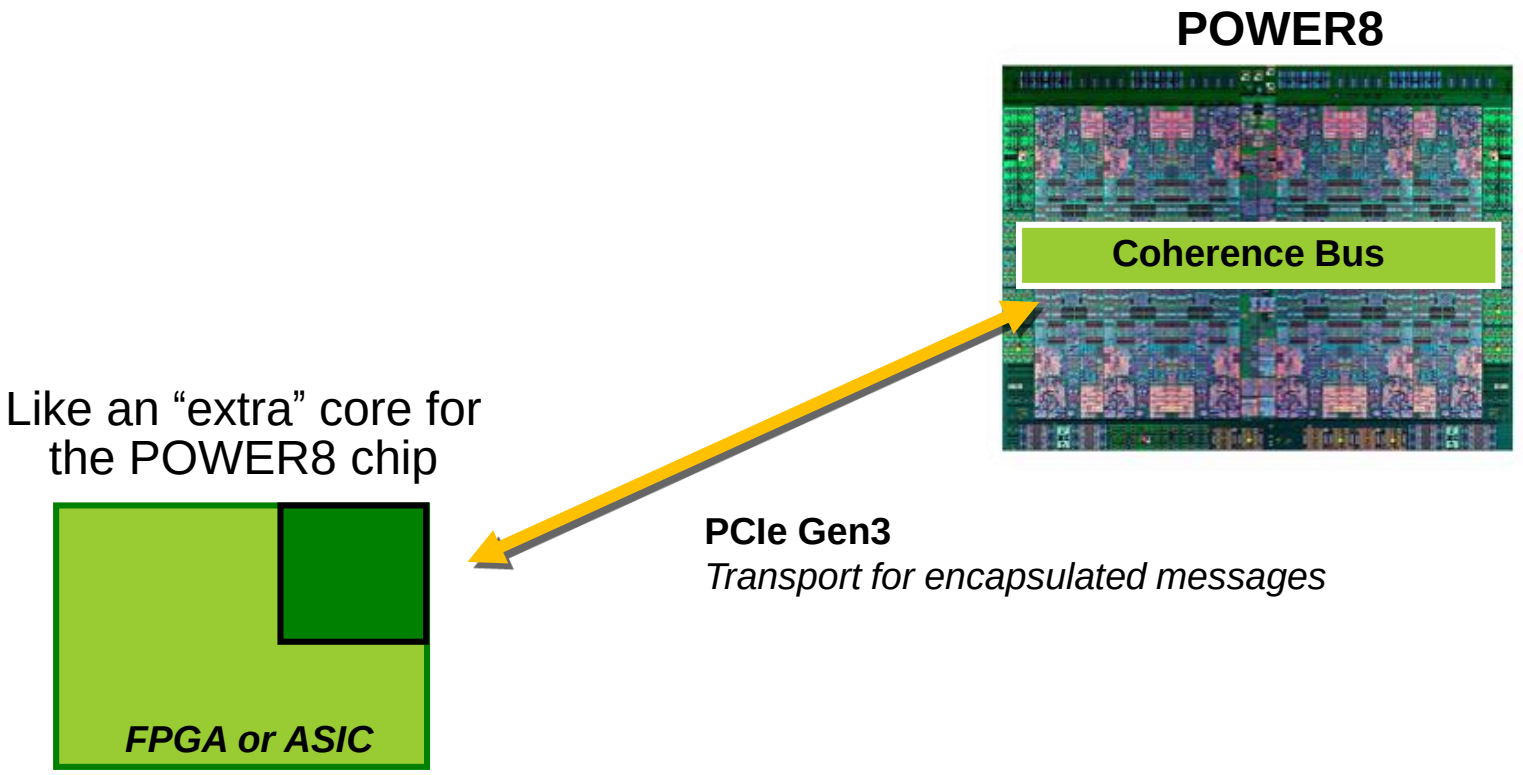

Customizable Hardware / Application Accelerator

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

First CAPI Solution Example

redislabs

High speed, fast response, social application – example Twitter

Enabled by in-memory NoSQL, distributed hash tables

Was initially implemented on x86 servers, but limited DRAM memory meant LOTS of servers resulting in a costly, complex infrastructure

24 : 1

One 2-socket, 2U server PLUS one FlashSystem Drawer replaced 24 x86 servers

Much lower cost of acquisition
Much smaller foot print, less energy
Much lower operational cost

Example: CAPI Attached Flash Optimization

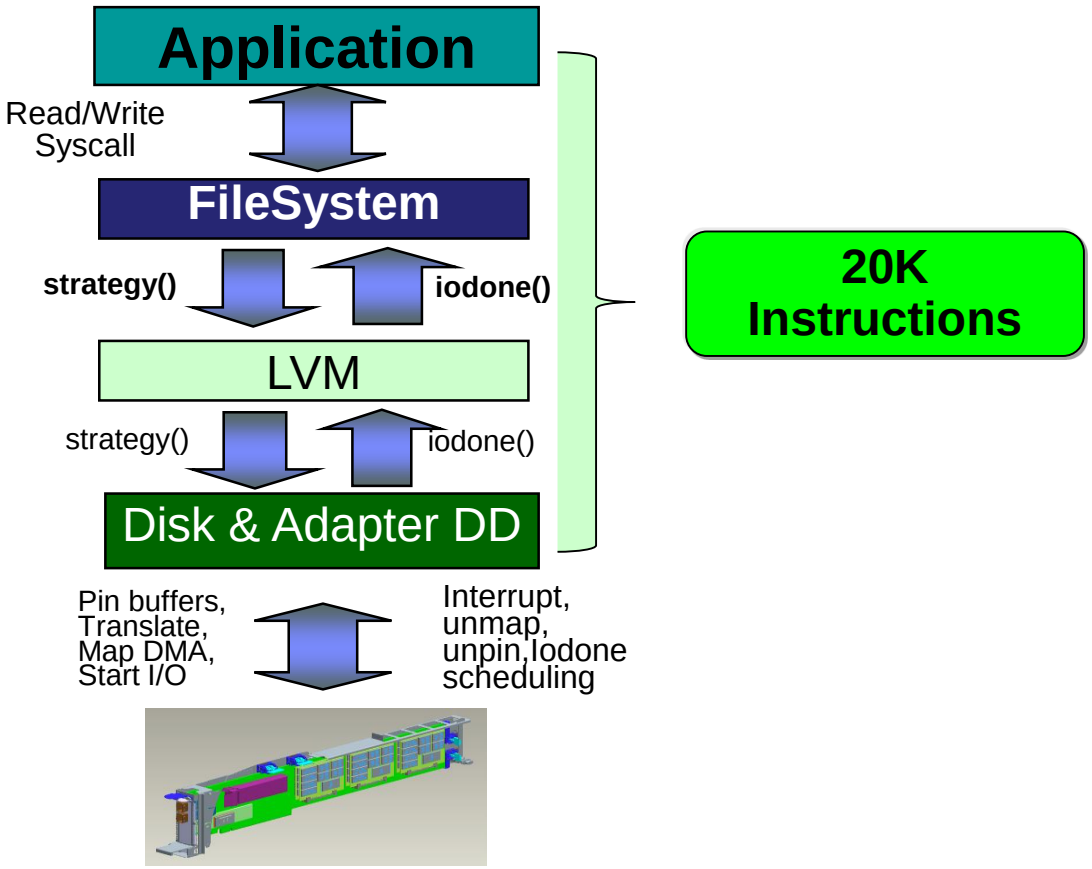

Example: CAPI Attached Flash Optimization

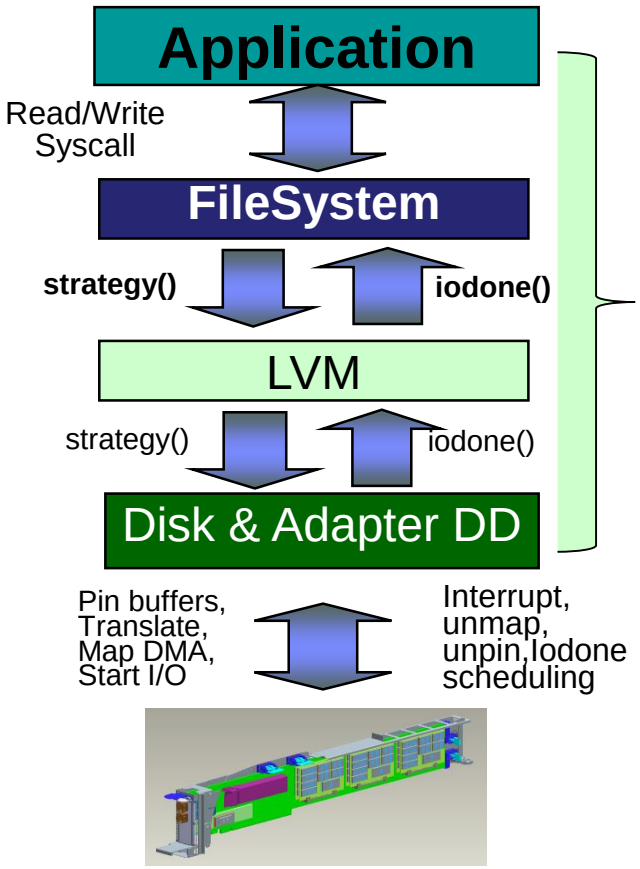

20K Instructions

< 500 Instructions

- **Attach flash memory to POWER8 via CAPI coherent Attach**

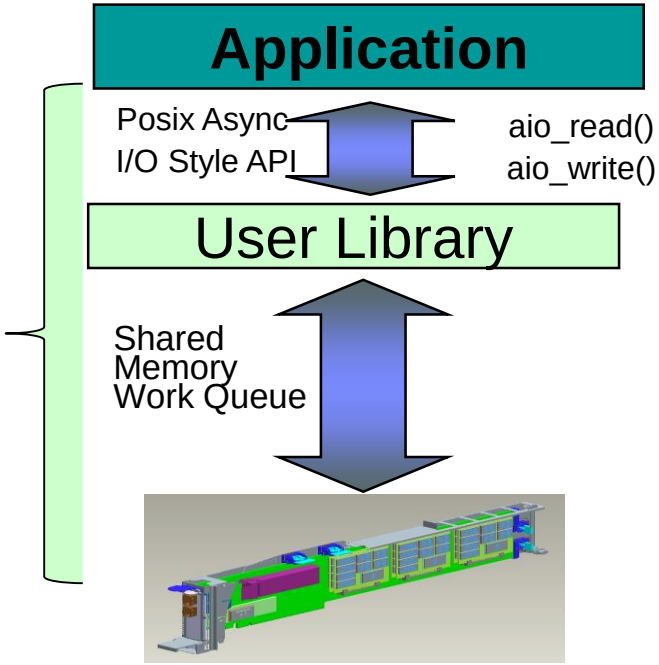

- **Issues Read/Write Commands from applications to eliminate 97% of instruction path length** **CAPI Flash controller Operates in User Space**
- **Saves 10 Cores per 1M IOPs**

POWER8 Leapfrogs

Memory Bandwidth

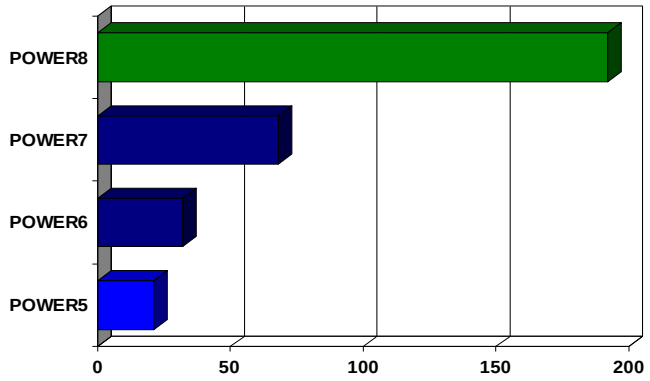

I/O Bandwidth

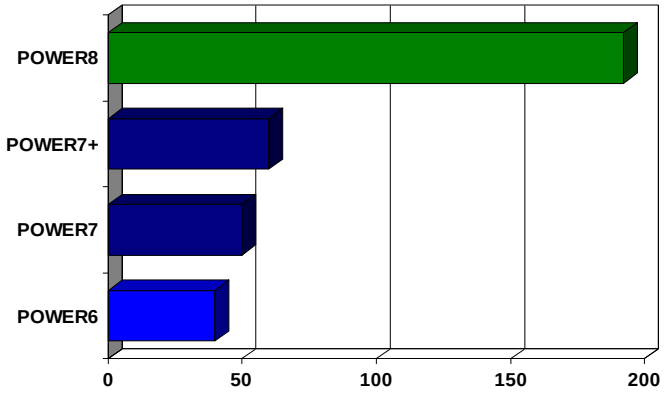

PCIe Gen3

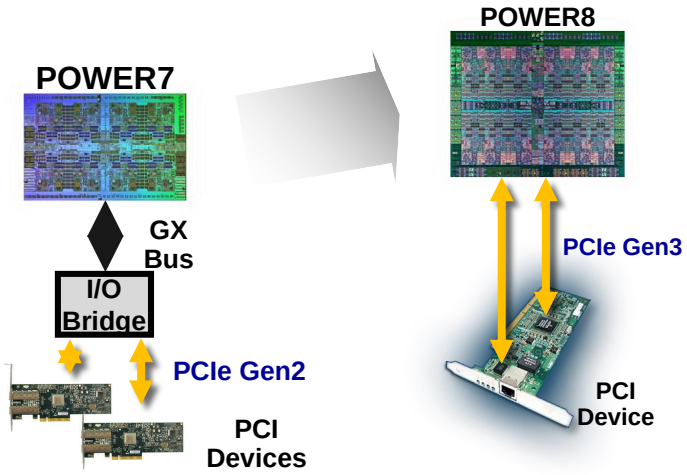

PLUS ....

- CAPI
- Accelerators
- Transactional Memory
- Scaleability
- Smart use of energy
- … and more

Scale-out CPW Comparisons

- 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%

- S814 (1 socket)
  - 4-core 3.0 GHz 39,500
  - 6-core 3.0 GHz 59,500
  - 8-core 3.7 GHz 85,500

+50% ~ GHz

- 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%

CPW

POWER7+ 770

- 16-core 3.8 GHz 110,000
- 32-core 3.8 GHz 191,500
- 48-core 3.8 GHz 290,500
- 64-core 3.8 GHz 379,300

- 12-core 4.2 GHz 90,000
- 24-core 4.2 GHz 154,800
- 36-core 4.2 GHz 242.600
- 48-core 4.2 GHz 306,600

POWER7+ 780

- 16-core 4.4 GHz 123,500
- 32-core 4.4 GHz 214,000
- 48-core 4.4 GHz 326,100
- 64-core 4.4 GHz 424,400

- 32-core 3.7 GHz 209,500
- 64-core 3.7 GHz 414,900
- 96-core 3.7 GHz 622,300
- 128-core 3.7 GHz 829,800

- 9

+87%

+50% ~ GHz, fewer core
+80% per core ~ GHz

+80%
~ GHz

+150%

E870

- 32-core 4.02 GHz 359,000
- 64-core 4.02 GHz 711,000

- 40-core 4.19 GHz 460,000
- 80-core 4.19 GHz 911,000

E880

- 32-core 4.35 GHz 381,000
- 64-core 4.35 GHz 755,000
- 96-core 4.35 GHz 1,144,000
- 128-core 4.35 GHz 1,523,000

- 48-core 4.02 GHz 518,000
- 96-core 4.02 GHz 1,034,000
- 144-core 4.02 GHz 1,551,000
- 192-core 4.02 GHz 2,069,000

Measured using SMT8
SMT4 would be somewhat lower

CPW

POWER7 795

- 48-core 3.7 GHz 288,500
- 96-core 3.7 GHz
- 144-core 3.7 GHz
- 192-core 3.7 GHz

- 64-core 4.0 GHz 399,200
- 128-core 4.0 GHz
- 192-core 4.0 GHz
- 256-core 4.0 GHz

- 32-core 4.25 GHz 218,400
- 64-core 4.25 GHz
- 96-core 4.25 GHz
- 128-core 4.25 GHz

+74%
~ GHz

E880

- 32-core 4.35 GHz 381,000
- 64-core 4.35 GHz 755,000
- 96-core 4.35 GHz 1,144,000
- 128-core 4.35 GHz 1,523,000

- 48-core 4.02 GHz 518,000
- 96-core 4.02 GHz 1,034,000
- 144-core 4.02 GHz 1,551,000
- 192-core 4.02 GHz 2,069,000

Measured using SMT8
SMT4 would be somewhat lower

“Single-thread-oriented” Workloads and POWER8 Technology

Good news

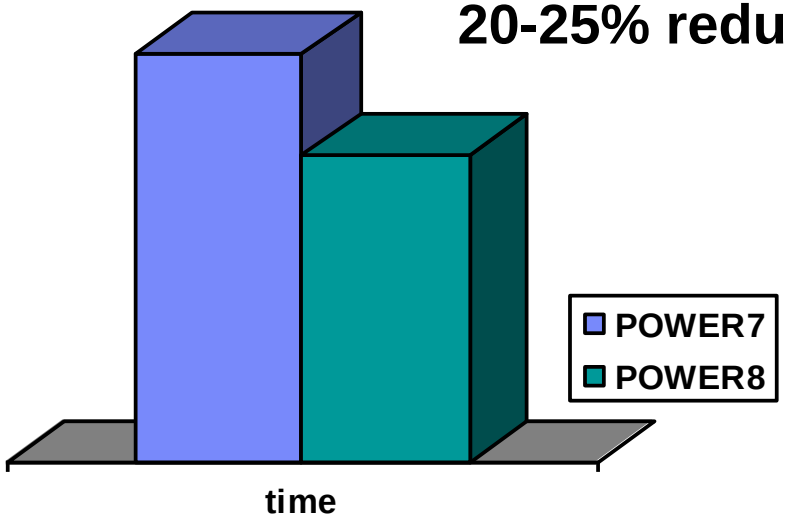

Because per core performance improving ..

Compared to POWER7/POWER7+, POWER8 chips with similar GHz run 20-25% faster from a wall clock perspective …… ASSUMING NOT I/O BOUND

A New Generation of IBM Power Systems

Designed for Big Data

Superior Cloud Economics

Open Innovation Platform

OpenPOWER Consortium

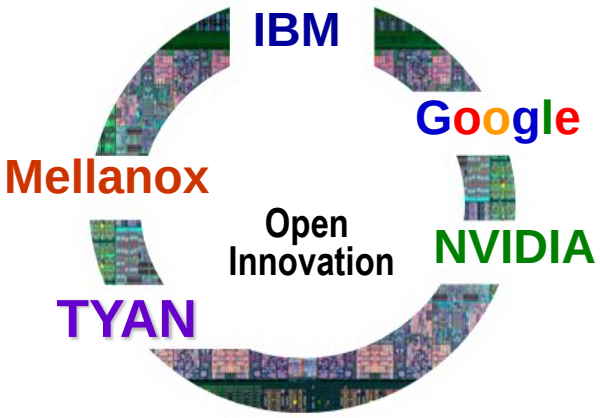

- Collaborative innovation for highly advanced systems
- Produce open hardware, software, firmware and tools
- Expand industry skills and investment for Power ecosystem
- Provide alternative architectures

Google® Mother Board using POWER8 Technology

Development board previewed by Google at April 2014 POWER8 announcement

Building collaboration and innovation at all levels

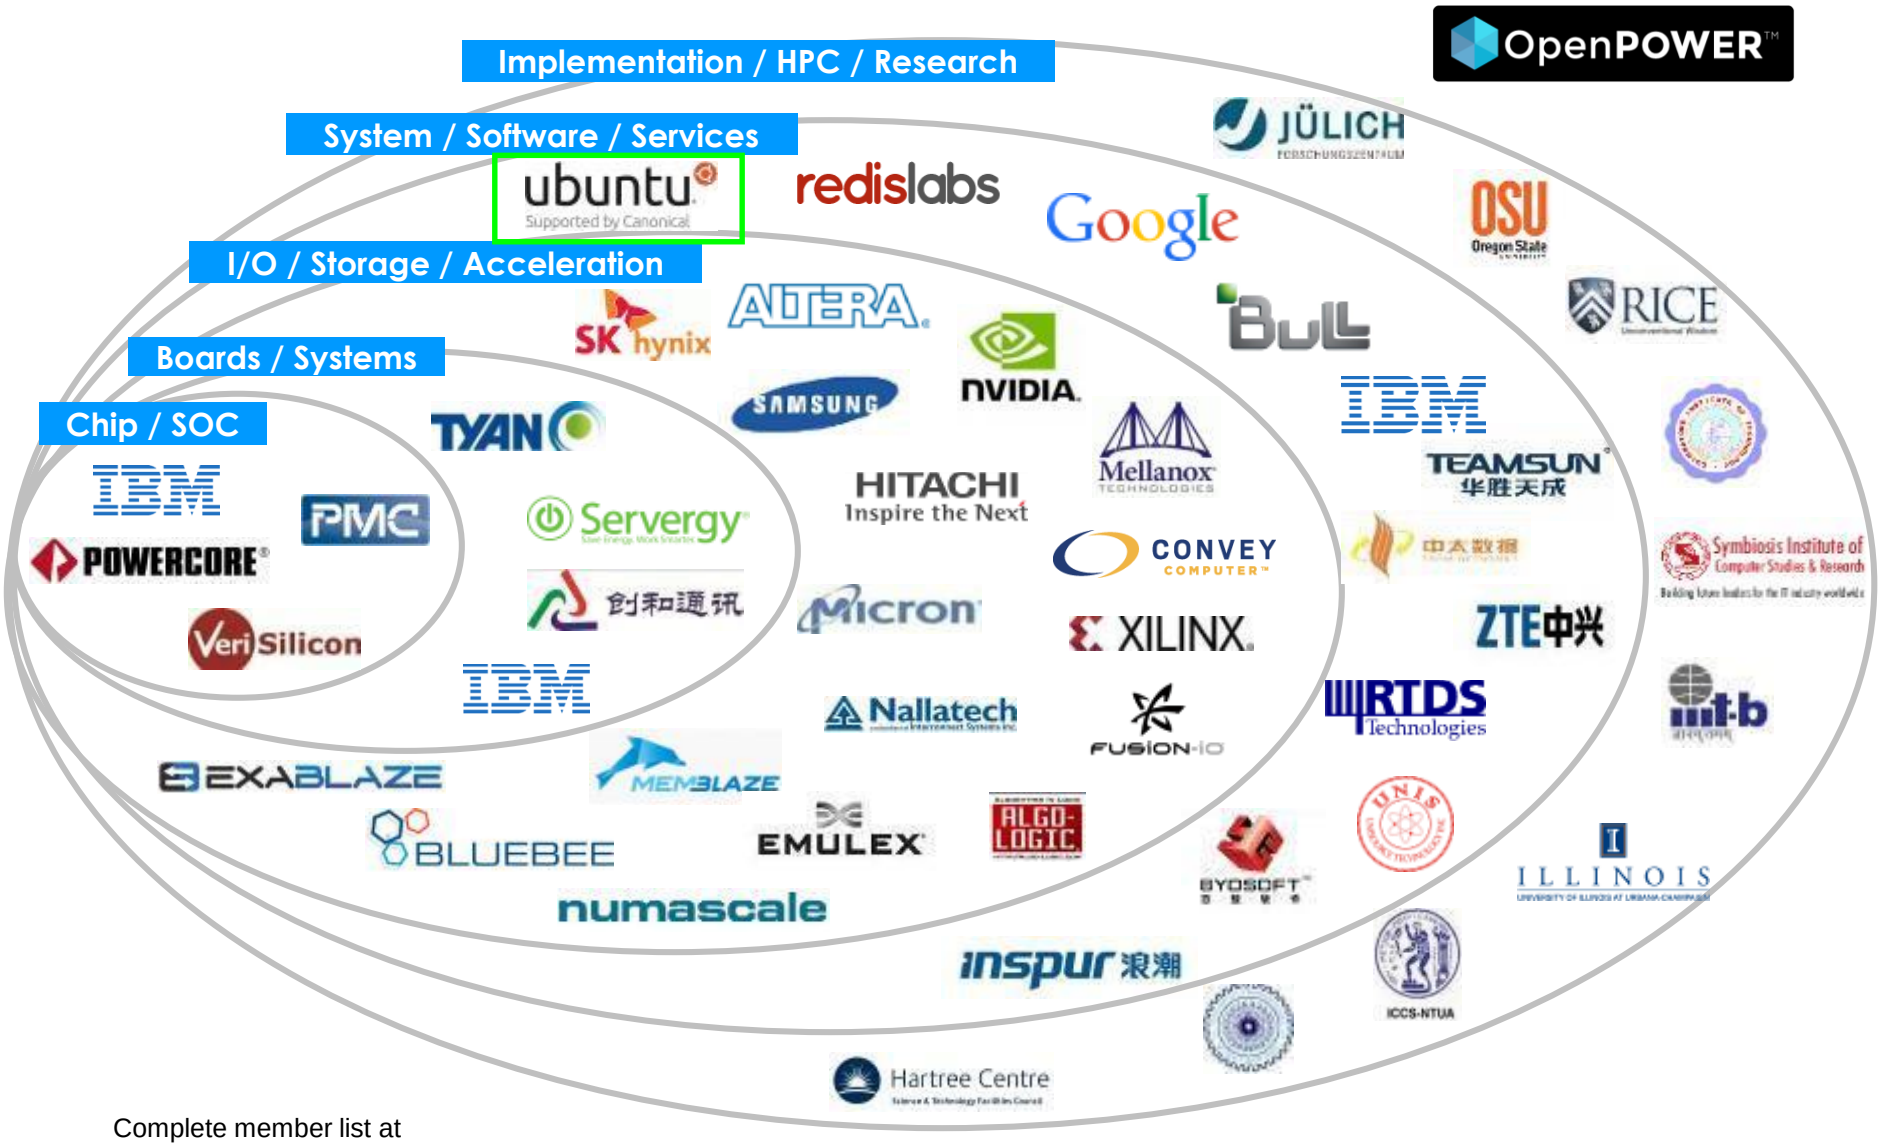

Complete member list at
www.openpowerfoundation.org

OpenPower Article in eWEEK - March 2015

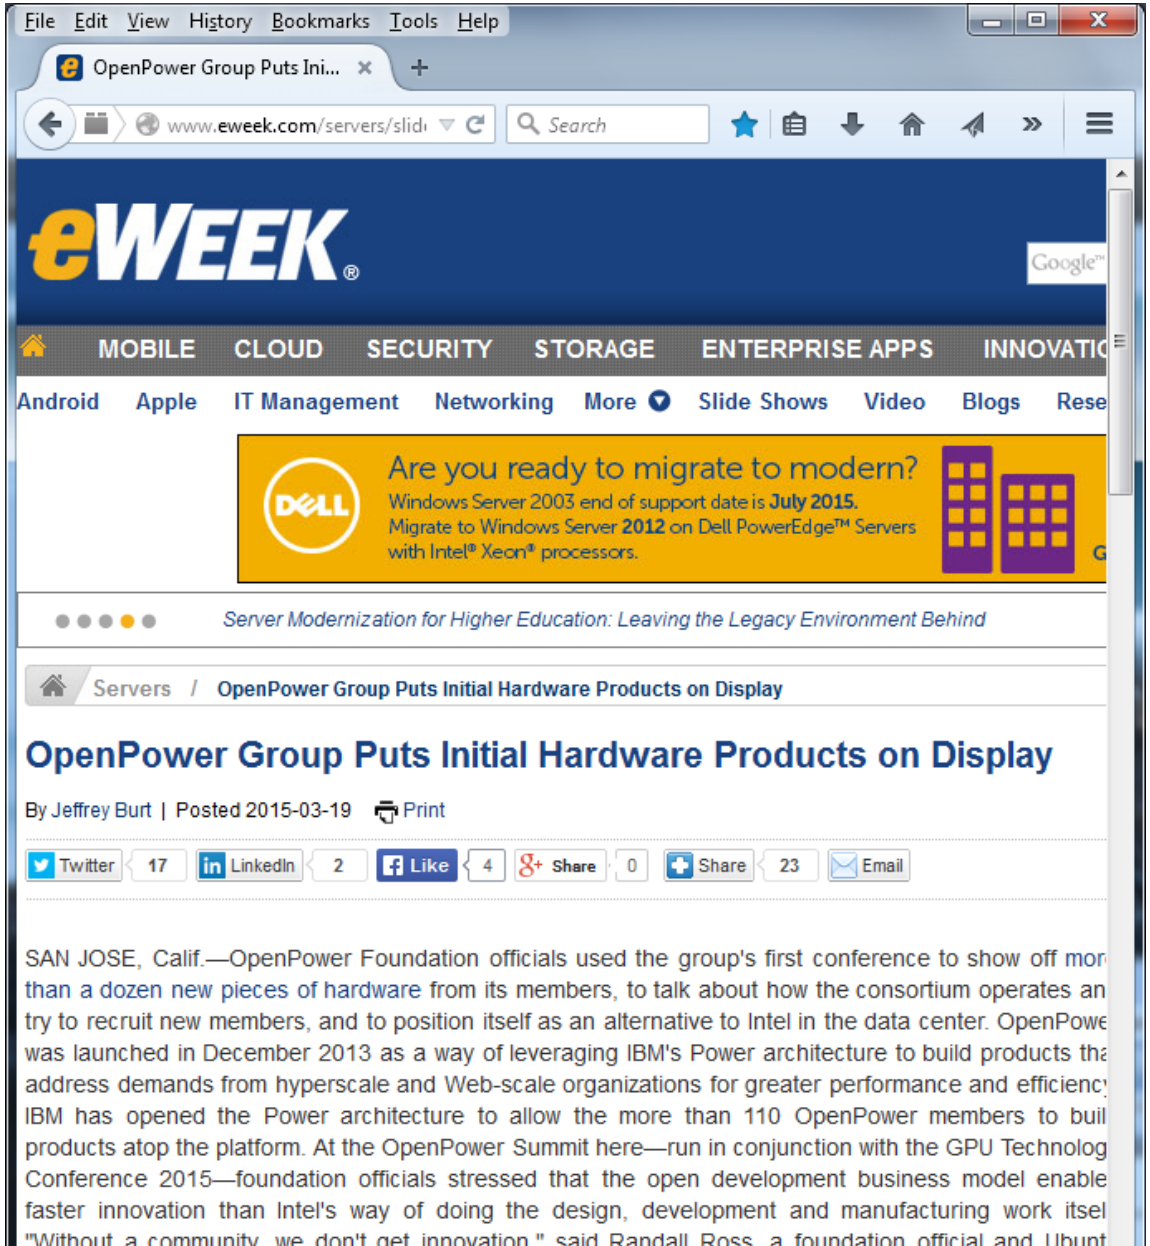

1Q 2014 Portfolio: POWER7/POWER7+

IBM PureFlex System

2Q 2014 Portfolio: POWER8/POWER7/POWER7+

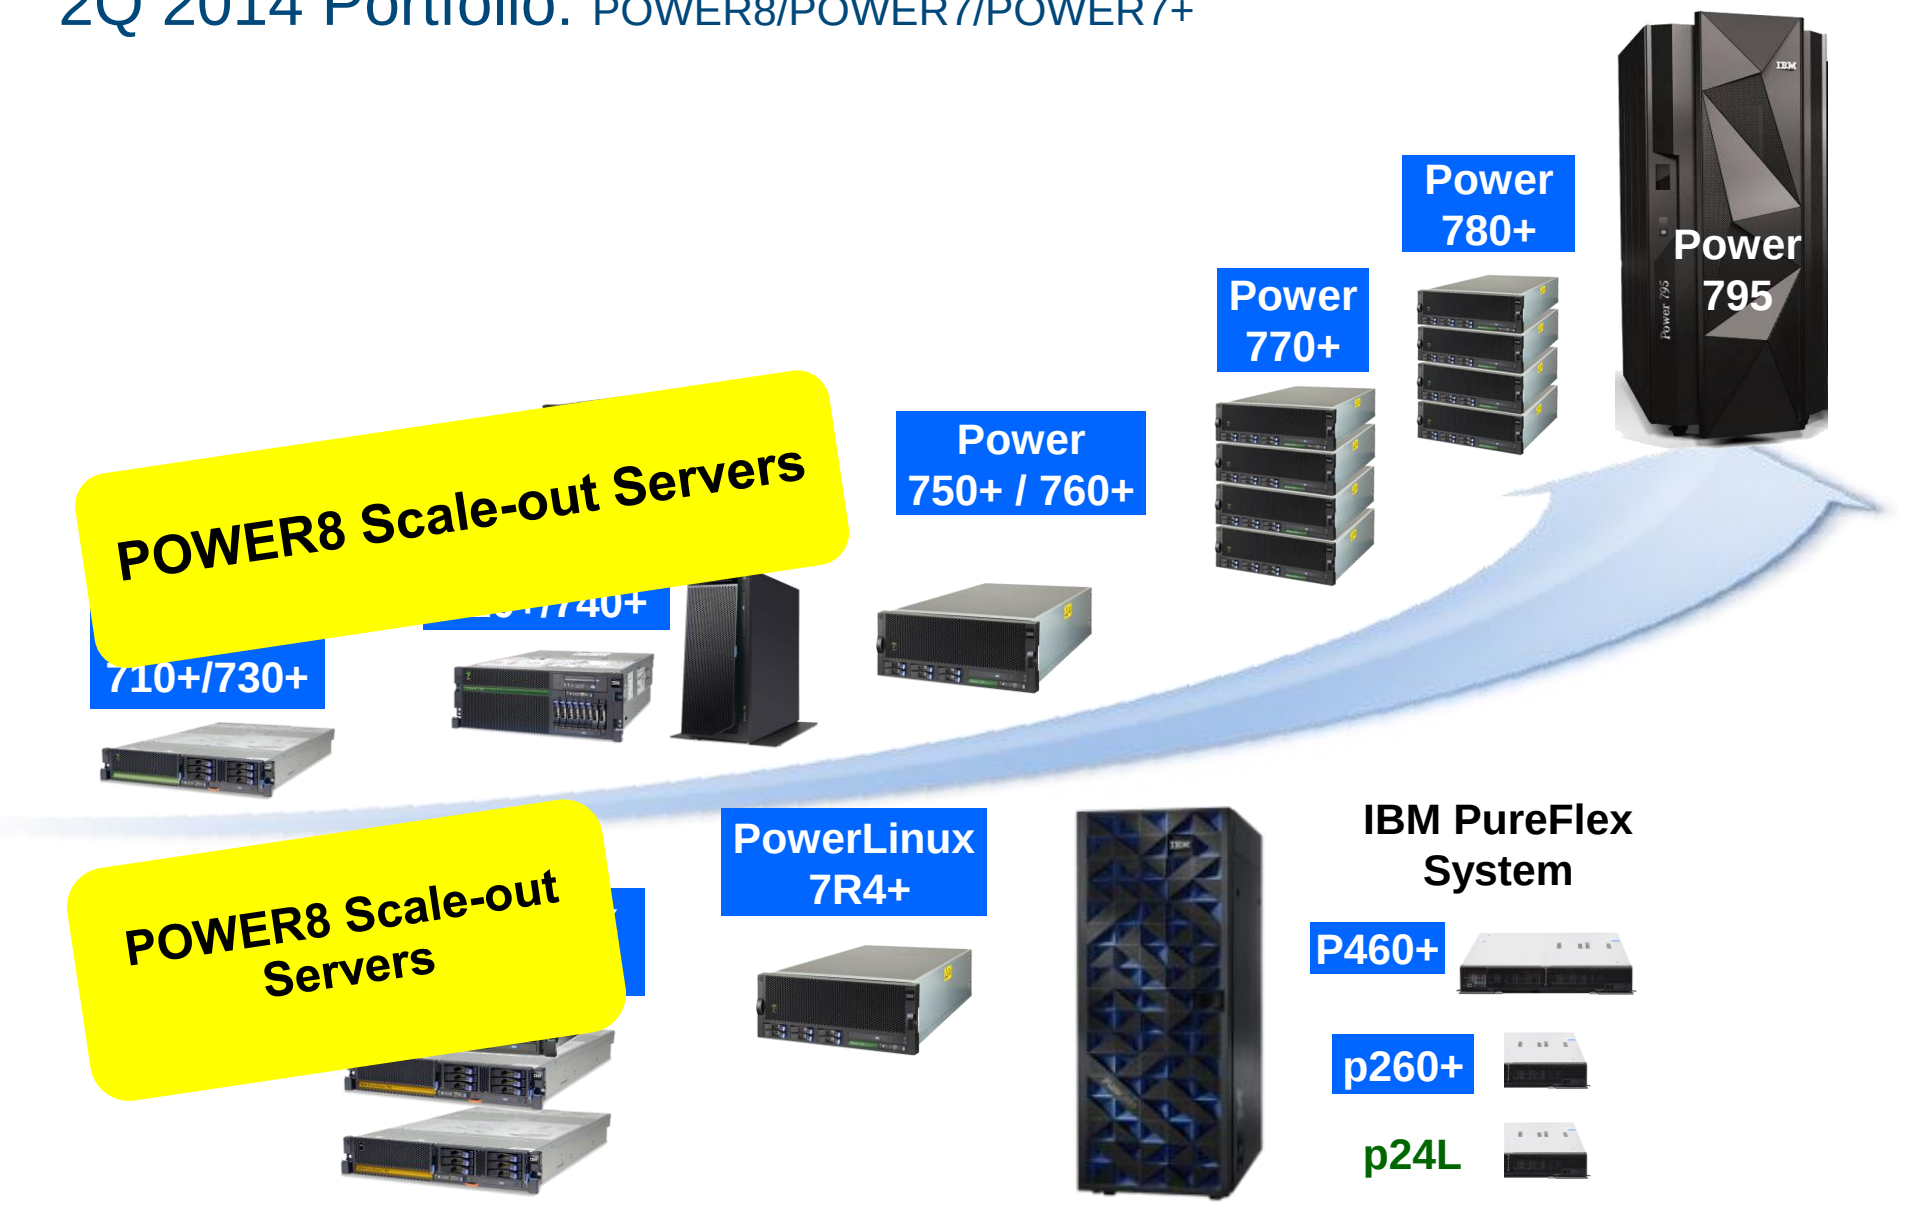

4Q 2014 Portfolio: POWER8/POWER7/POWER7+

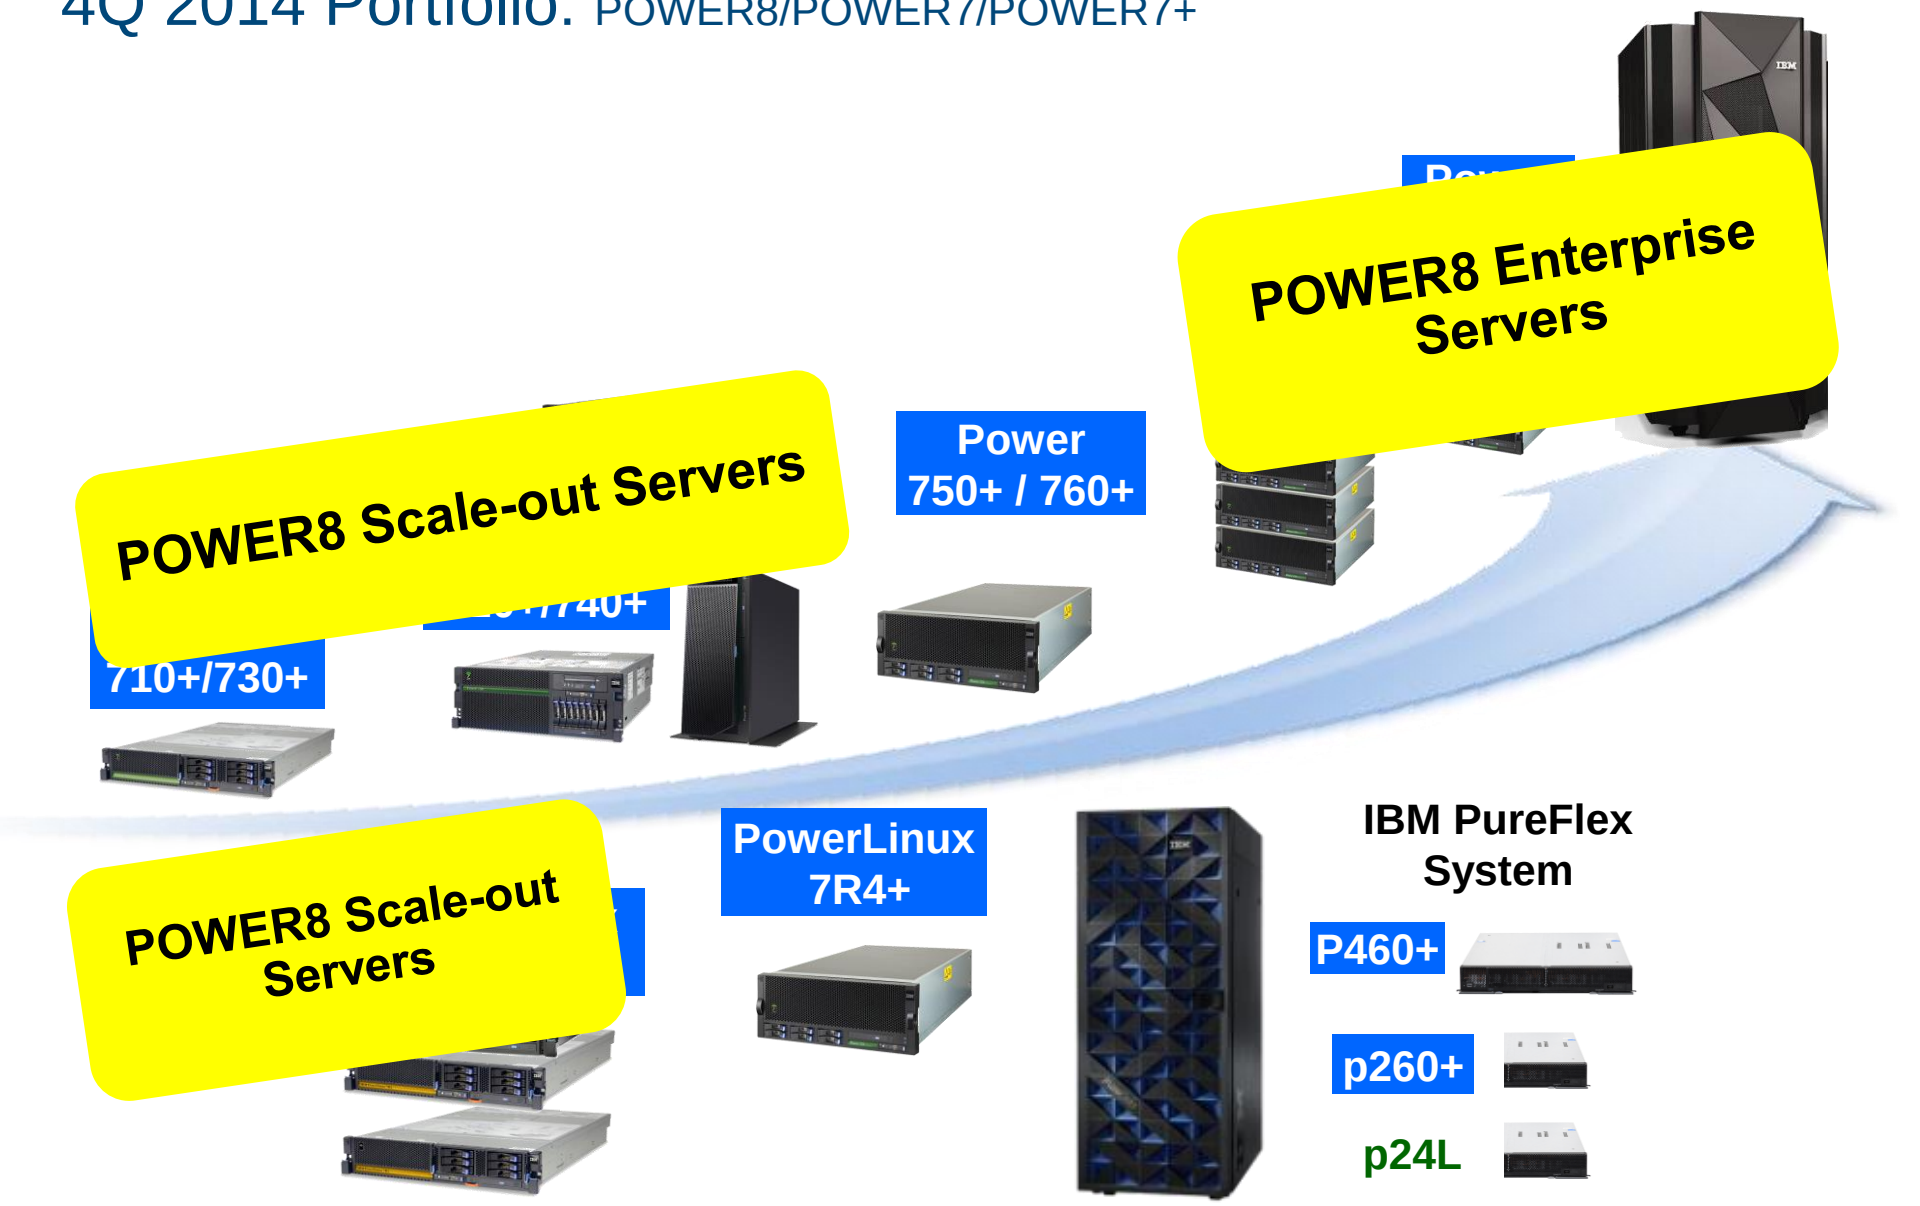

2Q 2015 Portfolio: POWER8 /POWER7/POWER7+

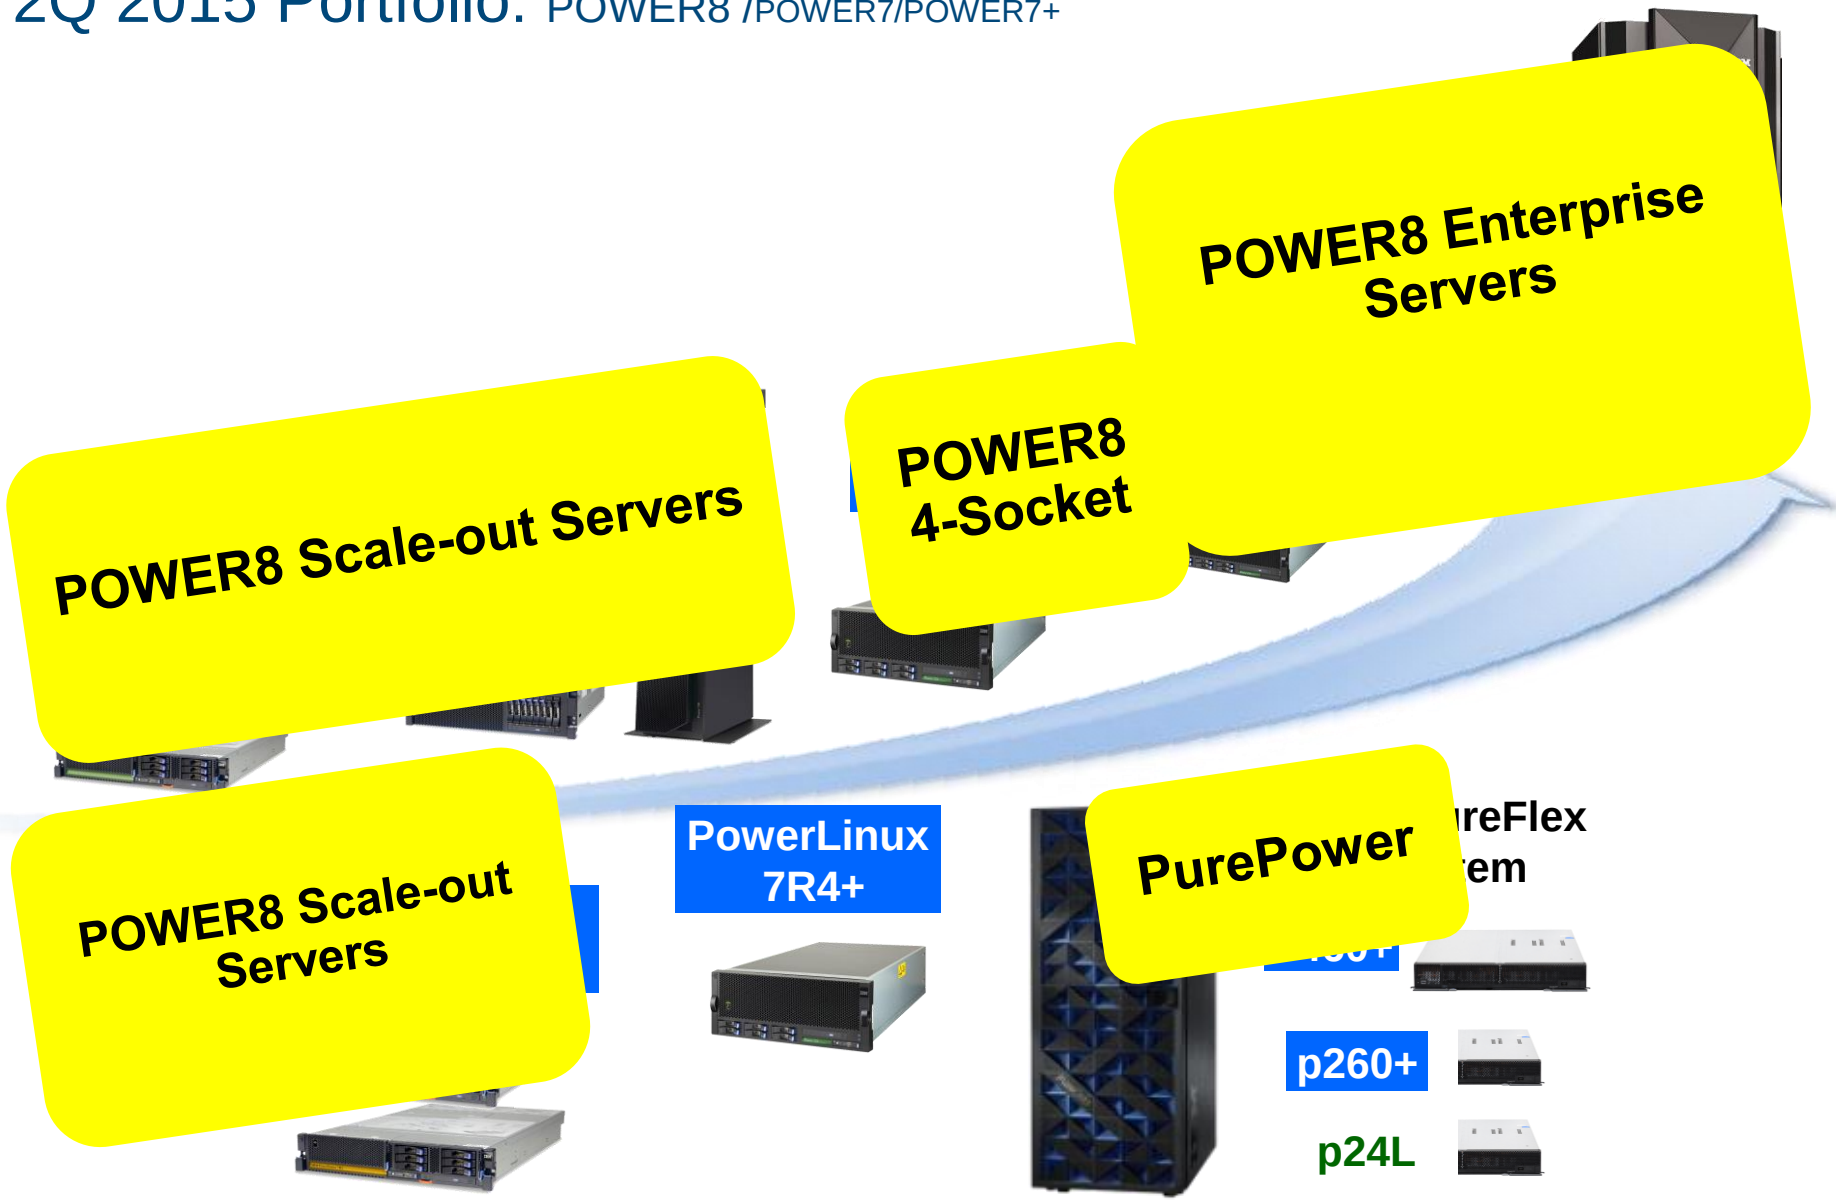

4Q2015 Portfolio: POWER8

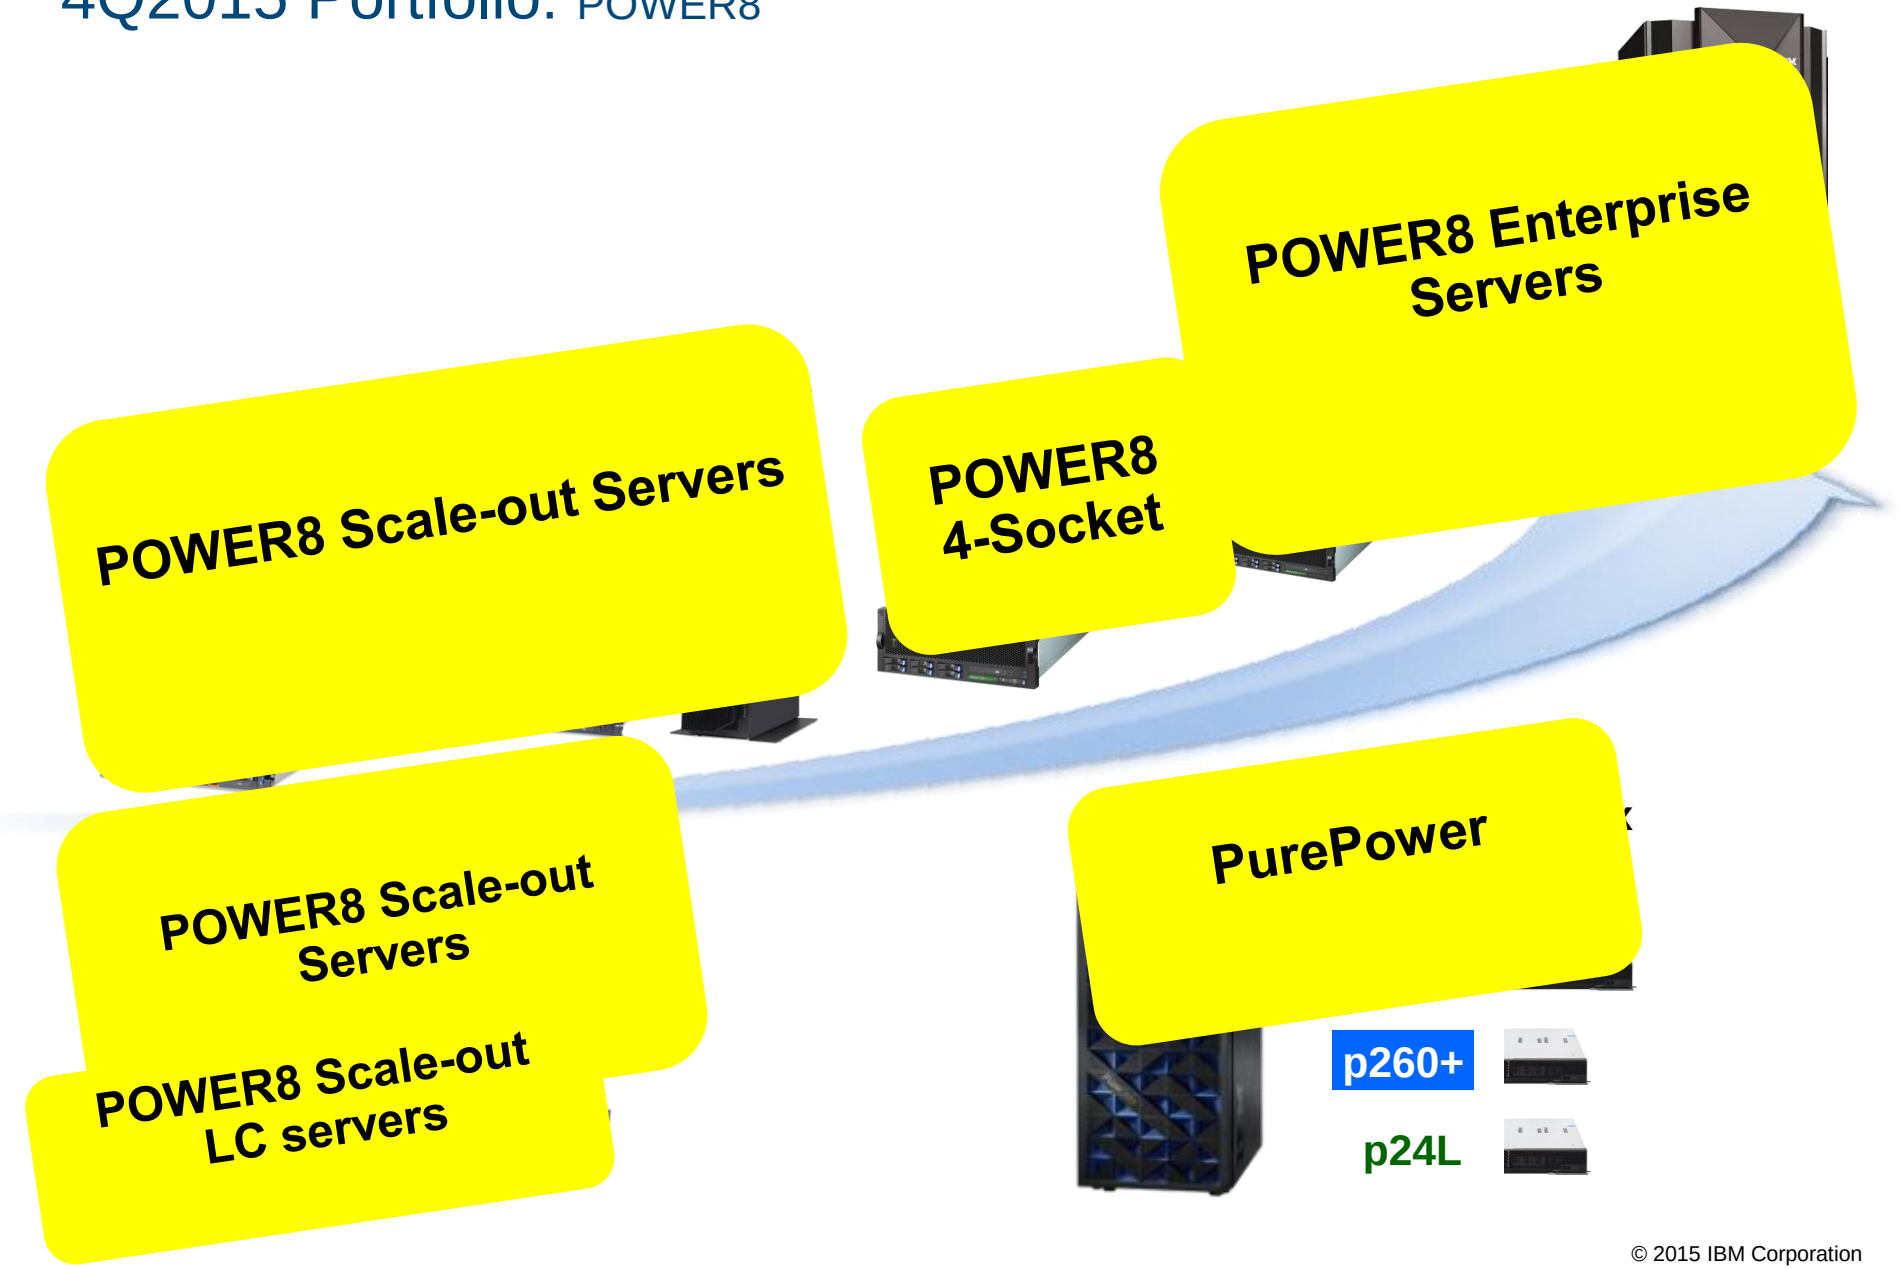

POWER8 Product Line 4Q 2015

Power E880

Power E870

Power E850

Power S814, S824

Power S822

Power S824L

Power S812L S822L

S812LC S822LC

PurePower
Built with Scale-out servers

POWER8 Product Line – IBM i Focus

Power E880

Power E870

Power E850

Power S814, S824

Power S822

Power S824L

Power S812L S822L

S812LC S822LC

PurePower
Built with Scale-out servers

Scale-out

Scale-out Servers (LC line not shown)

S824 (2 socket, 4U)

S824L (1 socket, 4U)

S814 (1 socket, 4U)

S812L (1 socket, 2U)

S822 (2 socket, 2U)

S822L (2 socket, 2U)

Linux Cluster Models Announcing October 2015

Bringing the benefits of the OpenPOWER Community to the IBM Power Systems Portfolio

New!

Power S812LC

Linux-only server with rich memory and storage
Great fit for applications such as Hadoop

2U, 1 socket up to 10 cores, up to 1TB memory, up to 14 disk/SSD (up to 112TB), 4 PCIe slots
Power KVM or bare metal

Power S822LC

Linux-only server with rich memory a focus on GPU capability

2U, 2 socket up to 20 cores, up to 1TB memory, up to two GPU, 5 PCIe slots
Power KVM or bare metal

Scale-out Servers with IBM i support (or AIX/Linux support)

S814 (1 socket, 4U)
- 4 core (P05 IBM i software tier)
- 6 or 8 core (P10 IBM i software tier)

S824 (2 socket, 4U)
- 6/12 or 8/16 or 24 core
- P20 IBM i software tier

S822 (2 socket, 2U)
- 6/8/10 or 12/16/20 core
- P10 IBM i software tier

POWER8 4U 2-Socket Server

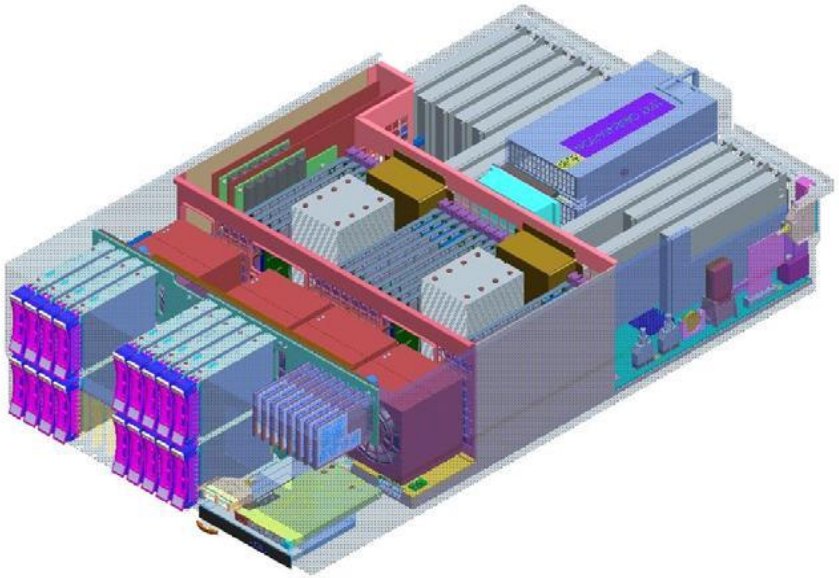

- **4U Form Factor**
- **1 to 2 Sockets (Up 24 Cores)**
- **Up to 2TB Memory**
  - **With 8GB DDR3 DIMMs**
- **11 PCIe Gen3 Slots / Hot Plug**
  - **Hot plug support**
  - **1 PCIe slot reserved for Ethernet**
- **Up to 18 SFF Bays**
  - **Optional RAID support**
- **Up to 8 1.8” SSDs**
- **Integrated PCIe Gen3 SAS controller w/ optional write cache**

2S2U Scale-out System

Power S822

- **Single Socket populated**
  - **Cores: 6 (3.8 GHz) or 10 (3.4 GHz) or 8 (4.1 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) or 16 (4.1 GHz)**
  - **Memory: Up to 1 TB ***
  - **PCIe Slots: 9 PCIe Gen3 LP (Hotplug)**

- **Ethernet: Choice of 1Gb or 10Gb PCIe adapter**
- **Integrated ports: USB (4), Serial (1), 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 (IBM i added Oct 2015)**

*** Half this memory maximum with 4.1GHz cores (no 64GB DIMMs) due to air cooling limitations. See following chart with additional cooling limitations and for SOD for water cooling.**

POWER8 4U Scale-out System

Power S814/S824

- **Form Factor: 4U or Tower***
- **One or two sockets ****
  - **Cores: 1 socket: 4–8, 2 socket: 6-24**
  - **GHz: 3.0 – 4.1**
- **Memory: Up to 1TB per socket*****
- **PCIe Slots: 7 or 11 PCIe Gen3 Full-high (Hotplug)**
  - **PCIe I/O drawer: adds up to 20 more slots (Hotplug)****
- **Integrated Storage**
  - **POWERFUL SAS controllers – up to 7GB cache**
  - **12 – 18 SFF Bays*** + 8 SSD Bays******
  - **Optional add 24 more SAS bays*****
  - **DVD**
- **Hypervisor: PowerVM**
- **OS: AIX, IBM i, Linux**
  - **IBM i 7.1 or later (P05 – P20 software tier)**

3 Yr Warranty

*** Tower only for 4-core or 6-core S814**
**** See model specifics for valid combinations of sockets, cores, GHz**
***** 4-core S814 has entry level server limitation**
****** 8 SSD bays only with 2-socket server**

POWER8 1S4U Scale-out System

Power S814

- **Form Factor: 4U or Tower**
- **Single Socket**
  - **Cores:** **4 (3.0 GHz), 6 (3.0 GHz) or 8 (3.7 GHz)**
  - **Memory:** **Up to 512 GB (6-core or 8-core)**
  - **Slots:** **7 PCIe Gen3 Full-high (Hotplug)**

- **PCIe Gen3 I/O Drawer adds 10 PCIe slots (6-core or 8-core)**
- **Ethernet: Quad 1 Gbt in PCIe slot**
- **Integrated ports: USB (4/5), Serial (1), HMC (2)**
- **Internal Storage**
  - **DVD**
  - **12 SFF Bays -- Split Backplane: 6 + 6**
  - **or 18 SFF Bays with 7GB write cache**

- **Hypervisor: PowerVM**
- **OS: AIX, IBM i 7.1 or later (P05/P10 software tier), Linux**

3 Yr Warranty

POWER8 2S4U Scale-out System

Power S824

- **Single Socket populated**
  - **Cores:** 6 (3.8 GHz) or 8 (4.1 GHz)
  - **Memory:** Up to 1TB
  - **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 2TB
  - **Slots:** 11 PCIe Gen3 full-high (Hotplug)

3 Yr Warranty

- *** PCIe Gen3 I/O drawer adds up to 20 PCIe slots**
- **Ethernet: Quad 1 Gbt in PCIe slot**
- **Integrated ports: USB (4/5), Serial (1), HMC (2)**
- **Internal Storage**
  - **DVD**
  - **12 SFF Bays -- Split Backplane: 6 + 6**
  - **or 18 SFF bays & 8 SSD bays with 7GB write cache**
- **Hypervisor: PowerVM**
- **OS: AIX, IBM i 7.1 or later (P20 software tier), Linux**

POWER8 4U 2-Socket Comparison

	Power 740	POWER8
Processor	POWER7+	**POWER8 4U**
Sockets	2	**2**
Cores	12 / 16	**16 / 20 / 24**
Maximum Memory	1 TB @ 1066 MHz	**2 TB @ 1600 MHz**
Memory Cache	No	**Yes**
Memory Bandwidth	68 GB/sec	**200 GB/sec**
Memory DRAM Spare	No	**Yes**
IO Expansion Slots	Dual GX++	**4 PCIe x16 G3 (SOD)**
PCIe slots	5 PCIe x8 FH 4 PCIe x8 HH (opt)	**7 PCIe x8 FH** **4 PCIe x16 FH**
PCIe Hot Plug Support	No	**Yes**
IO bandwidth	60 GB/sec	**192 GB/sec**
Base* Ethernet ports	Four 1 Gb	**Four 1Gb or Two 10 Gb**
SSD-only bays	N/A – use SFF bays	**Up to 8 1.8-inch bays**
SFF bays	6	**12 or 18**
Integrated SAS controller	PCIe Gen1 base	**PCIe Gen3 base**
SR-IOV Support	No	**Yes**
Service Processor	Generation 1	**Generation 2**

*** Base uses one PCIe slot**

2S2U Scale-out System

Power S822

- **Single Socket populated**
  - **Cores: 6 (3.8 GHz) or 10 (3.4 GHz) or 8 (4.1 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) or 16 (4.1 GHz)**
  - **Memory: Up to 1 TB ***
  - **PCIe Slots: 9 PCIe Gen3 LP (Hotplug)**

- **Ethernet: Choice of 1Gb or 10Gb PCIe adapter**
- **Integrated ports: USB (4), Serial (1), 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 (IBM i added Oct 2015)**

*** Half this memory maximum with 4.1GHz cores (no 64GB DIMMs) due to air cooling limitations. See following chart with additional cooling limitations and for SOD for water cooling.**

Power S822 and IBM i

- IBM i 7.2 TR3 or IBM i 7.1 TR11
- Special terms/conditions
- Only P10 software tier (P20 used on 2-socket S824)
- S822 with or without PurePower
- Max 2 cores per partition -- multiple partitions per server supported
- All IBM i I/O virtualized through VIOS. No non-VIOS I/O.
  - No Async/Bisync PCIe adapters (can use IBM i 7.2 with OEM LAN-attached box for Async including Fax/400)
  - No Crypto cards
  - If using disk/SSD in system unit or EXP24S I/O drawer for IBM i it must use 4k byte sectors, not 5xx byte sectors
- CBU feature #0444 not announced Oct 2015

Plus normal 2U limitations of smaller physical server vs 4U server -- smaller memory, fewer PCIe slots, fewer PCIe I/O drawers, subset of PCIe I/O card as full-high cards won't fit including any PCIe SAS adapters with write cache.

Minimum code levels required: IBM i 7.2 TR 3 or IBM i 7.1 TR 11, VIOS 2.2.3.52, FW 840, if using HMC then HMC 840

Enterprise

Power E850 / E870 / E880 Enterprise Characteristics

- ✓ **Increased performance and scale compared to previous generation servers**
- ✓ **Capacity on Demand**
  - **Elastic CoD**
  - **Utility CoD**
  - **Trial CoD**
- ✓ **Power Integrated Facility for Linux**
- ✓ **Active Memory Mirroring for Hypervisor ***
- ✓ **E870/E880: Power Enterprise Pools**
- ✓ **E870/E880: PowerVM Enterprise included**
- ✓ **E880: Large block no-charge days with initial purchase**

* optional for E850

Capacity on Demand for E870/E880

- Processor cores and/or memory activations

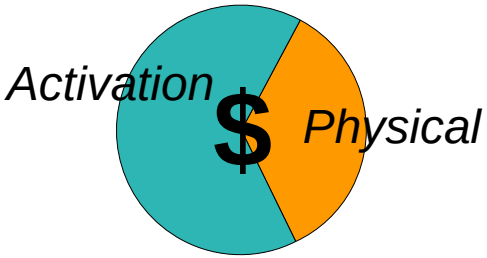

- Separating activation pricing from physical hardware
  - Makes redundant hardware more affordable
  - Eases budgeting – buy when needed
  - Reduces planned downtime outages for upgrades

- Elastic CoD / Utility Cod
  - Use activations “by the day” or “by the minute” as needed
  - Elastic: pay quarterly if use; Utility pay after usage
  - Elastic example for Power E870 – 4.19 GHz
    - 1 add’l core + 16GB memory
    - 1 add’l cores licensing for IBM i and for other IBM Power System’s controlled software
    - 1 add’l core’s worth of 5250
    - ONLY $196 per day

Prices are USA list price and are subject to change Reseller prices may vary.
Note if other GHz core are used, list price can vary slightly. Add’l memory is $1 per GB.
5250 inclusion assumes at least one 5250 Enterprise Enablement already on the server.

Power Enterprise Pools for E870/E880

Power E870 & 770+ Pool

- **Max 48 servers**
- **In the same country**
- **Running 7.8 firmware or later**
- **Running any supported OS**
- **Max 1000 partitions**

Power E880 & 780+ & 795 Pool

- **Max 32 servers**
- **In the same country**
- **Running 7.8 firmware or later**
- **Running any supported OS**
- **Max 1000 partitions**

Move processor and memory activations within a defined pool
Move quickly, dynamically, non-disruptively
Move at your convenience – without contacting IBM

PLUS also move specific Power software licensing entitlements

Power Integrated Facility for Linux

- For enterprise servers with processor cores and memory not permanently activated
  - Model E870 & Model E880
  - Model 770 & Model 780 & 795
- Package of activations & licensing
- Usable only for Linux workload
- Priced competitively vs. equivalent Linux stand-alone server
- Savings via lower priced activations, maintenance & licensing for Linux-only resource

Power IFL
Price/Performance Example
3yr TCA

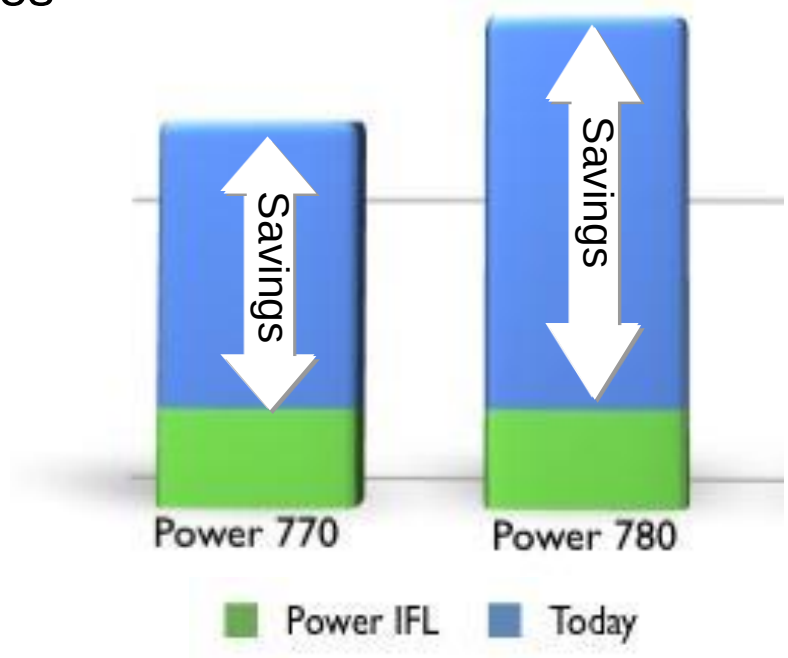

Example includes:

- ***4 cores activation***
- ***32 GB activations***
- ***4 PowerVM for Linux entitlements***
- ***4 RedHat Linux***
- ***3 yr 24x7 HWMA/SWMA***

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

E850

Why not provide IBM i support on Power E850 ?

750 vs S824 Hardware Comparison Highlights

- CPW: Given most IBM i clients who are using a Power 750 are using smaller configurations, there is lots of CPW growth in Power S824
- Memory: S824 has 4x more memory capacity – and it is 50% faster memory (1600MHz with L4 cache VS 1066MHz)
- PCIe:
  - S824 has newer, higher bandwidth PCIe Gen3 slots
  - Without considering PCIe I/O drawer, S824 has more PCIe slots
  - With PCIe I/O drawer, S824 has a very good number of high bandwidth slots (31), though less than the number of 750 2-4 socket maximum
  - S824 supports SR-IOV NIC
- Integrated Storage
  - S824 has optional integrated SAS controller with 7GB effectively write cache, more than 40X 750's 175MB.
  - S824 has up to 26 SAS bays in system unit for HDD/SSD. 3X more than 750
  - S824 can have up to 50 SAS bays run by integrated SAS controllers 6X more than 750

POWER7 750 (B mdl) vs POWER8 S824

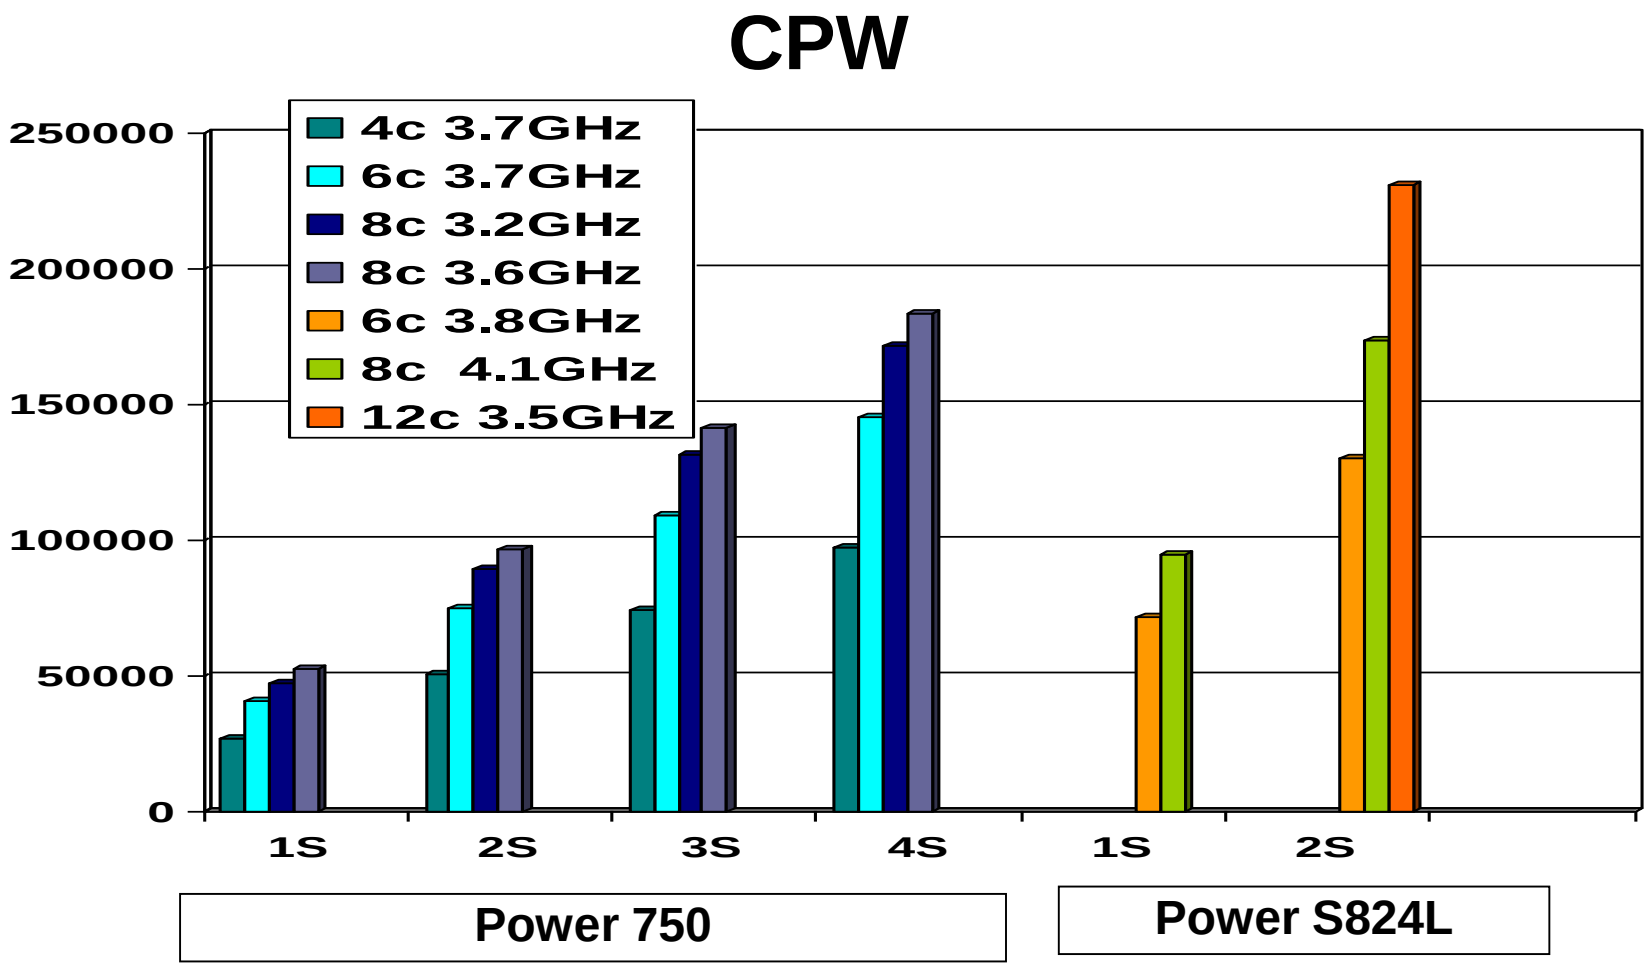

Except perhaps for the 32-core Power 750 – lot's of room for CPW growth.

Power E870 & E880 Servers

Power E870:

- Up to 80 cores
- 256 to 8TB Memory
- 1 or 2 nodes per system

Power E880:

- Up to 192 cores @ 4.02 GHz
- 256 to 16TB Memory
- 1 to 4 nodes per system

AIX · IBM i for Business · Linux

PowerVM

- ✓**Enterprise RAS** even for 1-node system
  - System Control Unit (“midplane”)
- ✓**Medium software tier**
- ✓**24x7 1-year warranty**

POWER8 Enterprise Solutions

Architecture

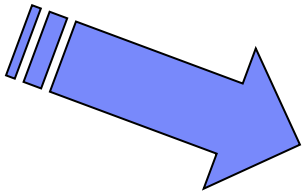

Packaging

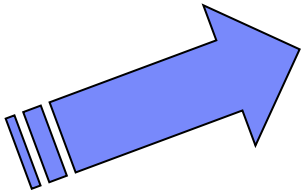

POWER8 Enterprise Solutions

No Primary Node

Midplane
- **Service Processors**
- **Clocks**
- **Oscillators**

Large Memory

Midplane

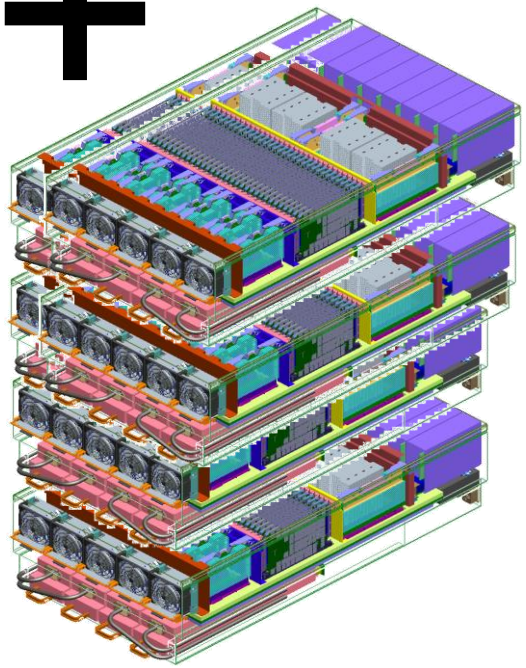

19” Rack
Modular design
Up to 4 CEC drawers
Blindswap IO Adapters

POWER8 Processor Enclosure (CEC Drawer)

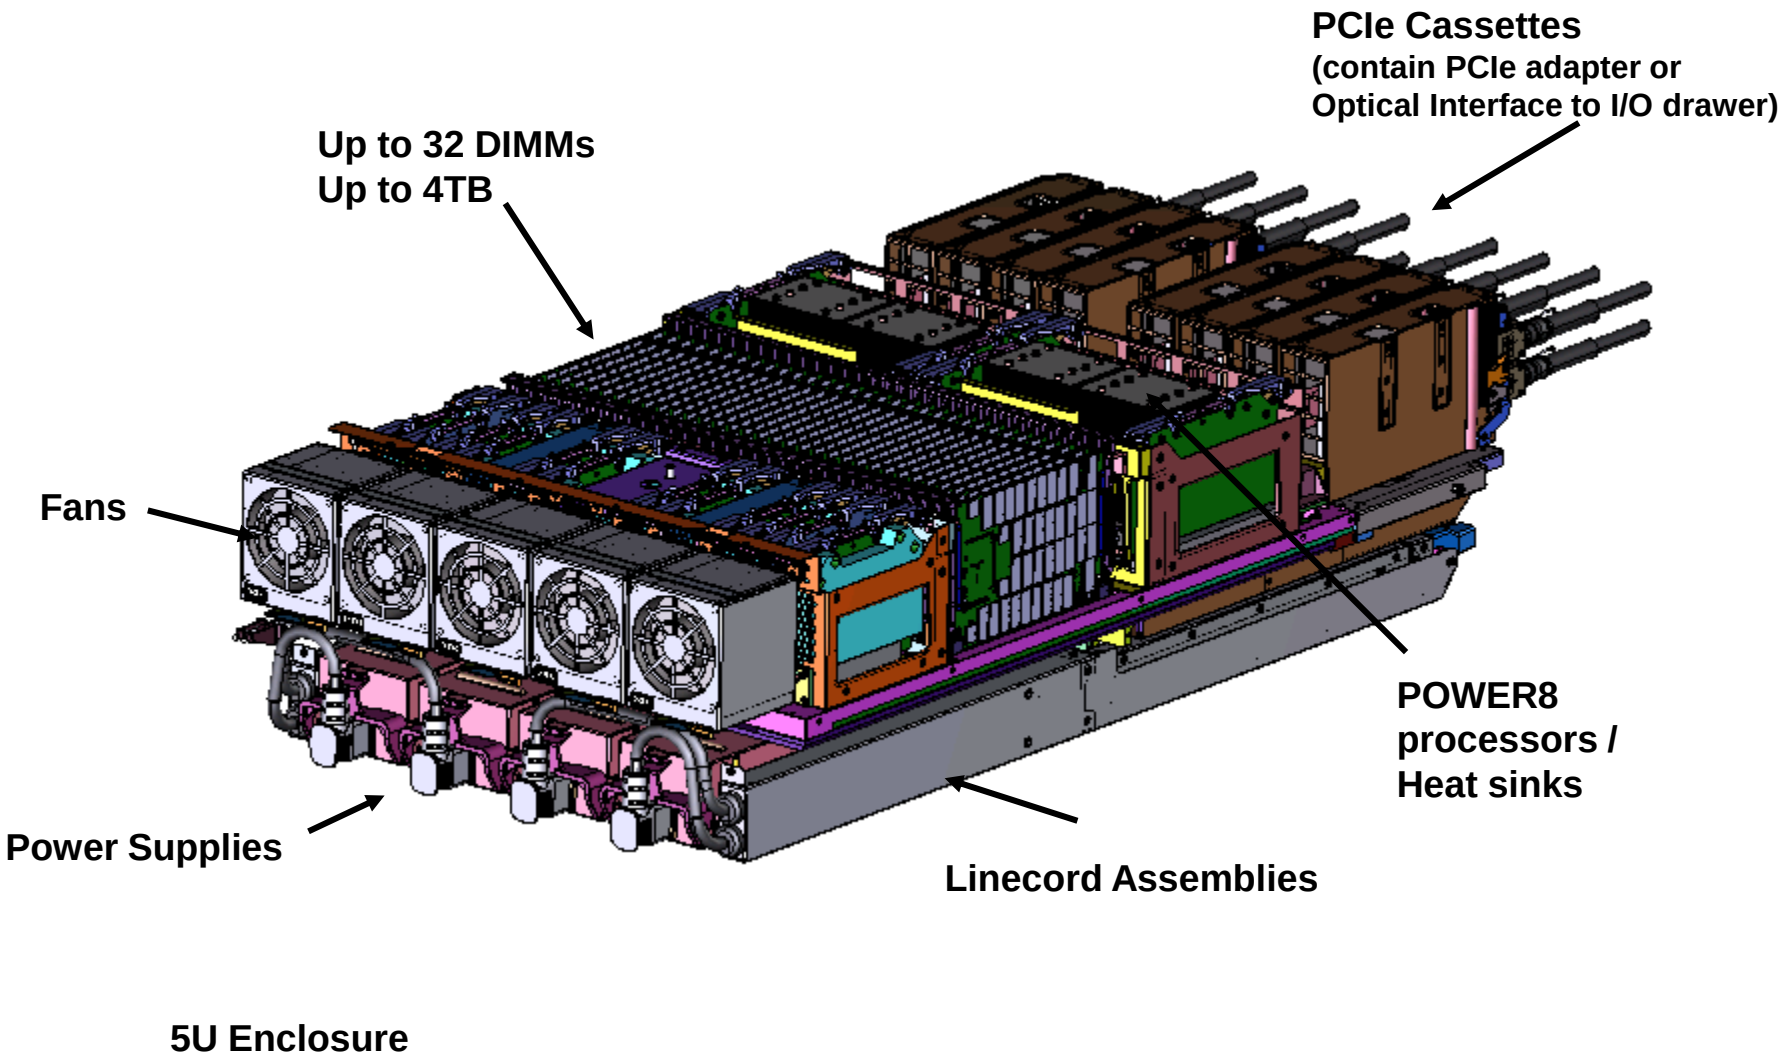

Power E870: 1 or 2 Nodes

12U in 19” rack

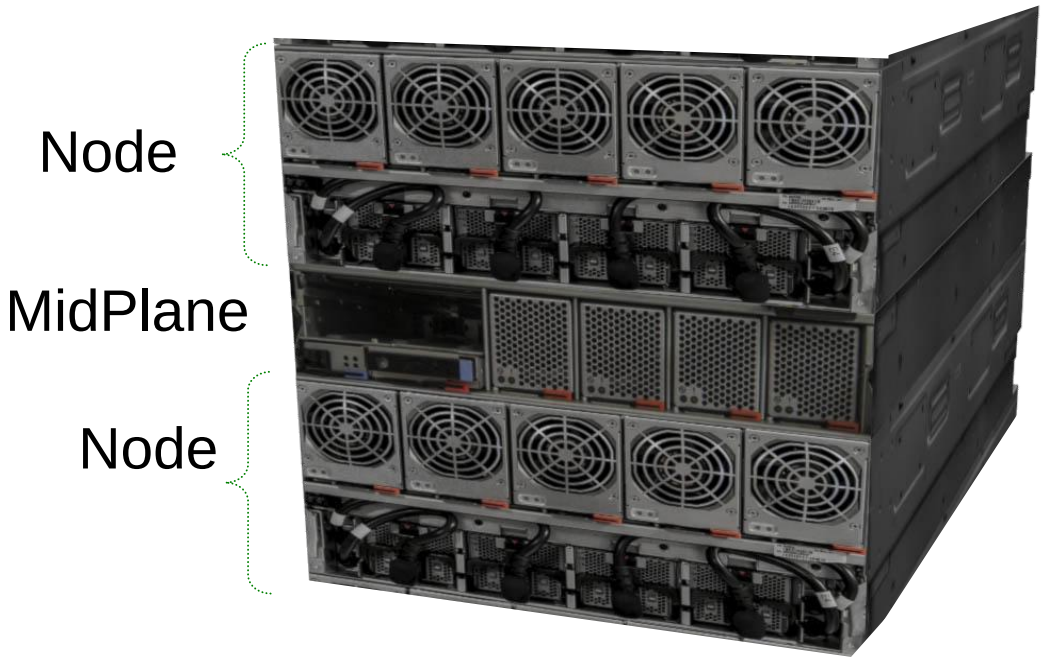

Power E880: 1, 2, 3 or 4 Nodes

22U in 19” rack

3- and 4-node configurations planned availability June 2015

Power Scale Perspective

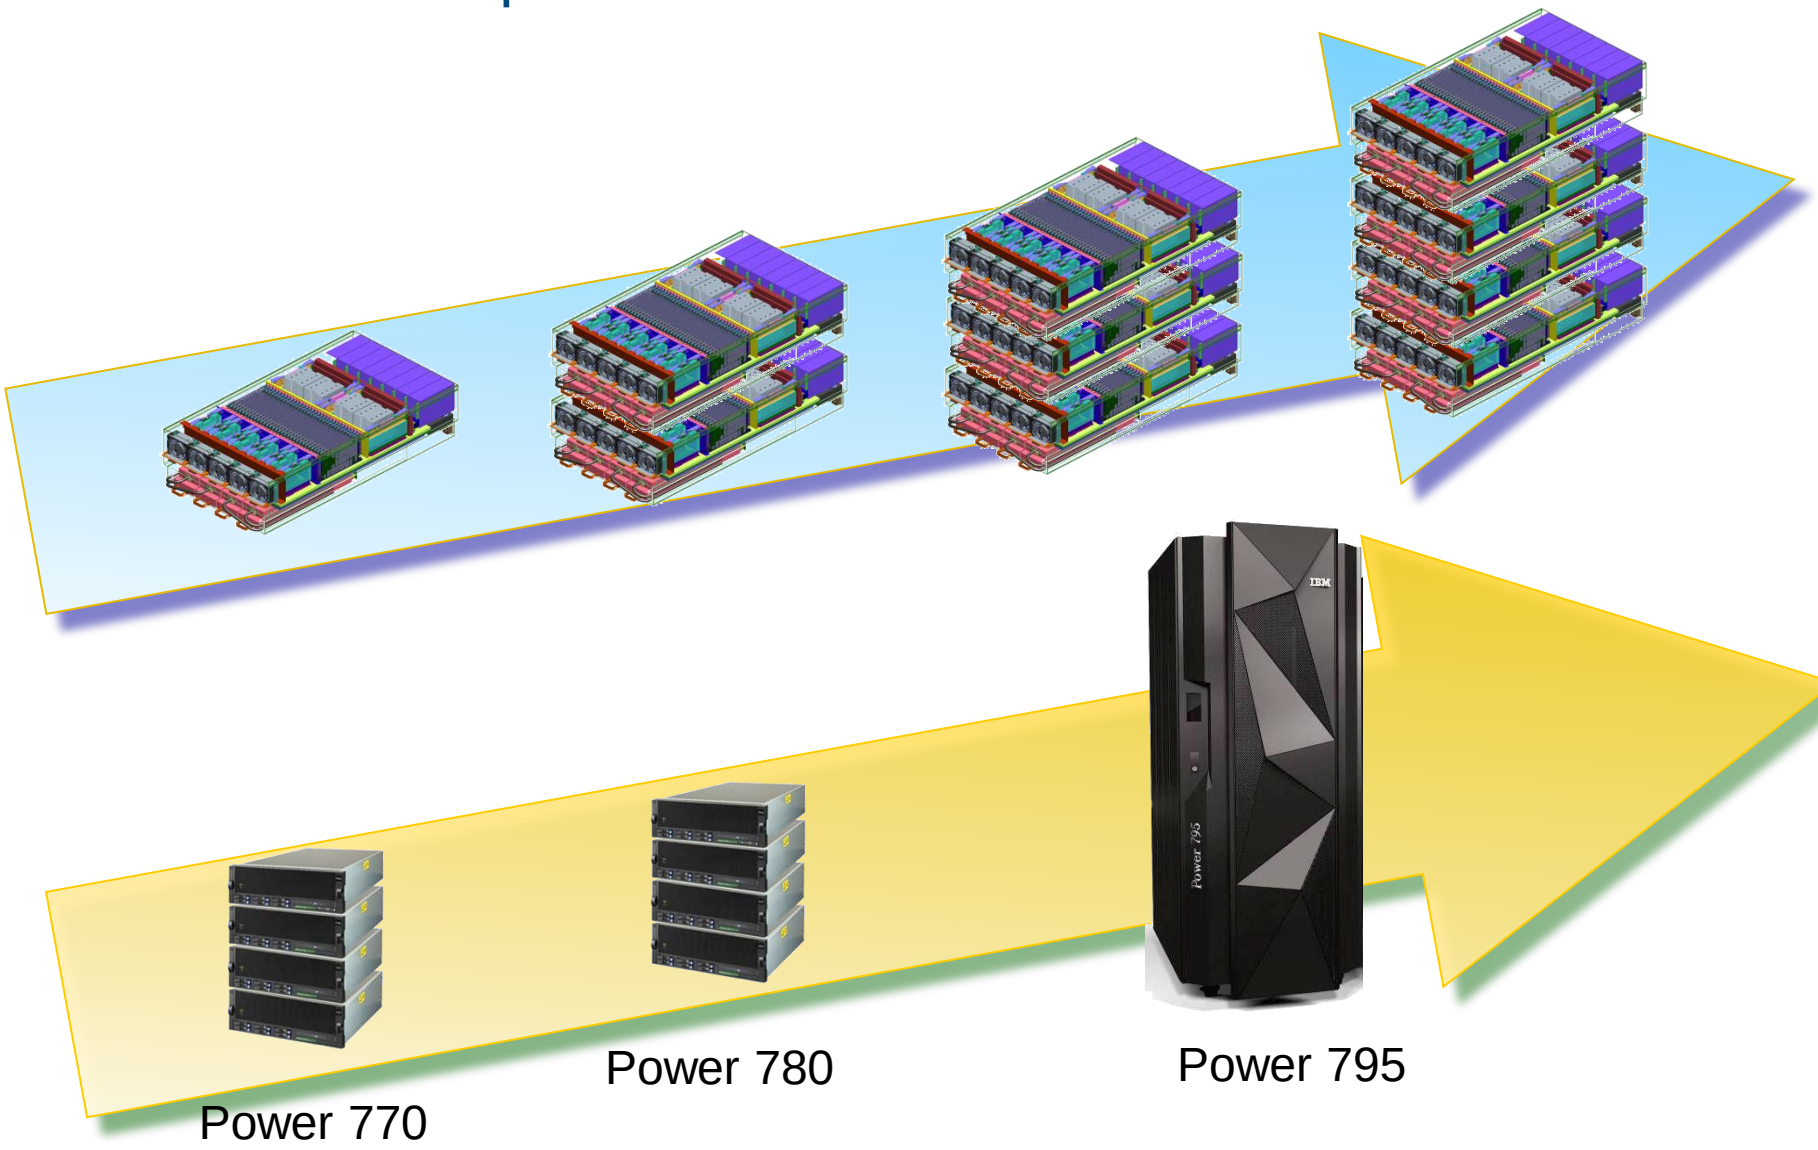

73

POWER8 E870 Compares: 2-Node/Drawer System

	9117-MMD Power 770	POWER8 Enterprise 2 CEC Node-drawer
CPU Sockets per Node/Drawer	4	4
Max processor nodes/Drawers	4	2
Max number sockets	16	8
Max Cores	64	64 or 80
Max Frequency	3.8 GHz	4 or 4.19 GHz
Max Memory	1 TB per CEC Drw	4 TB per CEC Drw
Memory per core	64 GB	62 or 50 GB
Memory Bandwidth (peak)	272 GB/s per CEC Drw	922 GB/s per CEC Drw
I/O Bandwidth (peak)	80 GB/s per CEC Drw (GX)	256 GB/s per CEC Drw (PCIe Gen3)
Max PCIe I/O drws	16 (4 per Node)	8 (4 per Node)
Max PCIe I/O Slots	160 - in IO drws 24 - internal	96 – in IO drws 0 - internal
IBM i levels supported	IBM i 6.1 and later	IBM i 7.1 and later

*** SOD = statement of direction – represents IBM plans shared publically. As always plans subject to change.**

POWER8 E880 Compares: 4-Node/Drawer System

	9117-MMD Power 770	9179-MHD Power 780	POWER8 Enterprise 4 CEC Node-drawer
CPU Sockets per Node/Drawer	4	4	4
Max processor nodes/Drawers	4	4	4
Max number sockets	16	16	16
Max Cores	64	128	192
Max Frequency	3.8 GHz	3.7 GHz	4.35
Max Memory	1 TB per CEC Drw	1 TB per CEC Drw	4 TB per CEC Drw
Memory per core	64 GB	32 GB	128 or 85 GB
Memory Bandwidth (peak)	272 GB/s per CEC Drw	272 GB/s per CEC Drw	922 GB/s per CEC Drw
Memory Bandwidth per core (peak)	17 GB/sec	8.5 GB/sec	19 GB/sec
I/O Bandwidth (peak)	80 GB/s per CEC Drw (GX)	80 GB/s per CEC Drw (GX)	256 GB/s per CEC Drw (PCIe Gen3)
Max PCIe I/O drws	16 (4 per Node)	16 (4 per Node)	16 (4 per Node)
Max PCIe I/O Slots	160 - in IO drws 24 - internal	160 - in IO drws 24 - internal	192
IBM i levels supported	IBM i 6.1 and later	IBM i 6.1 and later	IBM i 7.1 and later

PCIe I/O Drawer

POWER8 PCIe Gen3 IO Drawer

- 4U in standard 19” rack
- 6 or 12 PCIe Gen3 slots
  - Full-High, Full-Length
  - 1/3rd x16 slots and 2/3rd x8 slots
- Dual fiber optic cables connect to POWER8 CEC
  - Connects to x16 slot in CEC

IO Bandwidth (Comparing I/O Drawers)

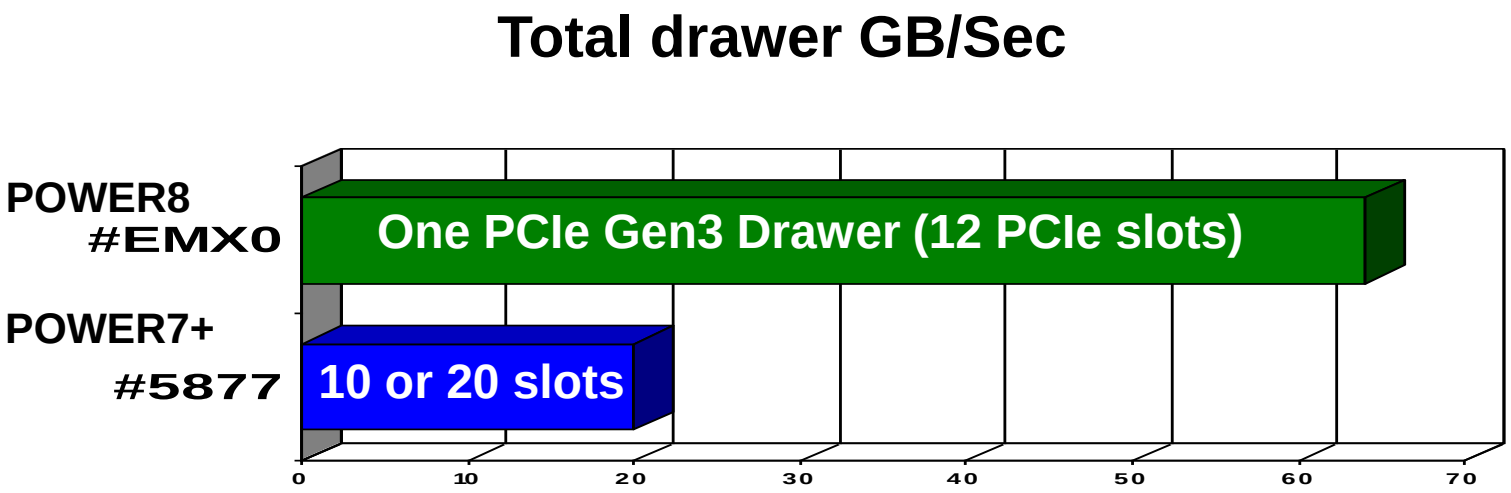

POWER8 PCIe-attached Gen3 I/O drawer has two fan-out modules and each fan-out module has 32GB/s

POWER7 12X-attached PCIe I/O drawer = #5877 or #5802
One or two #5877 or #5802 can share a single GX++ slot's 20GB/s bandwidth

I/O Loop (GX+) Technology Transitions

POWER5	POWER6	POWER7	POWER8 Not GX based
HSL / RIO w/ PCI-X	HSL / RIO w/ PCI-X 12X SDR w/ PCI-X 12X DDR w/ PCIe	--- 12X SDR w/ PCI-X 12X DDR with PCIe	--- --- --- PCIe3

Model Upgrades (Same Serial Number)

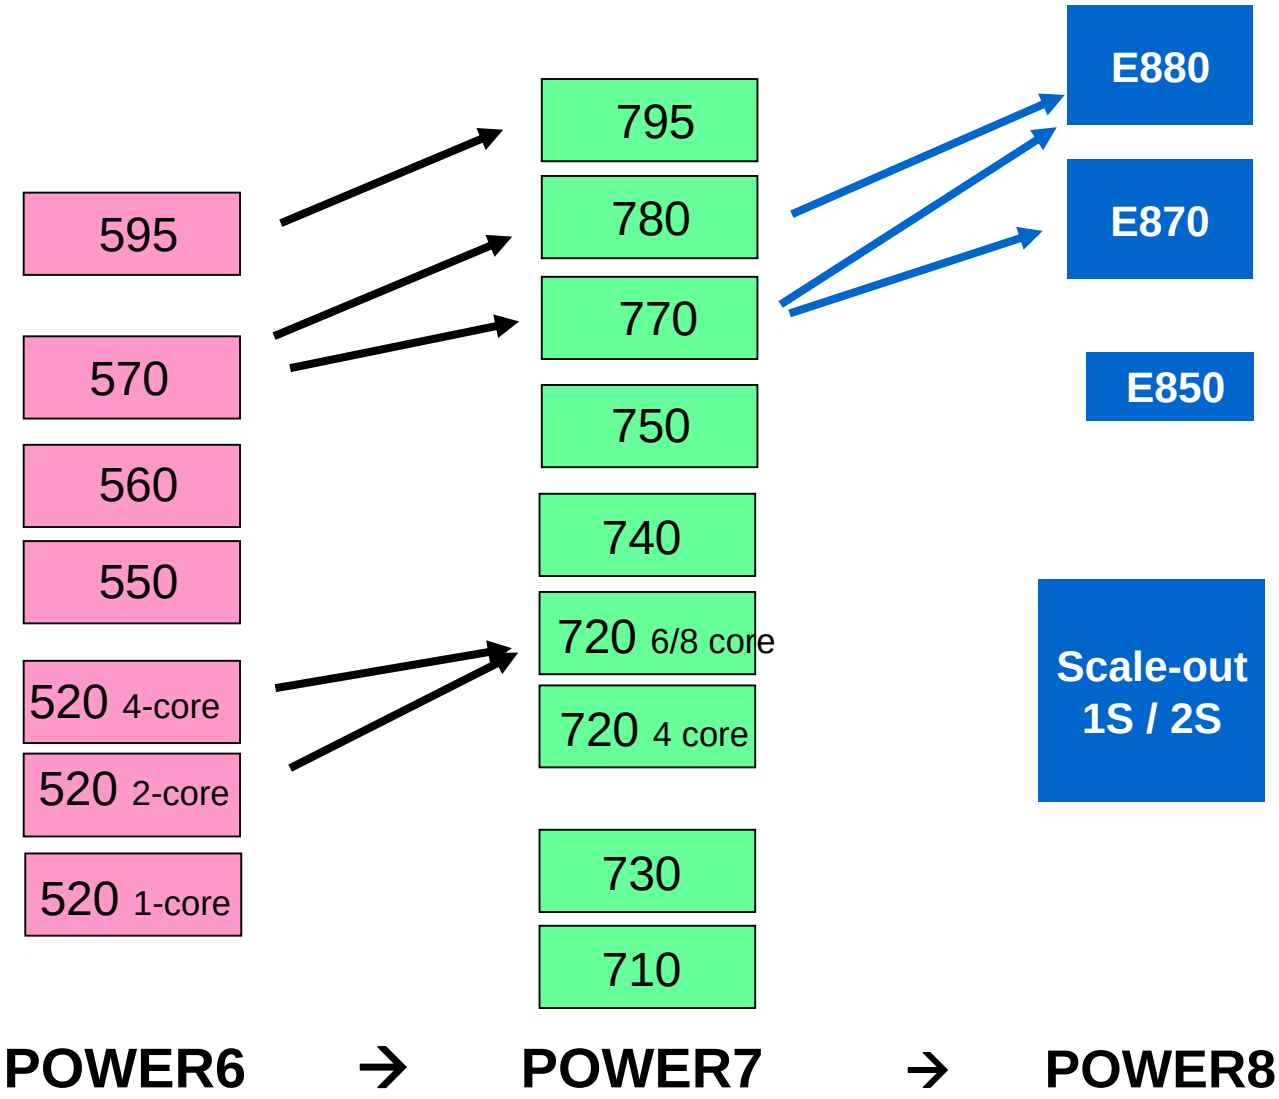

770/780 Model Upgrades (Same Serial Number)

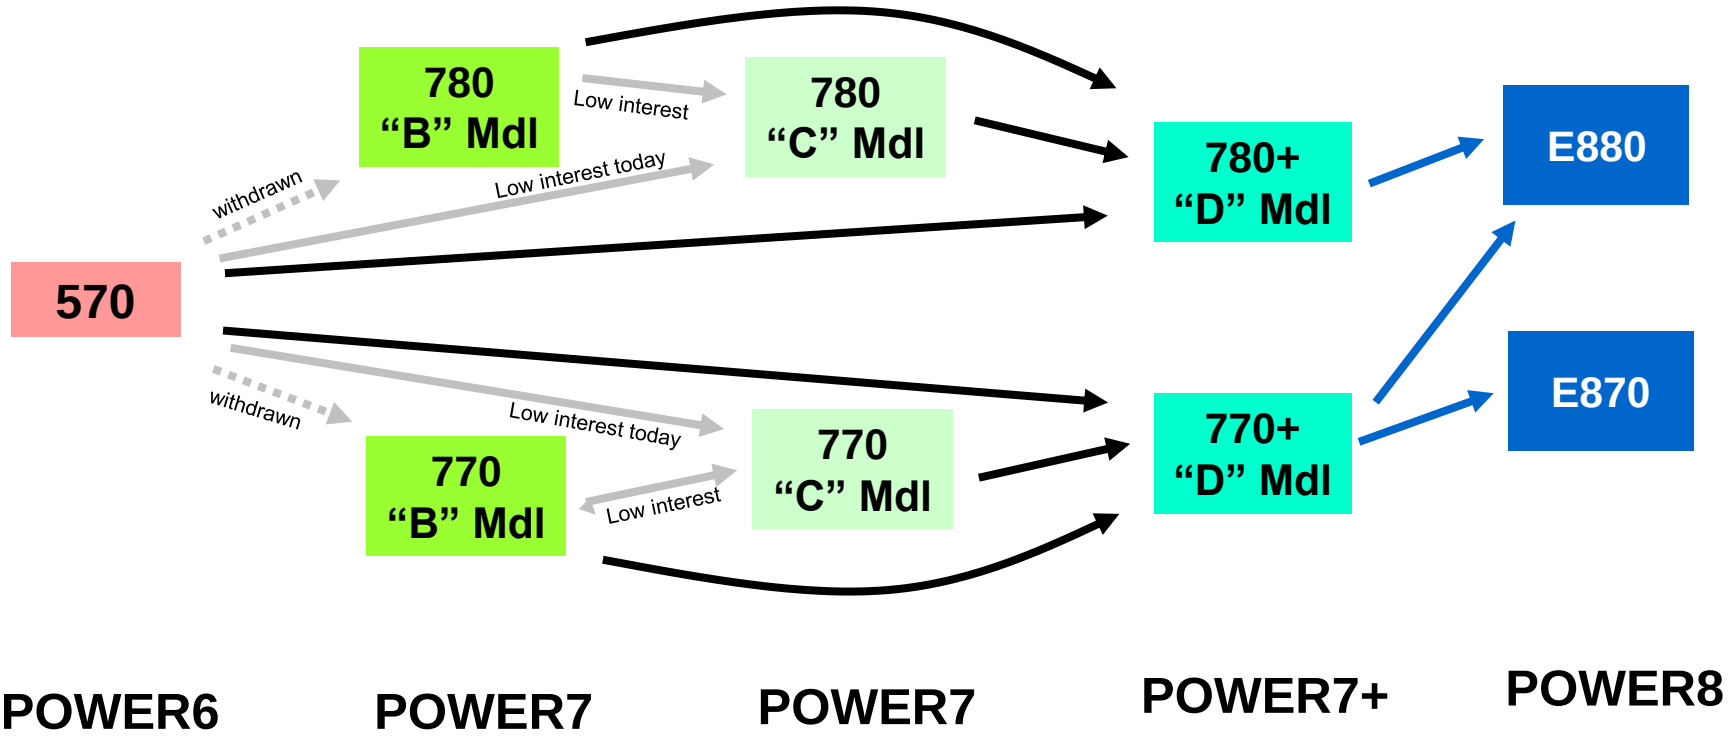

DAS or SAN ?

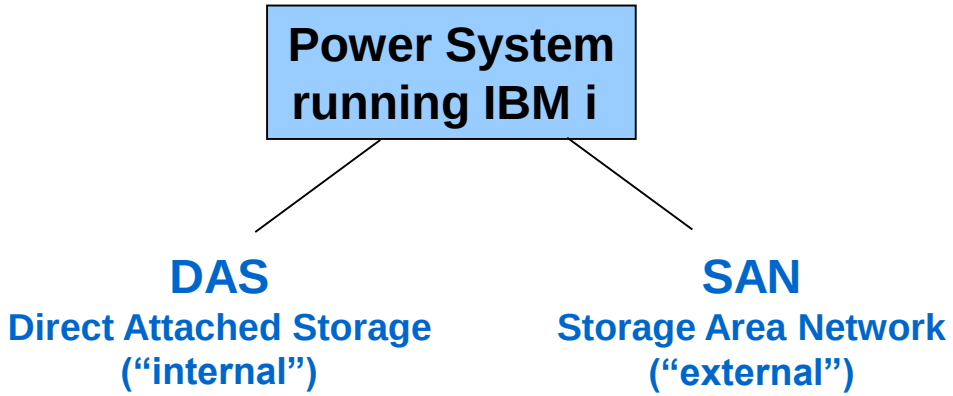

- Both options are strategic
- Both options have their strengths
- Can use both options on the same server
- Application independent
  - Ignoring operational options

DAS & SAN - Two Good Options

- **Both options are strategic**
- **Both options have their strengths**
- **Can use both options on the same server**

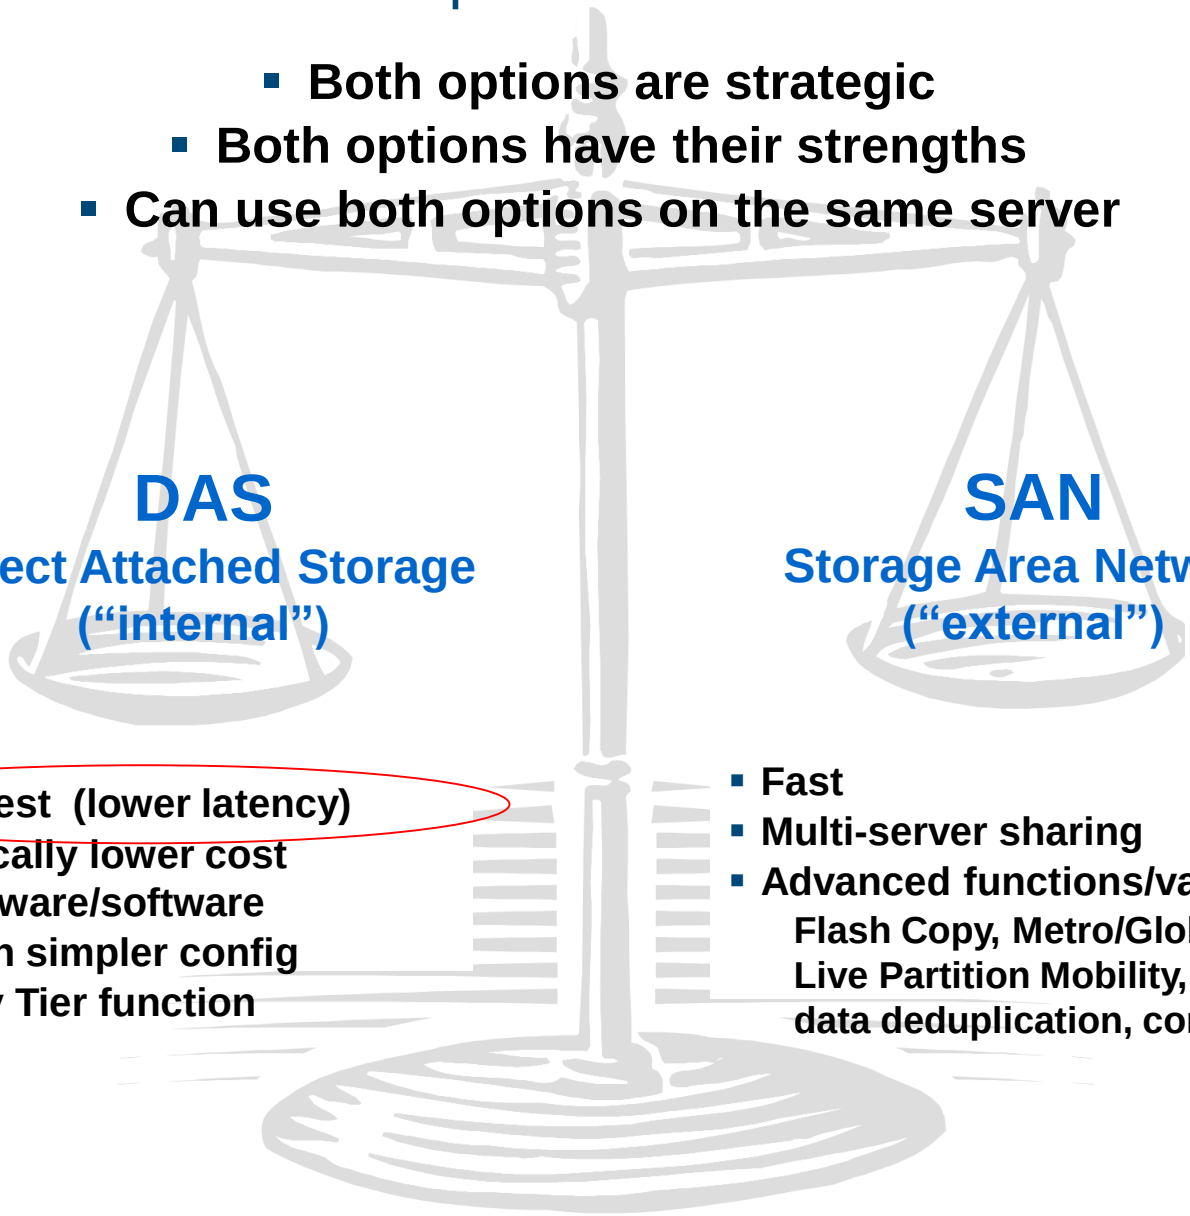

DAS
Direct Attached Storage ("internal")

- **Fastest (lower latency)**
- **Typically lower cost hardware/software**
- **Often simpler config**
- **Easy Tier function**

SAN
Storage Area Network ("external")

- **Fast**
- **Multi-server sharing**
- **Advanced functions/values**
  Flash Copy, Metro/Global Mirror, Live Partition Mobility, Easy Tier, data deduplication, compression

POWER8 SAS SFF HDD Options as of Nov 2015

SFF-3 HDD	Block size	Formatted with 512 or 4096 byte sectors *	Formatted with 528 or 4224 byte sectors
10k	5xx	300 GB *	283 GB
10k	5xx	600 GB *	571 GB
10k	5xx	1.2 TB *	1.1 TB
10k	4k	600 GB *	571 GB
10k	4k	1.2 TB *	1.1 TB
10k	4k	1.8 TB *	1.7 TB
15k	5xx	146 GB *	139 GB
15k	5xx	300 GB *	283 GB
15k	5xx	600 GB *	571 GB
15k	4k	300 GB *	283 GB
15k	4k	600 GB *	571 GB

* POWER8 SFF HDD shipped from IBM formatted in 528 or 4224 byte sectors. They can be reformatted to 512 or 4096 by the client if the extra protection is not desired.

4k drives (HDD or SSD) can NOT be reformatted to 5xx drives.

4k drives and 5xx drives can NOT be mixed on the same array. True for both HDD and SSD.

POWER8 SAS SFF HDD Options as of Nov 2015

SFF-3 HDD	Block size	Formatted with 512 or 4096 byte sectors *	Formatted with 528 or 4224 byte sectors
10k	5xx	300 GB *	
10k	5xx	600 GB *	571 GB
10k	5xx	1.2 TB *	1.1 TB
10k	4k	600 GB *	571 GB
10k	4k	1.2 TB *	
10k	4k	1.8 TB *	1.7 TB
15k	5xx	146 GB *	
15k	5xx	300 GB *	283 GB
15k	5xx	600 GB *	571 GB
15k	4k	300 GB *	
15k	4k	600 GB *	

Smaller capacity drives approaching end of marketing life

4k more strategic than 5xx

Smaller capacity drives approaching end of marketing life

4k more strategic than 5xx

* POWER8 SFF HDD shipped from IBM formatted in 528 or 4224 byte sectors. They can be reformatted to 512 or 4096 by the client if the extra protection is not desired.

4k drives (HDD or SSD) can NOT be reformatted to 5xx drives.

4k drives and 5xx drives can NOT be mixed on the same array. True for both HDD and SSD.

eMLC SSD Generations

2.5-inch (SFF)	IO OPERATIONS PER SECOND (IOPS)			Throughput (MB/s)		Latency - Response Time (ms)
SSD	Random Read	Random Write	Random Mixed Read/Write	Read	Write	Random Mixed Read/Write
Gen 1 177GB	15 k	4 k	11 k (70%-30%)	170 MB	64 MB (24-123)	0.31 ms (70%-30%)
Gen 2 387GB	39 k	22 k	24 k (70%-30%)	340 MB	375 MB	0.25 ms (70%-30%)
Gen 3 387GB 775GB	80 k	49 k	58 k (60%-40%)	340 MB	450 MB	0.15 ms (60%-40%)
For grins … 15k rpm HDD	0.12 - 0.4 k	0.12 - 0.4 k	0.12 - 0.4 k	~175 MB	~175 MB	8.3 – 2.5 ms

Note these are single drive specific measurements reflecting sustained drive workloads (not burst). The values assume 528 byte sectors running RAID-0 with no protection. Hypothetically if measured with unsupported 512 byte sectors, values would be higher. The values are highly workload dependent. Factors such as read/write mix, random/non-random data, drive cache hits/misses, data compressibility in the drive controller, large/small block, type of RAID or mirroring protection, etc will change these values. The choice of which SAS controller/adapter is running the drive can also impact these values.. These values produced by a server with plenty of processor, memory and controller resources to push this much I/O into the SSD. Most client system applications don't push SSD nearly this hard. Latency measurements using OLTP1 60/40 random 4k transfers.

SSD Driving SAS Adapter Technology

Comparing large cache SAS Adapters

980k IOPs

400k IOPs

80k IOPs

PCIe SAS – PCIe Gen 3
1Q 2014
Designed for newest SSD
800k-ish IOPS in Gen1 slot *
Nearly 1M IOPS in Gen2 slot
Up to 96 HDD
Up to 48 SSD

*** restricted by PCIe Gen1 slot in which the adapter first announced**

PCIe SAS – PCIe Gen 2
4Q 2011
1st SAS adapter designed fo SSD
300-400k IOPS
Up to 72 HDD
Up to 24 SSD

PCI-X SAS
2Q 2009
Est 70-80k IOPS
Up to 48-60 HDD ….
Also can do some SSD

SAS adapter technology – designed by IBM Power Systems with Power Systems reliability, ruggedness, integrity, performance,

FlashSystem and IBM i Support

- Native support with IBM i 7.2 TR2 or later for FlashSystem 900 or
- Also supported when “behind” an SVC (including FlashSystem V9000 … “non-native”

SR-IOV Virtualization Enhanced Oct 2015

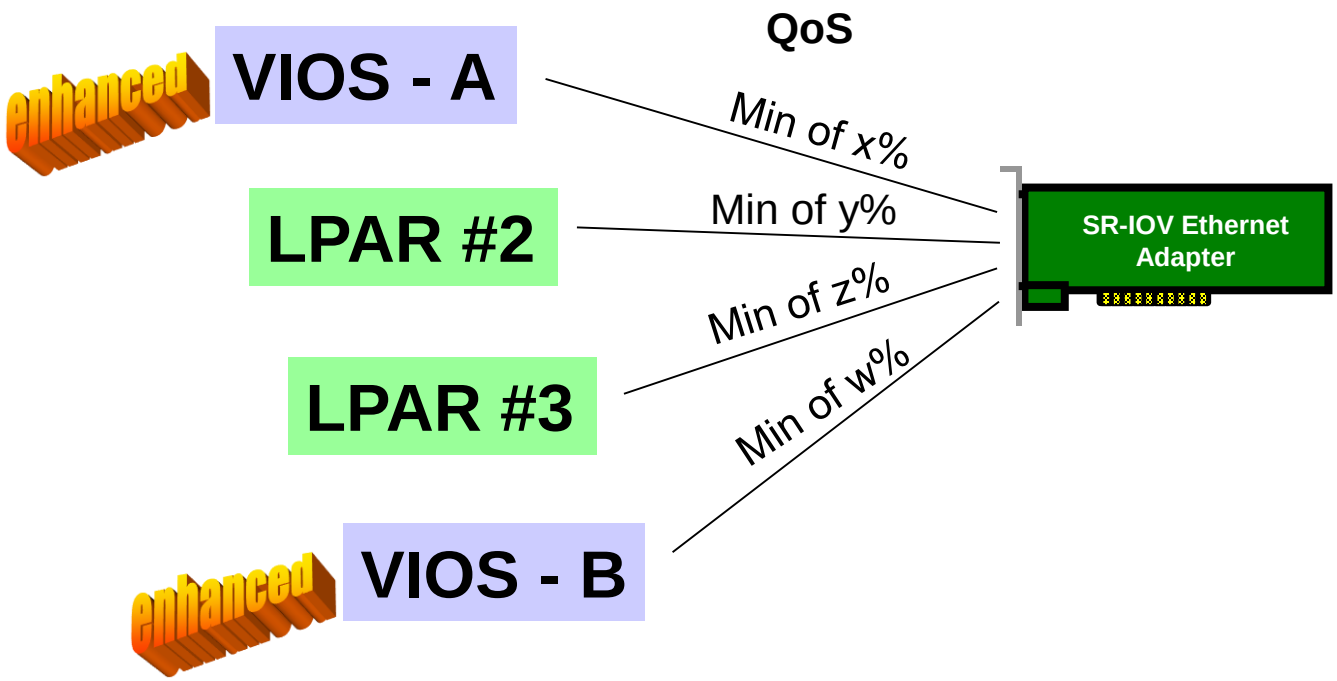

Power Hardware Update

Mark Olson

olsonm@us.ibm.com

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)