IBM Power Systems

This is Power on a Smarter Planet

Fabian Michel, CTA

Power Internal Storage & Flash Update

1Q 2014 Announce Content

Presented 27 March 2014

For COMMON Lux

The Most Powerful SAS RAID Adapters EVER

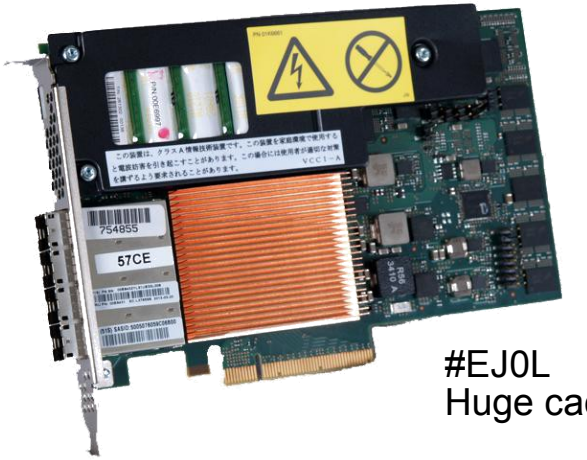

#EJ0L
Huge cache

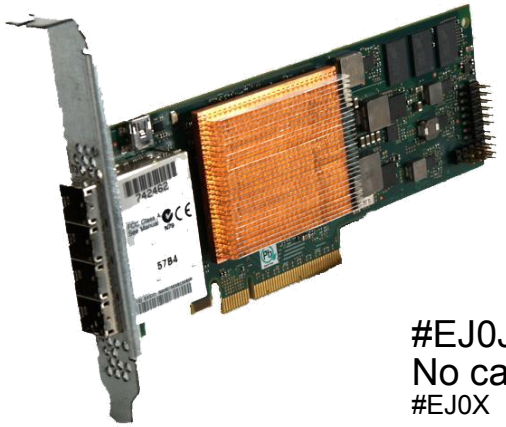

#EJ0J
No cache
#EJ0X

First PCIe Gen3 adapters introduced by Power Systems
Most powerful SAS adapters ever created by Power Systems *

*** IBM Development team not aware of a more powerful SAS adapter in the world**

- Supports 2X more SSD devices than Gen2 SAS adapter
- Up to 10X more performance than Gen1 SAS Adapter
- 2X more performance than Gen2 SAS adapter

#EJ0L effectively has 6X more write cache

Better prices than predecessor Gen2 SAS adapters

Introducing PCIe Gen3 SAS Adapter Technology

Comparing large cache SAS Adapters

PCIe SAS – PCIe Gen 3
1Q 2014
Designed for newest SSD
800k-ish IOPS *
Up to 96 HDD
Up to 48 SSD

*** restricted by PCIe Gen1 slot**

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,

Why Are These SAS Adapters a Important?

Answer varies

- For an AIX/Linux boot drive where all data is on a Fibre-Channel attached SAN …. not so critical except where 6Gb bandwidth is used

- For an AIX / IBM i / Linux server with data contained in the system unit or an I/O drawer which is a feature code of the server (#5803 or #5887 EXP24S), it can be a big deal
  - Performance:
    - SAS = lower latency than previous SAS adapters – continues to be better than SAN such as V7000 or DS8k or SVC over Fibre Channel
    - Huge write cache is further boost to lower latency
    - Higher bandwidth of PCIe Gen3 when supported in PCIe Gen2/3 slots
  - Price performance: lower price than #5913
  - Save PCI slots: support more devices per card, save slots and possibly save 12X-attached PCIe I/O drawers (#5803/5877) and expansion racks

- Future - This IBM-designed/owned technology can be leveraged in future Power servers

12 GB Write Cache Adapter #EJ0L

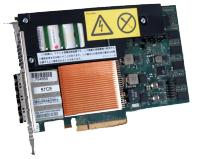

Effectively 6X more cache than PCIe2 SAS adapter

- **12 GB for PCIe3**
- **1.8 GB for PCIe2 #5913/ESA3**

- Effectively 12 GB
  - Physically 3 GB cache
    - Physical 3GB is 66% larger than the previous physical 1.8GB
  - Adapter uses compression to provide 12 GB cache
    - Data's compressibility will cause this to vary, but for typical workloads about 4X typical compression yields 12GB effectively

Adapter Write Cache Value

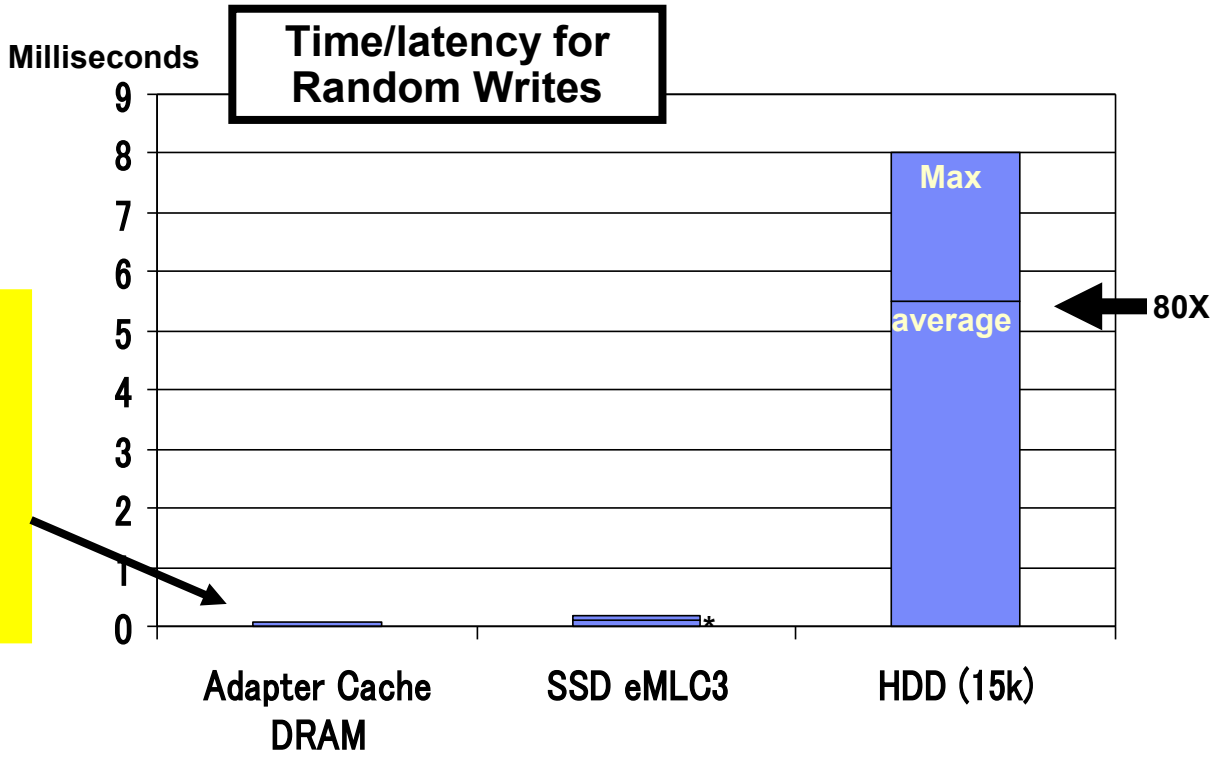

Adapter cache is

- **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.

IBM

Adapter Write Cache Value

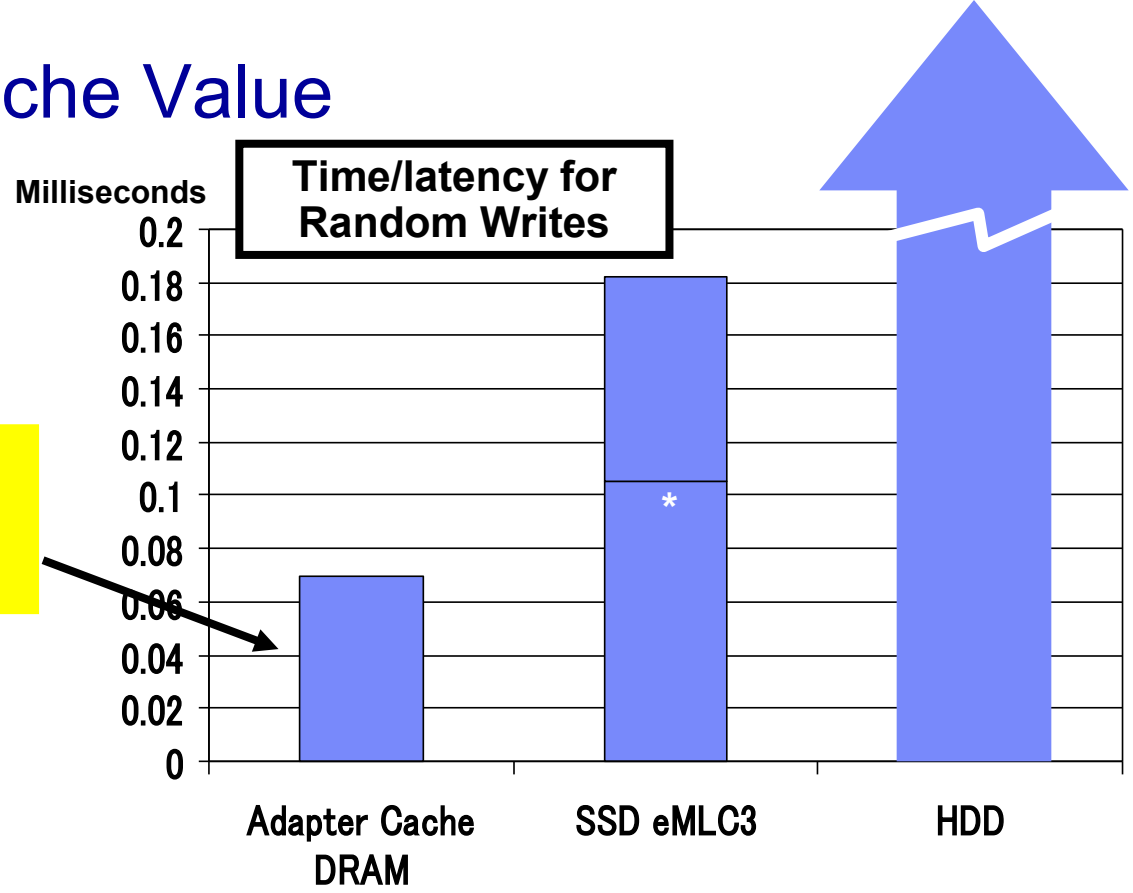

Adapter cache is

- **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.

PCIe Gen3 SAS Adapters Better Match Newer SSD

Oct 2013 announcement: Refreshed Solid State Drives (SSDs) with faster technology – 3rd Generation eMLC technology

387 GB

775 GB

Up to 2X more IOPs than previous

Up to 40% better latency than previous

Earlier SAS adapters support the newer SSD, but can not run as many SSD per adapter.

IBM USA suggested list prices as of Oct 2013 and are subject to change without notice; Reseller prices may vary.

PCIe3 SAS Adapter Support – Huge Cache #EJ0L

January 2014 announcement only for the Power 795

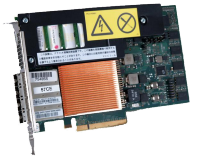

#EJ0L

Models	"B" models	"C" models	"D/T" models
PowerLinux 7R1	-	n/a *	n/a *
PowerLinux 7R2	-	No	Not Jan 2014
PowerLinux 7R4	-	-	Not Jan 2014
Power 710	n/a *	n/a *	n/a *
Power 720	No	No	Not Jan 2014
Power 730	No	No	Not Jan 2014
Power 740	No	No	Not Jan 2014
Power 750	No	-	Not Jan 2014
Power 760	-	-	Not Jan 2014
Power 770	No	Not Jan 2014	Not Jan 2014
Power 780	No	Not Jan 2014	Not Jan 2014
Power 795	Y Jan 2014	-	-
Power ESE	-	-	Not Jan 2014

* **physically impossible, no full high PCIe Slots**

IBM

PCIe3 SAS Adapter Support – Zero Cache #EJ0J

January 2014 announcement only for the Power 795

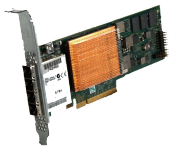

#EJ0L

Models	"B" models	"C" models	"D/T" models
PowerLinux 7R1	-	No	Not Jan 2014*
PowerLinux 7R2	-	No	Not Jan 2014
PowerLinux 7R4	-	-	Not Jan 2014
Power 710	No	No	Not Jan 2014*
Power 720	No	No	Not Jan 2014
Power 730	No	No	Not Jan 2014
Power 740	No	No	Not Jan 2014
Power 750	No	-	Not Jan 2014
Power 760	-	-	Not Jan 2014
Power 770	No	Not Jan 2014	Not Jan 2014
Power 780	No	Not Jan 2014	Not Jan 2014
Power 795	Y Jan 2014	-	-
Power ESE	-	-	Not Jan 2014

* **would need low profile feat code if announced**

IBM

PCIe Gen3 SAS OS Support

For #EJ0J, #EJ0L, & EJ0X

- OS
  - AIX V7.1 with the 7100-03 Technology Level and Service Pack 1 or later
  - AIX V6.1 with the 6100-09 Technology Level and Service Pack 1 or later

  - IBM i 6.1 with machine code 6.1.1 or later
  - IBM i 7.1 or later

  - Red Hat Enterprise Linux 6.4 for POWER, or later
  - SUSE Linux Enterprise Server 11 Service Pack 3, or later

  - VIOS 2.2.3.1 or later

#EJ0L Big cache

PCIe3 12GB Cache SAS RAID Adapter - #EJ0L

Premier SAS PCIe adapter – IBM technology designed for SSD

- Paired, full-high, single-slot adapters
- PCIe Gen3 adapter Up to 4G**B**/sec transfer
  (GB/s limited by PCIe Gen1 Slot Jan 2014)
- Four 6Gbs SAS ports
  - Each port with four x4 lanes
- Effectively 12GB write cache
- No batteries to maintain
  - Built in protection by flash memory
- Supports SAS HDD and/or SSD
  - All expected protection options
  - RAID, mirroring, etc
  - RAID array sizes from 3 - 32

Paired adapters for redundancy and performance

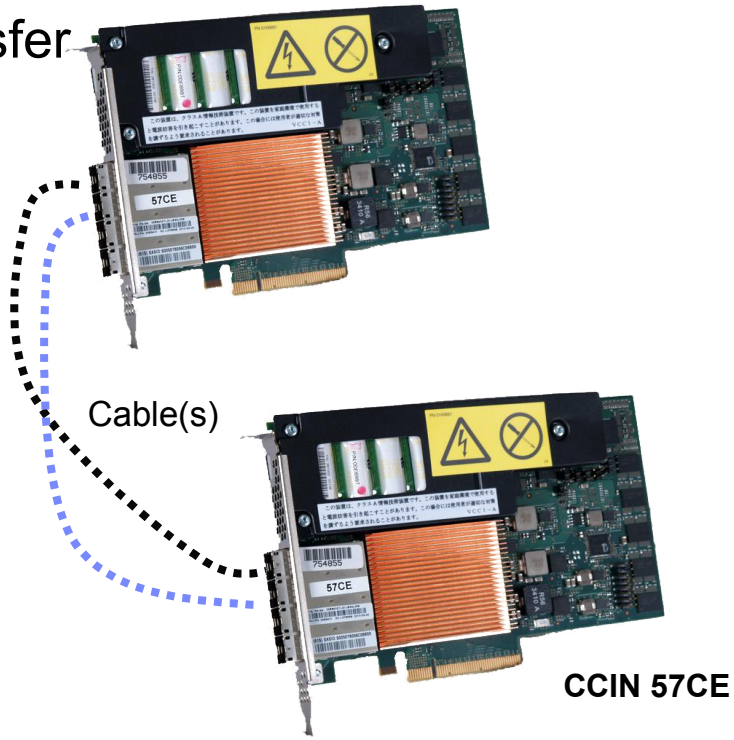

CCIN 57CE

PCIe3 vs PCIe2 Large-cache SAS PCIe Adapters

	#5913	New #EJ0L	
PCIe card technology	PCIe Gen2	PCIe Gen 3	
Write cache	1.8 GB	12 GB	6X better
Read cache	0	0	
SAS ports	3 6Gb ports	4 6Gb ports	33% more
# PCI slots per adapter	1	1	
Two cards required -- paired	Yes	Yes	
Max HDD per pair	72	98	33% more
Max SSD per pair	24	48	100% more
Mix SSD & HDD	Y	Y	
Performance	base	~ 2X	~2X better
USA List price on Power 795	$ 9,824 each $19,648 / pair	$ 9,169 each $18,338 pair	~7% better PLUS slot savings

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

#EJ0J ø cache

PCIe3 SAS RAID Adapter - #EJ0J (zero cache)

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

- Same chip technology as #EJ0L, but zero cache
  - No cache = physically smaller & lower cost & optional pairing

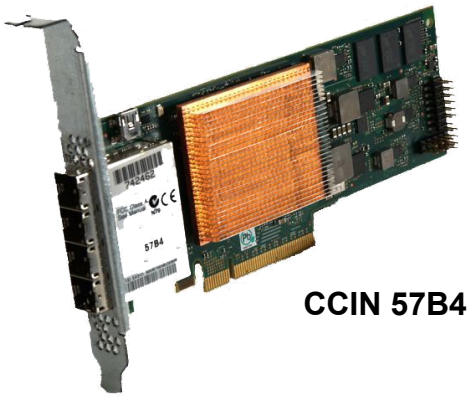

CCIN 57B4

- PCIe Gen3 adapter
- Up to 4GB/sec transfer (GB/s limited by PCIe Gen1 Slot Jan 2014)
- Four 6Gb SAS ports
  - Each port with four x4 channels

- Supports SAS HDD and/or SSD
  - All protection options – including RAID-5 and RAID-6
  - RAID, mirroring, etc
  - RAID array sizes from 3 - 32

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

PCIe3 vs PCIe2 Zero-cache SAS Adapters

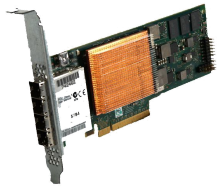

	#ESA1/ESA2	New #EJ0J
PCIe card technology	PCIe Gen2	PCIe Gen 3
Write cache	0	0
Read cache	0	0
SAS ports	2 6Gb ports	4 6Gb ports
# PCI slots per adapter	1	1
Two cards required -- paired	Optional	Optional
Max HDD per pair	0 (SSD only)	96 *
Max SSD per pair	24	48
Mix SSD & HDD	n/a	No
Performance - IOPS	base	~2X base
USA List price on Power 795	$ 8,121	$2,332

100% more

LOTS% more

100% more

2X better

only 30%

* use caution if write workload is significant, as no write cache

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

PCIe3 Zero-cache vs PCIe1 Small-Cache Adapter

* #EJ0J for modest write workload – has no write cache

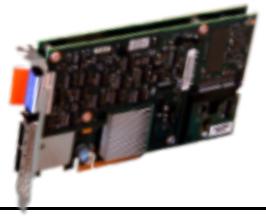

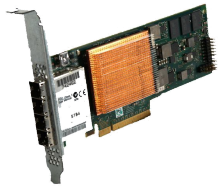

	#5805	New #EJ0J
PCIe card technology	PCIe Gen1	PCIe Gen 3
Write cache	380 MB	0
Read cache	0	0
SAS ports	2 3Gb ports	4 6Gb ports
# PCI slots per adapter	1	1
Two cards required -- paired	Required	Optional
Max HDD per pair	48	96 *
Max SSD per pair	9	48
Mix SSD & HDD	N	N
Performance - IOPS	base	**
USA List price on Power 795	$ 2,880	$2,332

"200"% more
2x ports
2x bandwidth

save slot

LOTS* more HDD or SSD

only 30%

* use caution if write workload is significant, as no write cache
** hard to compare because of write cache impact, especially to HDD

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

PCIe3 vs PCIe2 vs PCIe1 Zero-cache SAS Adapters



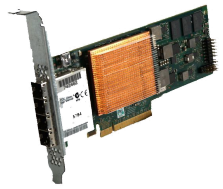

	#5901/5278	#ESA1/ESA2	New #EJ0J
PCIe card technology	PCIe Gen1	PCIe Gen2	PCIe Gen 3
Write cache	0	0	0
Read cache	0	0	0
SAS ports	2 3Gb	2 6Gb ports	4 6Gb ports
# PCI slots per adapter	1	1	1
Two cards required -- paired	Optional	Optional	Optional
Max HDD	48	0 (SSD only)	96 *
Max SSD	0 (HDD only)	24	48
Mix SSD & HDD	n/a	n/a	N
Performance - IOPS	Est 1/4th base	base compare	~2X PCIe2
RAID-5/6 supported	No	Yes**	Yes**
USA List price on Power 795	$991	$ 8,121	$2,332

* use caution if write workload is significant, as no write cache
** function supported, but without write cache use judiciously

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

Excellent SAS Reference (being updated)

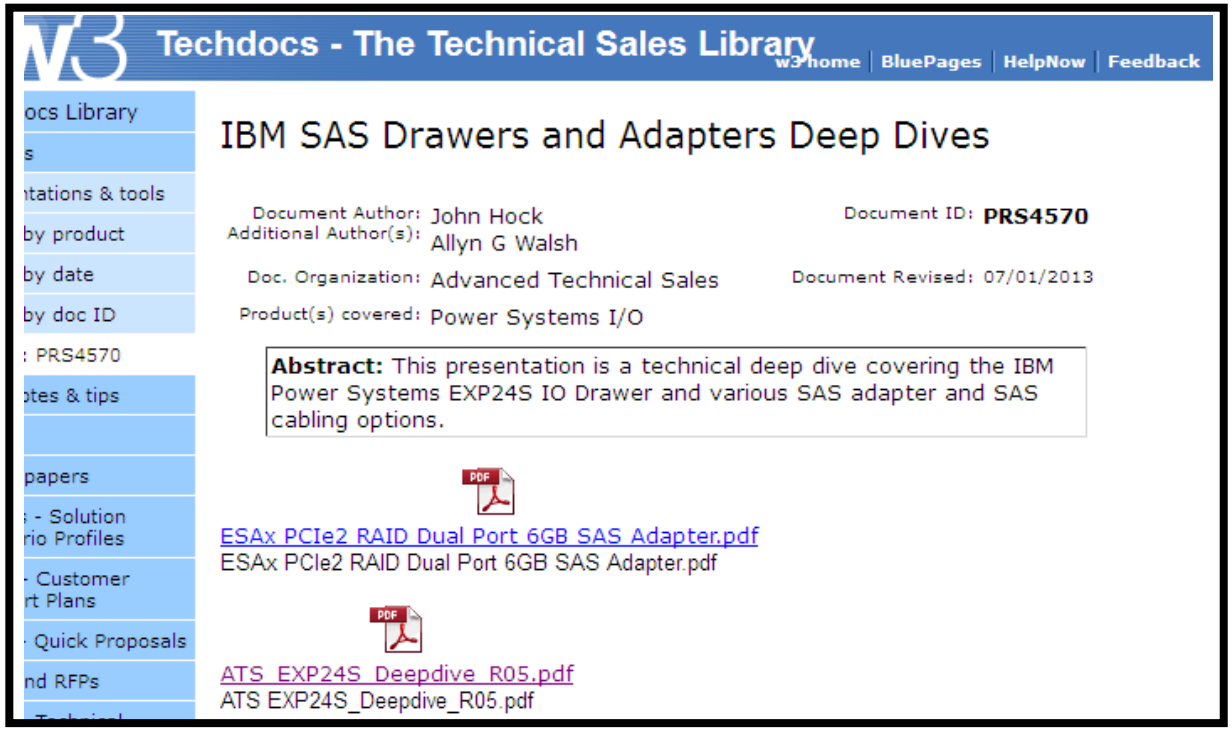

Techdoc PRS4570

- TechDoc contains excellent, detailed material
  - http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4570 (IBMers)
  - http://partners.boulder.ibm.com/src/atsmastr.nsf/WebIndex/PRS4570 (Business partners)
  - http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4570 (Clients)
- However, material not yet updated with #EJ0J and #EJ0L material. Estimating update by mid March 2014. Use announcement letter and charts in this deck including detailed charts at end of presentation until then.

- See also IBM InfoCenter for planning information

#EJ0X Tape

PCIe3 SAS Tape Adapter - #EJ0X

LTO-5/LTO-6 SAS tape adapter

- Supports full bandwidth of LTO-5/LTO-6

- PCIe Gen3 adapter
- Up to 4GB/sec transfer (GB/s limited by PCIe Gen1 Slot Jan 2014)
- 4 SAS ports
  - 6Gbps ports

- DVD and other SAS tape drives not supported

- Physically is same card as #EJ0J, but has different feat code to help IBM config tools like eConfig understand card's use and appropriate cabling.
  eConfig does not know how to swap usage to disk/SSD. Use no-charge RPQ to adjust IBM feature records for any "re-purposing".

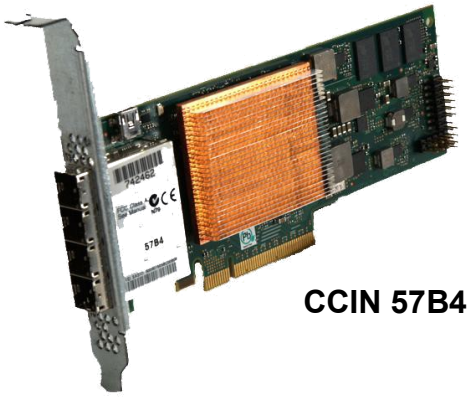

CCIN 57B4

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

Attaching both disk & tape to the same adapter NOT supported

PCIe3 SAS Tape Adapter Support

January 2014 announcement only for the Power 795

Models	"B" models	"C" models	"D/T" models
PowerLinux 7R1	-	No	Not Jan 2014*
PowerLinux 7R2	-	No	Not Jan 2014
PowerLinux 7R4	-	-	Not Jan 2014
Power 710	No	No	Not Jan 2014*
Power 720	No	No	Not Jan 2014
Power 730	No	No	Not Jan 2014
Power 740	No	No	Not Jan 2014
Power 750	No	-	Not Jan 2014
Power 760	-	-	Not Jan 2014
Power 770	No	Not Jan 2014	Not Jan 2014
Power 780	No	Not Jan 2014	Not Jan 2014
Power 795	Y Jan 2014	-	-
Power ESE	-	-	Not Jan 2014

* would need low profile feat code if announced

PCIe3 vs PCIe1 SAS Tape Adapters

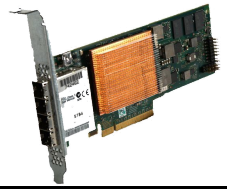

	#5901/5278	New #EJ0X
PCIe card technology	PCIe Gen1	PCIe Gen 3
SAS ports – quantity per card	2	4
SAS ports – Gb/sec	3Gb	6Gb
Max tape drives per card	2	8
SAS/SATA DVD supported	Y	N
LTO-2, LTO-3 (SCSI vs SAS)	N	N
LTO-4	Y	N
LTO-5, LTO-6	Y, but max bandwidth limited	Y, with full bandwidth
DAT160	Y	N
Attach disk and tape on the same card at the same time	Not supported	Not supported
Models supported on	all	Just 795 today
Feat code if re-purpose to disk	#5901/5278	#EJ0J
USA List price on Power 795	$991	$2,332

2X more
2X more
4X more

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

Refreshed PCIe2 large cache SAS adapter

REVIEW
Oct 2013 Announcement

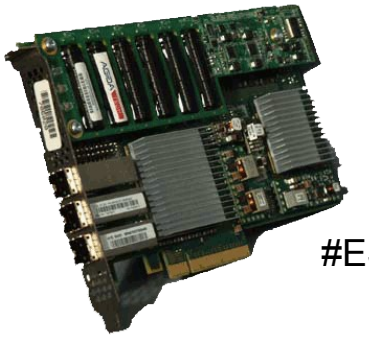

#ESA3

Re-engineered PCIe2 1.8GB RAID SAS Adapter → feature # ESA3

~20% more energy efficient – lower power draw, less cooling

Same performance as #5913 PCIe2 1.8GB RAID SAS Adapter

$600/$786 lower price per pair adapters than existing #5913 Adapter

Supported on POWER7+ 720, 740, 770, 780. SOD for more models

$600 for Power 720/740 IBM USA suggested list prices as of Oct 2013. $786 for Power 770/780 IBM USA suggested list price as of Oct 2013. Subject to change without notice; Reseller prices may vary.

Expanded PCIe2 Large Cache SAS adapter Usage

Jan 2014 Announcement

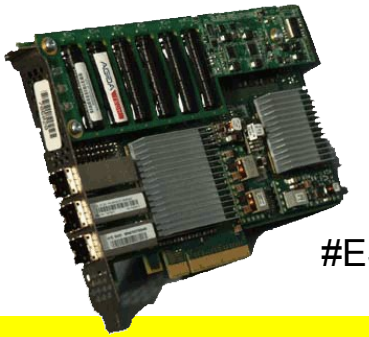

#ESA3

Adding support for the:

- Power 720, 740, 770, 780 "C" models (8202-E4C, 8205-E6C, 9117-MMC, 9179-MHC)
- Power 750/760 "D" models (8408-E8D, 9109-RMD)
- Power 730 "C and D" model (8231-E2C, 8231-E2D) in a 12X-attached PCIe I/O drawer
- PowerLinux 7R2 (8246-L2S, 8246-L2T) in a 12X-attached PCIe I/O drawer
- PowerLinux 7R4 (8248-L4T)
- Power ESE (8412-EAD)

In addition to existing support on the Power 720, 740 770, 780 “D” models

Re-engineered #5913 PCIe2 1.8GB RAID SAS Adapter → → feature # ESA3
~20% more energy efficient – lower power draw, less cooling
Same performance as #5913 PCIe2 1.8GB RAID SAS Adapter
$600/$786 lower price per pair adapters than existing #5913 Adapter

$600 for Power 720/740 IBM USA suggested list prices as of Oct 2013. $786 for Power 770/780 IBM USA suggested list price as of Oct 2013. Subject to change without notice; Reseller prices may vary.

Expanded PCIe2 Large Cache SAS adapter Usage

#ESA3

Added Jan 2014

Models	"B" models	"C" models	"D" models
PowerLinux 7R1	-	n/a **	n/a **
PowerLinux 7R2	-	Y* Jan 2014	Y* Jan 2014
PowerLinux 7R4	-	-	Y Jan 2014
Power 710	n/a **	n/a **	n/a **
Power 720	Not Jan 2014	Y Jan 2014	Y (2013)
Power 730	No	Y* Jan 2014	Y* Jan 2014
Power 740	Not Jan 2014	Y Jan 2014	Y (2013)
Power 750	Not Jan 2014	-	Y Jan 2014
Power 760	-	-	Y Jan 2014
Power 770	Not Jan 2014	Y Jan 2014	Y (2013)
Power 780	Not Jan 2014	Y Jan 2014	Y (2013)
Power 795	No	-	-
Power ESE	-	-	Y Jan 2014

*** requires full high PCI slot in 12X-attached I/O drawer**
**** no full high PCIe slots available**

When Use #EJ0L vs #ESA3 vs #5913

#1

1. First option -- Use #EJ0L when you can.
   (better performance, pricing, function)
2. If first option not available -- Use #ESA3 where you can
   (energy savings and price over #5913)
3. Otherwise use #5913

Today there is little overlap.

- #EJ0L is only on Power 795
- #ESAS is only on POWER7+ D models POWER7 C models

In the future there will more overlap. Use the selection criteria above when there is overlap

Easy Tier Server Adds Flash Adapter 90 Support

POWER7/7+ Server w/ AIX

DS8870 Storage

News:
Support added Dec 2013 by DS8k Easy Tier Server for Flash Adapter 90

- Accelerates applications by caching frequently accessed data from IBM DS8870 SAN to Flash Adapter 90 or SSD on server
- Customers with I/O intensive applications can recognize:
  - Up to 5x improvement in application performance
  - Improved end user experience, numbers of tractions, faster data insights, increased business opportunities and revenue
  - Save on SW licensing, floor space, and energy costs with improved processor and SAN utilization
- Supports DS8870 System Storage, POWER7/POWER7+ servers with AIX, and EXP30 and EXP24S SSD drawers and IBM Flash Adapter 90

IBM Flash Adapter 90 & Easy Tier Server Support

Classic benefits of flash memory technology

- Deliver extremely fast application and database performance
- Gain better and faster business insights
- Increase revenue with faster transactions and better end user experience
- Improve processor utilization and application efficiency. Do more with less capital expense, rack space, and licensing costs

In a PCIe Card that

- ✓ Does not require rack space for an I/O drawer
- ✓ Low entry price point
- ✓ Delivers low latency, high IOPS and high bandwidth

Easy Tier Server now supports Flash Adapter 90

Significantly accelerate applications by caching frequently accessed data from IBM DS8870 SAN to Flash Adapter 90

Capacity	900GB eMLC
Size	Full Height, Half Length
IOPS	325K
BW	700 MB/s
Latency	**Read: 96 µs, Write: 37 µs**

IBM Flash Adapter 90 Supported Configurations

#ES09 IBM Flash Adapter 90 (PCIe2 0.9TB) supported on

- POWER7+ 720 & 740 – in system unit or #5802/5877 I/O drawer
- POWER7+ 730 – in #5802/5877 I/O drawer
- PowerLinux POWER+ 7R2 – in #5802/5877 I/O drawer
- Planned support of additional POWER7+ servers in 2014

Supported Operating Systems

- AIX V7.1 TL3 or later
- AIX V7.1 TL2 and Service Pack 4 or later
- AIX V7.1 TL1 and Service Pack 9 or later
- AIX V6.1 TL 9 or later
- Red Hat Enterprise Linux 6.5 or later for IBM POWER

Dec 2013 annc

Other config rules

- Up to 4 adapters per LPAR or server
- Maximum 4 adapters per System unit
- Maximum 4 adapters per #5802/5877 I/O drawer

#ES09
Full high PCIe adapter

Here a use case as described in the mentioned White Paper:

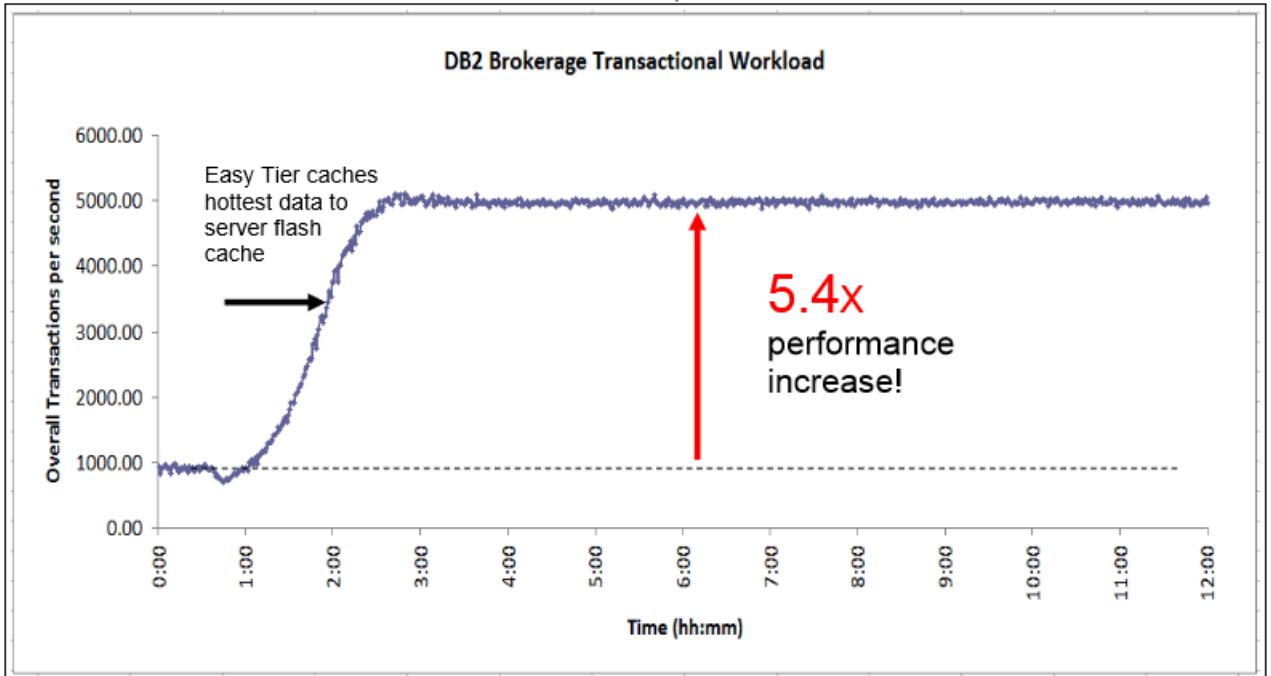

Figure 2 Overall DB2 Brokerage Transactions per second in DS8870 with Easy Tier Server

Easy Tier Server is an extension to the internal Easy Tier of the DS8000, that goes up to the server. The Power server uses **local flash caching to accelerate its I/Os** (read cache), and all what is cached locally with the very short latencies at the server, doesn't need to go through the SAN. Yet what is being cached is to some degree determined by the Easy Tier logic within the DS8870, which is in steady "contact" with the local server flash cache. To some degree, the local server caching can act autonomously and on a short-term timescale, and with small granularities, to optimise itself also.
This is one of the prominent **System p synergy** items that we have for the IBM DS8000 product line. Before the Flash Adapter 90 was available, only the bigger EXP30 Ultra SSD drawer could be used for local flash caching in Power servers.
http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102300
http://www.redbooks.ibm.com/abstracts/redp5013.html?Open

How did we arrive here with flash?

- **The revolution started** at the **Flash Ahead Initiative** in NYC, April ’13
  - IBM launched the FlashSystem product line
  - IBM has a $1B commitment to flash support and development
- We’ve arrived here by way of **three technology transformations and one epic problem**
  - Processor performance increases have followed Moore’s law
  - Unprecedented growth in concurrent users and transactions driven by the proliferation of mobile and internet attached devices as well as faster networks
  - Rapid adoption of flash in consumer electronics drove up flash density and drove down flash price thus enabling flash in the enterprise
  - Disk drives increased capacity at the expense of performance/capacity
- This demands an evolutionary solution, the next generation **IBM FlashSystem**

IT infrastructure challenges

CPU performance **up 10x** this last decade

Storage has grown capacity but unable to keep up in performance

Systems are now Latency & IO bound resulting in significant performance gap

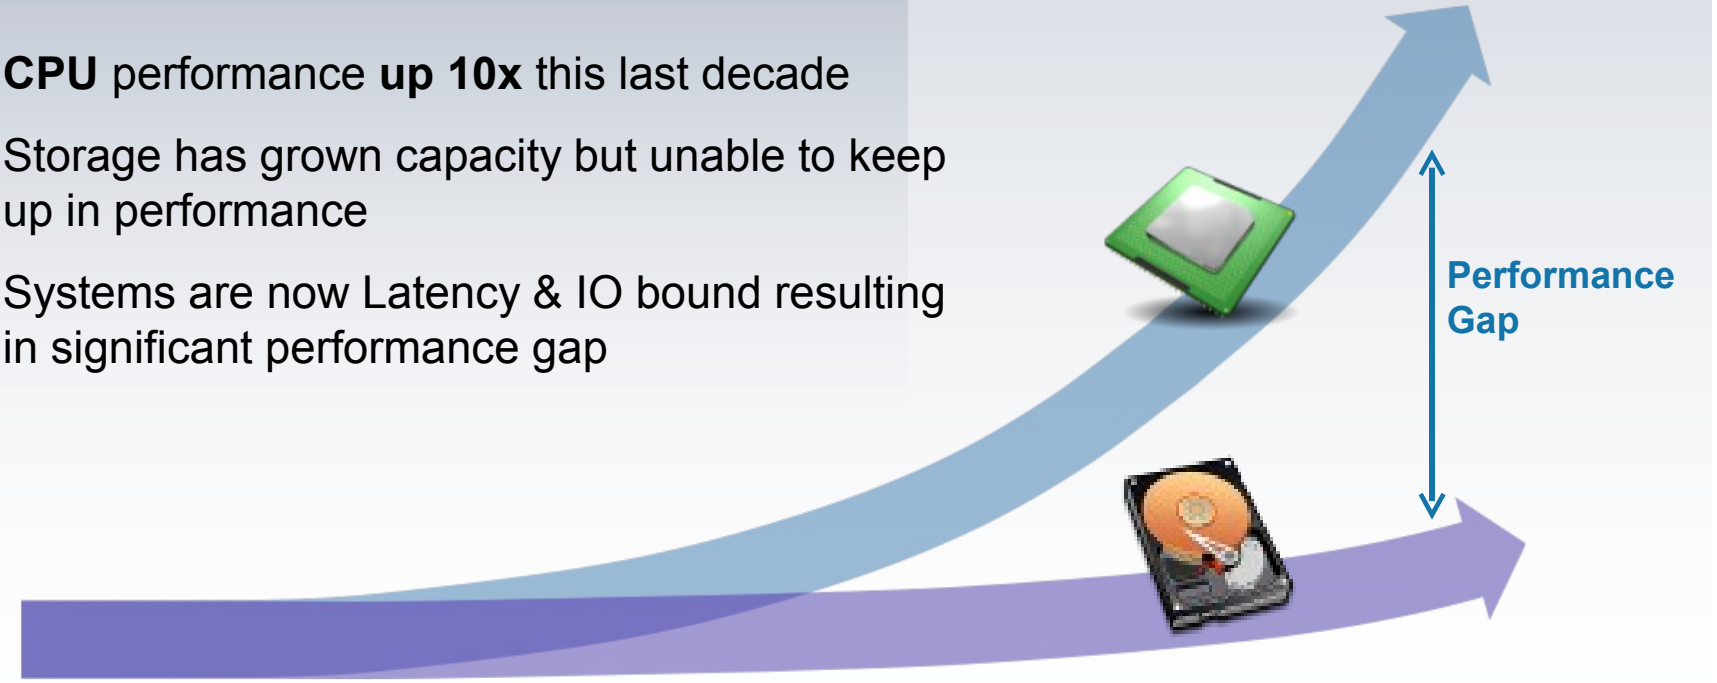

From 1980 to 2010, **CPU** performance has **grown 60% per year***

…and yet, **disk** performance has **grown ~5% per year** during that same period**

Client responses to performance gap

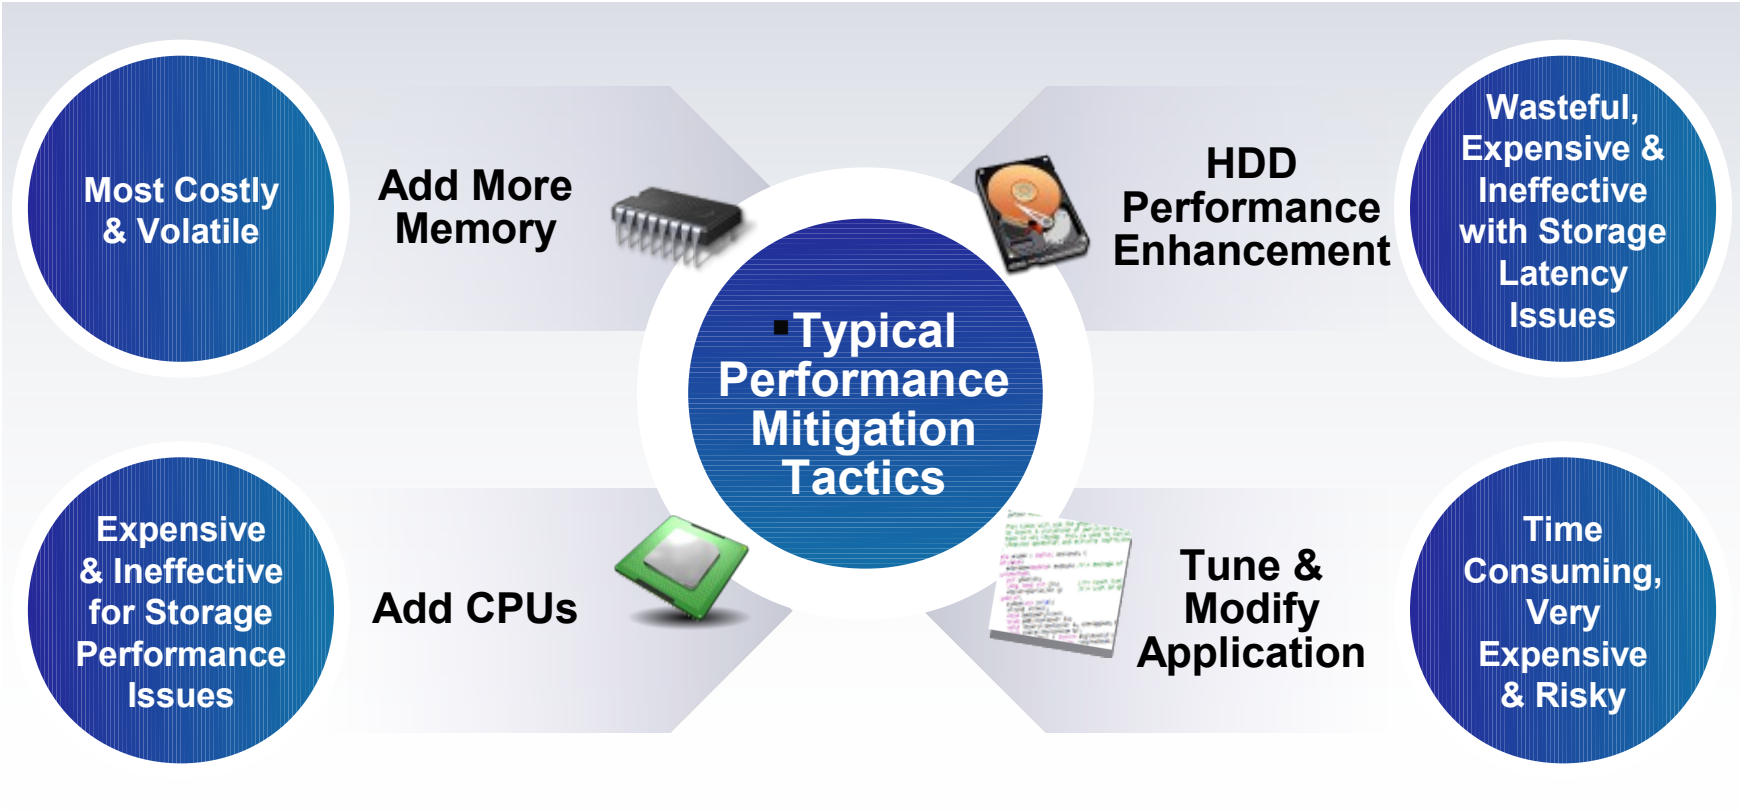

Our unlikely hero makes his entrance

Flash Memory

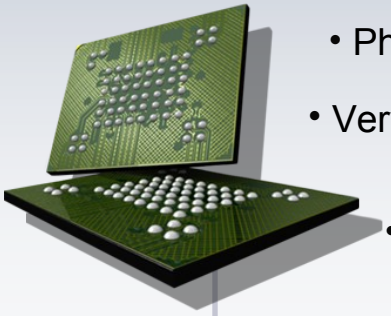

- Physically small and design flexible
- Very fast (low latency)
- Non volatile
- Very low power
- Inherently rugged

Like many heroes this one can be imperfect; a single flash chip…

- Cannot hold enough data
- Writes & erases slowly
- Is not reliable enough

IBM uses tools and technologies to drive a robust solution…

- Array the flash chips
- IBM MicroLatency™
- Variable Stripe RAID™
- System-level RAID/Two Dimensional RAID
- Redundant Components
- Serviceable Components
- Concurrent Code Load

Introducing the solution, the FlashSystem 840 replacing performance HDD in your data center

Data center optimized to deliver extreme performance, flexible capacity and total system protection

Performance at-a-glance

Minimum latency	
Write	90 µs
Read	135 µs
Maximum IOPS 4 KB	
Read (100%, random)	1,100,000
Read/write (70%/30%, random)	775,000
Write (100%, random)	600,000
Maximum bandwidth 256 KB	
Read (100%, sequential)	8 GB/s
Write (100%, sequential)	4 GB/s

Capacity Options

Flash module configuration	2 x 2 TB	4 x 2 TB	4 x 4 TB	8 x 2 TB	12 x 2 TB	8 x 4 TB	12 x 4 TB
Raw capacity	5 TB	11 TB	21 TB	21 TB	32 TB	42 TB	65 TB
RAID 0 usable capacity	4 TB	8 TB	N/A	16 TB	24 TB	32 TB	48 TB
RAID 5 usable capacity	N/A	4TB	8 TB	12 TB	20 TB	24 TB	40 TB

IBM FlashSystem 840

Macro Efficiency	• **2U** form factor- minimal footprint for best of breed ROI • **Low power** 625 watts • **Field upgradeable, granular capacity** - 4, 8, 12, 16, 20, 24, 32, 40, 48 • Reduce installation and management time with **intuitive standardized GUI**
Enterprise Reliability	• **Fully redundant** and **hot swappable** architecture: - Flash modules, power supplies, batteries, interfaces, fans • Maintain **business continuity** with non-disruptive maintenance and updates - **Concurrent code load**, highly serviceable design
MicroLatency™	• **Low Latency** 135/90 µs R/W • **Purpose-built**, **highly parallel** design • Maximize host **CPU efficiency** and **productivity**
Extreme Performance	• **1.1M IOPS** • **8 GB/s** Bandwidth • Multiple connectivity interfaces - 16Gb/8Gb **Fibre** Channel - 40Gb **QDR InfiniBand** - 10Gb **FCoE**

Data center optimized to deliver extreme performance, flexible capacity and total system protection

Boost IT efficiency

Macro Efficiency

driven by consolidation of hardware and software, deployment speed, efficient use of IT staff as well as power and cooling savings

98% Reduction

in Application Processing Time

50% Reduction

in TCO

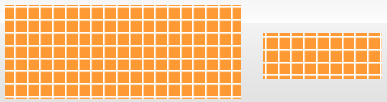

97% Reduction

in Physical Footprint

95% Reduction

in Power Consumption

A top price comparison retail site saved $1.84K in three ways.

$75K in floor space.

$1.8K in power.

$91K in cooling.

Source: Price comparison retail site

Source: IBM Client Experiences

The data below are based on average operating conditions that may or may not be representative of a particular customer's operating environment. The use case measurements are from TMS customers using the flash technology that has been integrated into IBM's systems.

Improve business uptime

Enterprise Reliability

durable and reliable designs that use enterprise class flash and patented data protection technology

- Superior protection with multiple RAID layers
- Advanced wear leveling and over-provisioned space
- Non disruptive maintenance and current code load

"IBM FlashSystem ticks all the boxes for us."

Reliability. ✓

Ease of setup. ✓

Source:
Technical Analyst, Rathbone Brothers Plc. case study

Superior Protection: *Beyond Disk RAID*

Variable Stripe Sizes
Read Disturb Mitigation
Automatic Read Sweeper
High-Speed Clock Recovery

Advanced Engineering = Less Maintenance

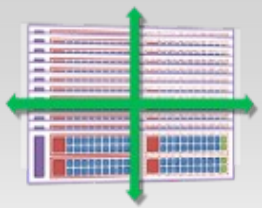

Protection Within ***And*** Across Flash Modules

Self-Recovering Flash Modules
Avoid system rebuilds

The IBM MicroLatency advantage

MicroLatency
deliver a microsecond response time to accelerate critical applications to achieve competitive advantages

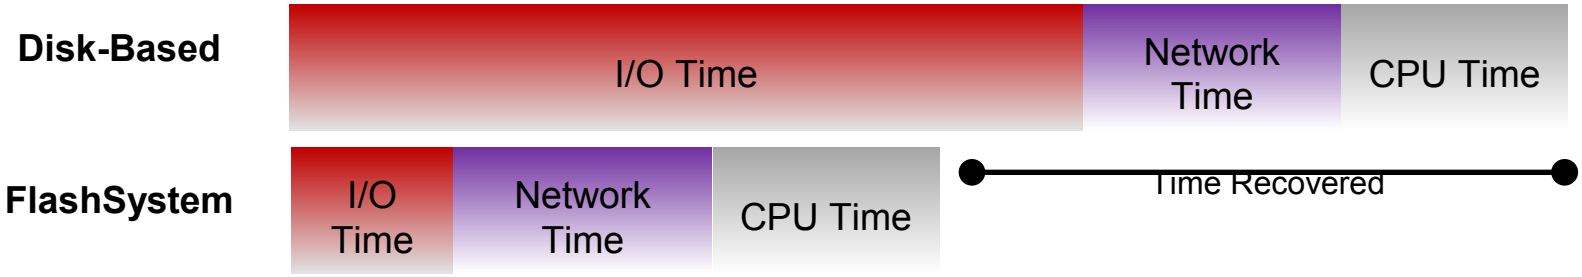

- Faster decision making
- Increase revenue
- Accelerate cost savings
- Eliminate wait time
- Scale performance with capacity

"You could crush scheduled and non-scheduled I/O requests."

Source: IT Mgr, SciQuest

1 microsecond : 1 second :: 1 second : 11.5 days

The value of performance

Extreme Performance
enable business to unleash the power of performance, scale, and insight to drive services and products to market faster

- Improved end-user experience
- Faster insights into critical applications

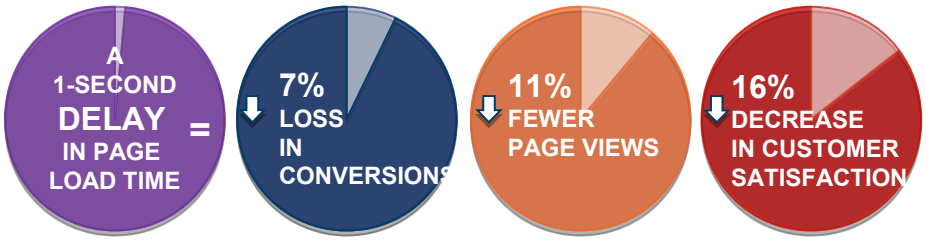

In dollar terms, this means that if your site typically earns $100,000 a day, this year you could lose **$2.5 million** in sales.

Source: Aberdeen Group

CCBCC cut data processing time by 75% without replacing a single server.

Source: Coca-Cola Bottling Co. Consolidated case study

Comprehensive flash offerings strategy

IBM FlashSystem addresses the two primary segmentations of the flash market as defined by IDC taxonomy:

Absolute Performance for extreme performance of throughput, IO, and latency.
Enterprise for full featured enterprise features sets married to flash performance

IBM FlashSystem

- Standalone all flash array
  - *Also a building block*
- Eschew features that increase latency, bottleneck IOPS or cap bandwidth
- Promote use of application-level services

IBM FlashSystem Solution

- Software defined storage
  - *FlashSystem Enterprise Performance Solution*
  - *Storwize V7000 FlashSystem Edition*
- FlashSystem as a storage enclosure
- Promote use of SAN services

FlashSystem Solution is planned Feb
Subject to change until announced

Help clients understand how to deploy: Cross IBM

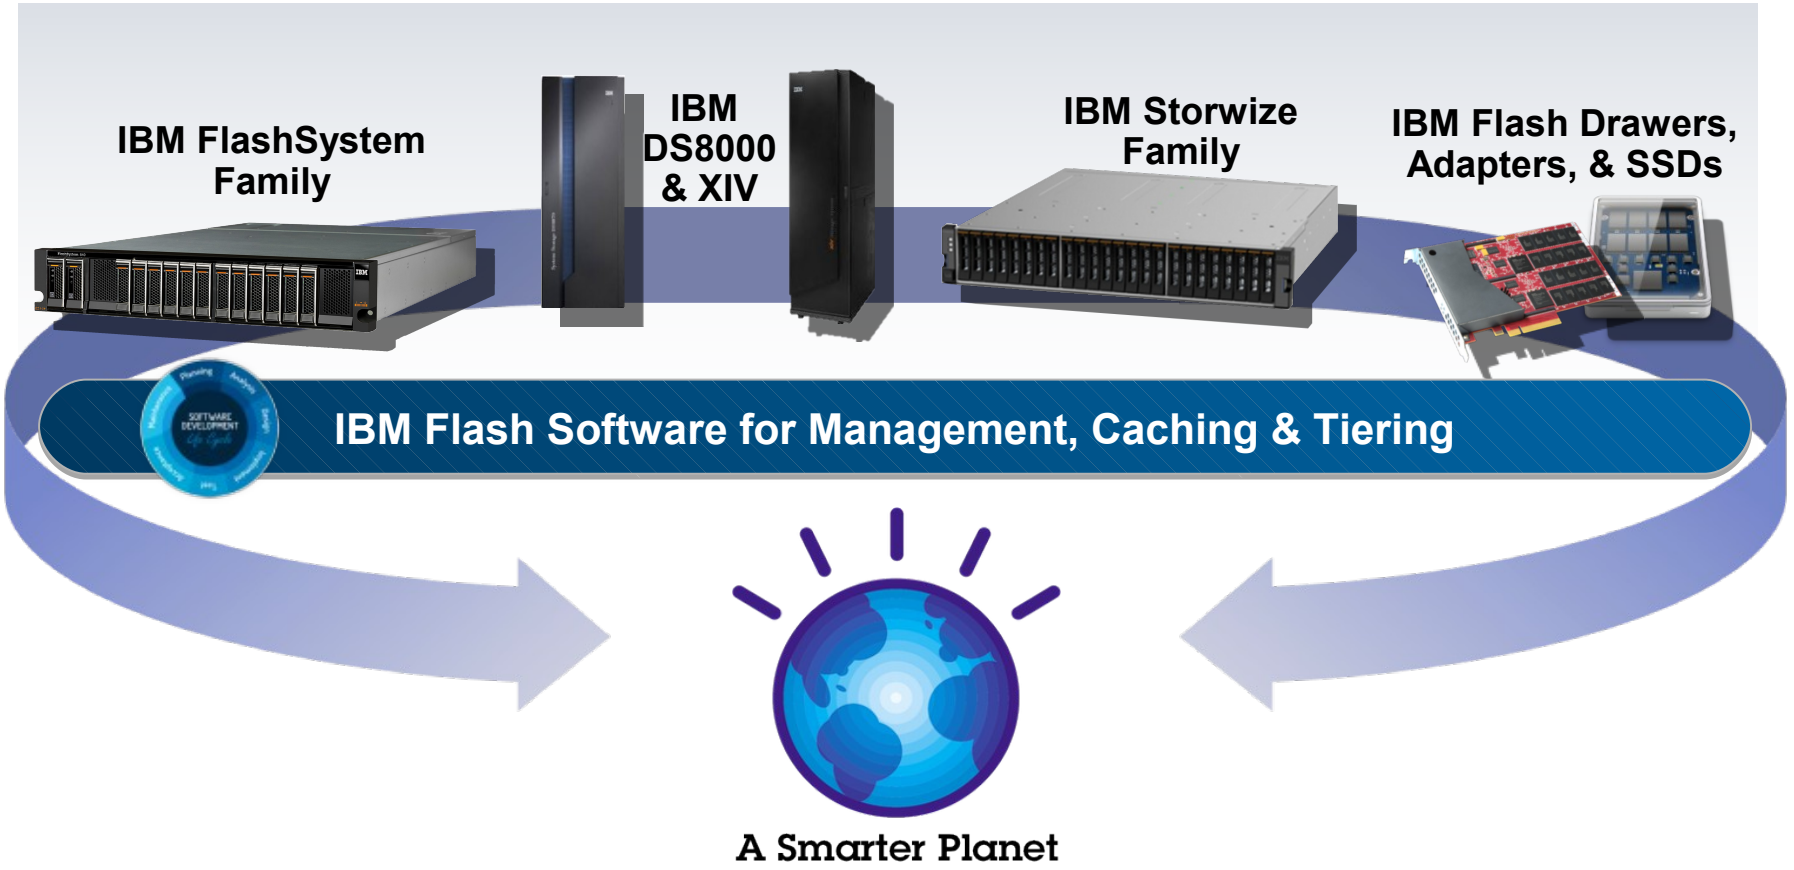

Why do this? Performance and Economics

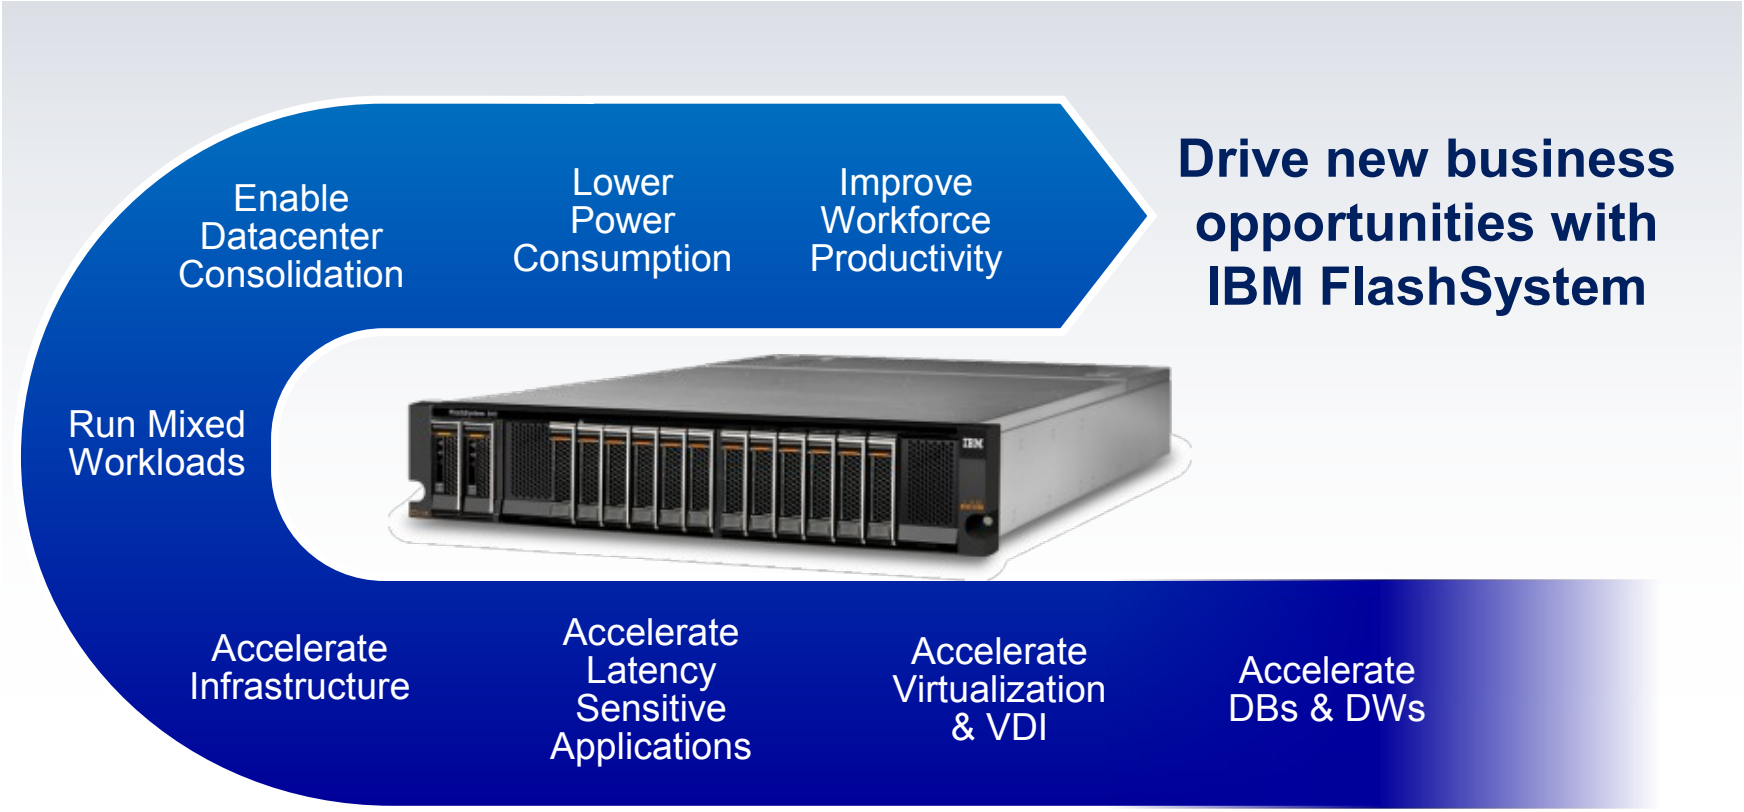

Summary

- **Flash is breaking new ground** and changing how we look at IT infrastructure and how we define performance.
- The **hard disk drives continues to have a future** in your data center…storing the data you rarely need to access.
- The **IBM FlashSystem family enables a new future** where IT is no longer constrained and pained by the need to deploy HDD for performance.

Learn more at **ibm.com/storage/flash**

Revised No-Charge SSD With POWER7+ 770 / 780

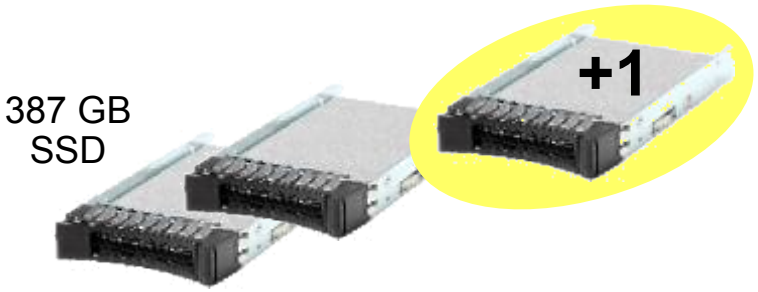

2 or 3 with new POWER7+ 770

$ 9,400 or $14,100 list price value

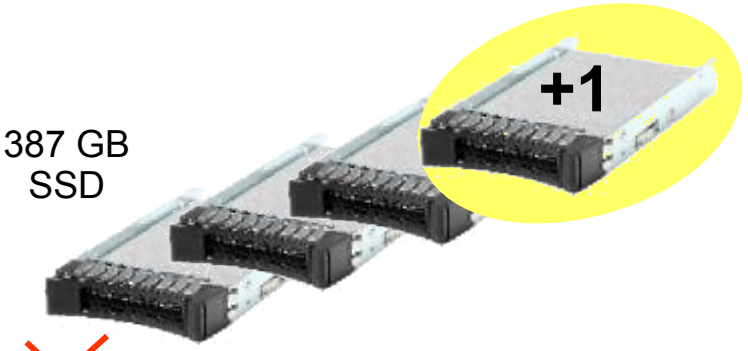

~~3~~ **4 with new POWER7+ 780**

$ 18,800 list price value

Leverage SSD to help unleash the performance capabilities of your Power 770/780 !

- Included with new POWER7+ 770 and 780 (or upgrades into a new 770/780 D model keeping same serial number)
- Flexibility to configure in system unit or I/O drawer
- eConfig will now include/require on all new/refreshed configurations run starting 11 Feb 2014 with the revised rules. Ordering/manufacturing starting week of 14 Feb.
- Transition period for 780 to accommodate existing proposals/orders.

Normal maintenance prices for SSD after 1 year warranty

IBM USA suggested list prices as of Jan 2014 for Power 770/780 and are subject to change without notice; Reseller prices may vary.

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.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, AIX 5L, Chiphopper, Chipkill, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Purpose File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, System i, System p, System p5, System Storage, System z, Tivoli Enterprise, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
UNIX is a registered trademark of The Open Group in the United States, other countries or both.
Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Itanium, Pentium are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.
Other company, product and service names may be trademarks or service marks of others.

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3, AIX 5L or AIX 6 were used. All other systems used previous versions of AIX. The SPEC CPU2006, SPEC2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
NotesBench	http://www.notesbench.org
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Microsoft Exchange	http://www.microsoft.com/exchange/evaluation/performance/default.asp
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised January 15, 2008

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX 5L were used. All other systems used previous versions of AIX. The SPEC CPU2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ABAQUS	http://www.abaqus.com/support/sup_tech_notes64.html select Abaqus v6.4 Performance Data
ANSYS	http://www.ansys.com/services/hardware_support/index.htm select "Hardware Support Database", then benchmarks.
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised January 15, 2008

Notes on performance estimates

- **rPerf for AIX**
- **rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.**
- **rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.**
- **All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.**
- **==========**
- **CPW for IBM i**
- **Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html**

Revised April 2, 2007