IBM POWER SYSTEMS
April 2011 Announcement Highlights

Rick Cowles
Power Systems FTSS
rcowles@us.ibm.com

Subject to Change Until Announced

Midsized businesses are the engines of a Smarter Planet.

90%

of the world's workforce are employed by small and midsized businesses, accounting for over 65% of global GDP.

IBM and our Business Partners provide solutions that help midsized businesses work smarter.

Subject to change until announced

The level of complexity continues to grow

81%

of US midmarket growth company CEOs anticipate greater complexity over the next five years.

42%

Feel prepared for it.

6%

Percent of available capacity used by the average commodity server.

30%

Number of servers in some organizations that sit unutilized.

70%

Percent of typical IT budgets devoted to managing, maintaining, securing and upgrading systems rather than building new capabilities, services and applications.

Capitalizing on Complexity. Insights from the Global Chief Executive Officer Study. IBM 2010
http://www-304.ibm.com/businesscenter/cpe/html0/199672.html

Midmarket IT Priorities and Goals

We are seeing dramatic shifts as our planet becomes smarter. These shifts are changing the way the world works. Nothing is changing more than IT. These changes are introducing new challenges for midsize businesses.

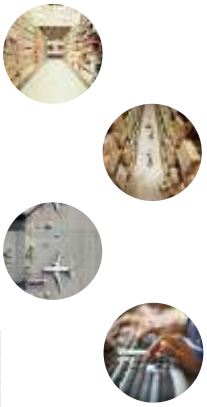

IT priorities for midsize business

- Ensure business continuity
- Improve system performance/speed
- Improve system uptime/availability
- Increase the utilization of existing IT infrastructure
- Replace aging servers/storage hardware

Midsize Businesses are looking for ways to:

- *Increasing Agility*
- *Managing Risk*
- *Drive Effectiveness & Efficiency*

*Source: IBM segment and audience profile for IT Managers in Unix and x86 midmarket segments

Power your planet.

Workload-Optimizing Systems

+

AIX® - the future of UNIX

Total integration with i

Scalable Linux® ready for x86 consolidation

Virtualization without Limits

- ✓ Drive over 90% utilization
- ✓ Dynamically scale per demand

Dynamic Energy Optimization

- ✓ 70-90% energy cost reduction
- ✓ EnergyScale™ technologies

Resiliency without Downtime

- ✓ Roadmap to continuous availability
- ✓ High availability systems & scaling

Management with Automation

- ✓ VMControl to manage virtualization
- ✓ Automation to reduce task time

Smarter Systems for a Smarter Planet.

Subject to change until announced

Power Systems is helping deliver higher quality services – by impacting the way humans communicate with computers

- **IBM Watson** represents the latest in a long line of groundbreaking innovations from IBM
- Watson can understand the meaning and context of human language, and rapidly process information to find precise answers to complex questions
- What's next for Watson?
  - Project with Columbia University and Maryland School of Medicine to provide healthcare and life sciences diagnostic assistance
  - Research agreement with Nuance Communications to develop and apply Watson to healthcare
  - Other fields of investigation range from enterprise knowledge management to IT help desk

To learn more about IBM Watson: http://www.ibm.com/watson/

Subject to change until announced

Power is Performance Redefined
Extending Watson's POWER7 DNA to new frontiers

BladeCenter PS703 & PS704 Express

Power 750 Express

Power 775

SDMC

Achieve superior economics for application server consolidation on Power blades and servers with new levels of scalability, energy and cost efficiency

Dramatically reduce project concept to delivery time for the toughest challenges in science with the world's most powerful & energy efficient supercomputer

Improve service delivery time via a unified, intuitive interface for physical and virtualized system resource management

New Power Systems Solutions with Rational and Sybase to dramatically reduce time to value

Rational. software

SYBASE

Subject to change until announced

POWER7 System Highlights

Balance System Design
- Cache, Memory, and IO

POWER7 Processor Technology
- 6th Implementation of multi-core design
- On chip L2 & L3 caches

POWER7 System Architecture
- Blades to High End offerings
- Enhances memory implementation
- PCIe, SAS / SATA

Built in Virtualization
- Mobility Memory Expansion
- VM Control

Green Technologies
- Processor Nap & Sleep Mode
- Memory Power Down support
- Aggressive Power Save / Capping Modes

Availability
- Processor Instruction Retry
- Alternate Process Recovery
- Hot Add & Services

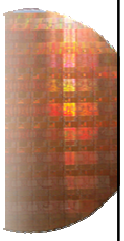

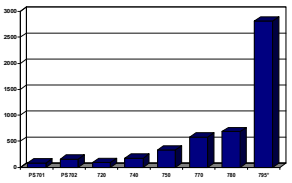

Subject to change until announced

2011 April – Selected Hardware Highlights

- New POWER7 Blades: PS703 & PS704
- Refreshed Power 750 & 755
- Higher max LPARs per server (750/770/780/795)
- Expanded use of #2319 Factory Deconfig for Entry Systems
- New Firmware 7.3 functionality
- HDD update
- New SFF I/O drawer
- New SAS-bay-based SSD
- Existing SSD update
- Enhanced 7016-1U2 Multi-Media Enclosure
- New 795 DVD/Tape Storage/Media drawer
- SODs
- Systems Director Management Console

Announce: April 12
Planned availability:
- May 20

eConfig/ordering support:
- April 12 for most things
- May 10 IBM/BNT Switches

IBM BladeCenter PS703 and PS704 Express
Smarter Computing with Power Blades

Higher blade density that delivers superior economics
- PS703 Express single-wide with 16-cores
- PS704 Express double-wide with 32-cores

Unprecedented performance on a Power blade that delivers new services faster
- Dual-stack 1.8" Solid State Disks
- Support for dual VIO Servers for improved reliability and I/O performance
- Greater performance with PCI-e Gen2 I/O bus

Improved mobility and management for delivering higher quality services
- Support for Live Partition Mobility between Power blades and rack servers
- Next-generation SDMC simplifies blade management
- Support for Active Memory Expansion enables more work to be done with existing server resources

PS703 & PS704 Express
- ✓ 16 or 32 cores
- ✓ Single or Double Wide
- ✓ 2.4GHz POWER7
- ✓ Up to 128GB of Memory (PS703)
- ✓ Up to 256GB of Memory (PS704)
- ✓ Solid State or Rotating Storage
- ✓ Supports multiple BladeCenter chassis (BCH, BCHT, BCS)

For more information: http://www.ibm.com/systems/power/hardware/

POWER7 PS703 Dual Socket

7891-73X

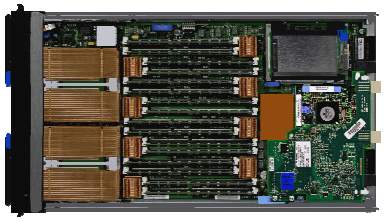

Architecture	16 Core @ 2.4 GHz
L2 & L3 Cache	On Chip
DDR3 Memory	Up to 128 GB / 16 DIMMs
1 SAS Bay	HDD: 0 - 1 (0-600GB) SSD: 0 – 2 (0-354 GB)
Daughter Card Options	CIOv & CFFh (PCIe Gen2 Support)
Integrated Options	Dual Port Gbt Ethernet Ethernet, USB

Fiber Support	Yes (via BladeCenter)
Media Bays	1 BladeCenter
Redundant Power	Yes BladeCenter
Redundant Cooling	Yes BladeCenter
Service Processor	Yes
Power & Thermal	POWER Save / Power Cap

Subject to change until announced

PS 703 Blade with SSD

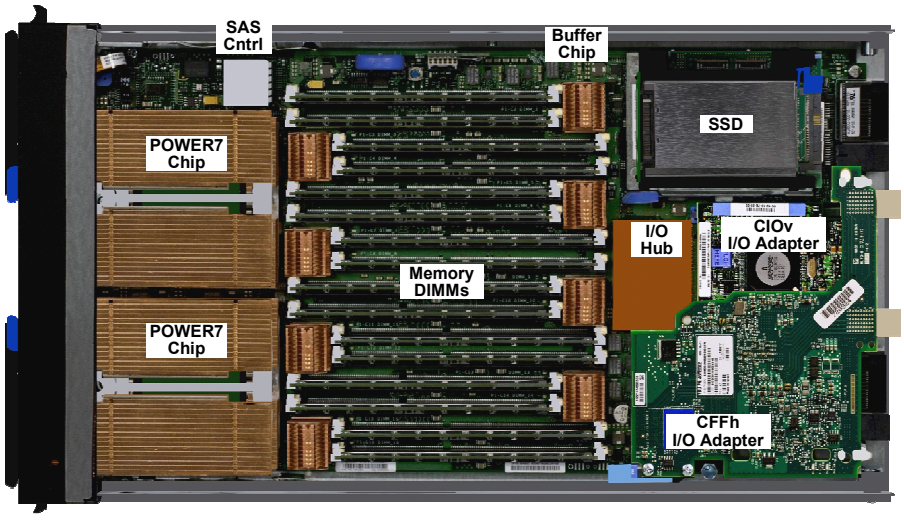

Subject to change until announced

POWER7 PS704 32 Cores Quad Socket

7891-74X

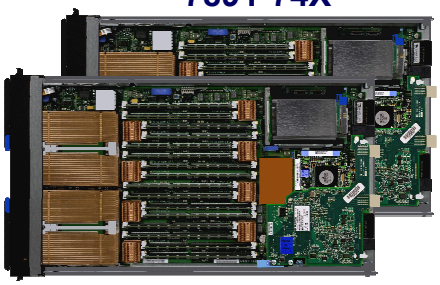

Architecture	32 Core @ 2.4 GHz
L2 & L3 Cache	On Chip
DDR3 Memory	Up to 256 GB / 32 DIMMs
2 SAS Bays	HDD: 0 - 2 (0-1200GB) SSD: 0 – 4 (0-708 GB)
Daughter Card Options	CIOv & CFFh / (PCIe Gen2 Support)
Integrated Options	Quad Port Gbt Ethernet Ethernet, USB

Fiber Support	Yes (via BladeCenter)
Media Bays	1 BladeCenter
Redundant Power	Yes BladeCenter
Redundant Cooling	Yes BladeCenter
Service Processor	Yes
Power & Thermal	POWER Save / Power Cap

Subject to change until announced

New POWER7 Blade SSD Option

PS703 (1 SAS bay) & PS704 (2 SAS bays)
- FC # 8207
- 177 GB

Bay can be for either HDD or SSD

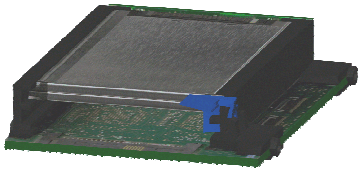

Interposer is required
- FC #4539
- Supports 1 or 2 SSD modules
- Provides the SAS to SATA conversion

SSD formatted with 528 byte sectors.

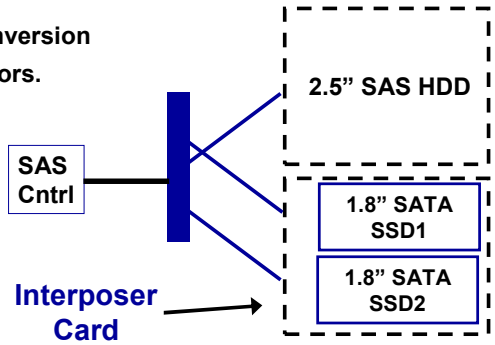

Subject to change until announced

PS Blade Performance Comparison

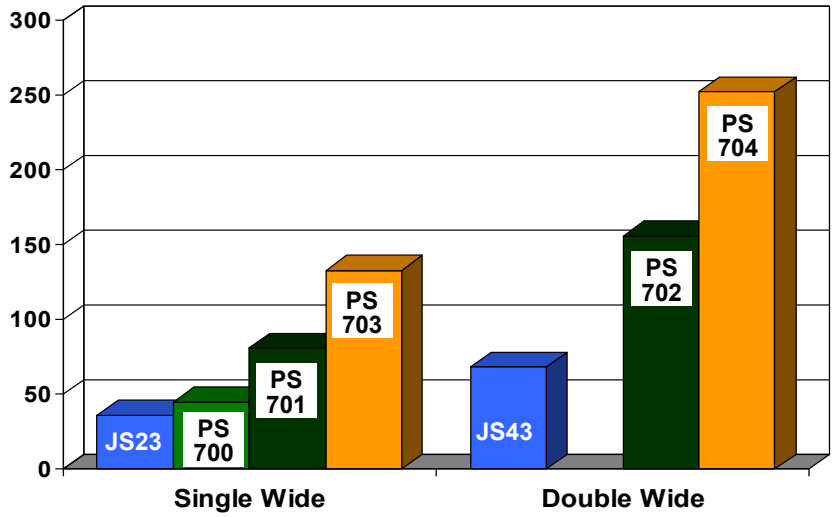

rPerf shown, CPW would be similar

Subject to change until announced

PS703 & PS704 IBM OS Support

AIX 5.3 with TL12 with Service Pack 4, or later
AIX 5.3 with TL12 with Service Pack 7, or later
AIX 6.1 with TL06 with Service Pack 5, or later
AIX 6.1 with TL05 with Service Pack 6, or later
AIX 6.1 with TL04 with Service Pack 10, or later
AIX 7.1 with Service Pack 3, or later

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

SUSE Linux Enterprise Server 11 Service Pack 1 for POWER, with current maintenance updates available from Novell *
Red Hat Enterprise Linux 5.6 for POWER, or later *
Red Hat Enterprise Linux 6.0 for POWER, or later *
* Users should also update their systems with the latest Linux for Power service and productivity tools from IBM's website: http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html

AIX and IBM i have the same software tier on the PS703/PS704 as the PS702 software tier

Subject to change until announced

PS703 and PS704 Options

Memory
- FC # 8196 8 GB (2x4GB RDIMMs) Memory
- FC # 8199 16 GB (2x8GB RDIMMs) Memory

Adapters / CIOv
- FC # 8240 Emulex 8 Gbt Fibre Channel Exp
- FC # 8241 QLogic 4 Gbt Fibre Channel Exp
- FC # 8242 QLogic 8 Gbt Fibre Channel Exp
- FC # 8243 Broadcom 2-Port Gbt Ethernet Exp
- FC # 8246 3 Gbt SAS Passthrough Expansion

Adapters / CFFh
- FC # 8252 QLogic Eth 4Gbt Fibre Exp. Card
- FC # 8271 QLogic 8 Gbt Fibre Chan / **Dual 1Gbt ENET** Exp Card
- **FC # 8272 2-Port QDR 40 Gbt/s Infiniband**
- FC # 8275 QLogic 2 port 10 Gb Converged
- FC # 8291 4-Port 1Gb Eth Expansion Card

Storage
- **FC # 8207 177 GB Solid State Drive**
- FC # 8274 IBM 300GB SAS 10K RPM SAS HDD
- FC # 8276 IBM 600GB SAS 10K RPM SFF

Power 750 Express 8233-E8B

~~1 to 4 Sockets~~
~~6 or 8 Cores per Socket~~
~~3.0 to 3.55 GHz~~
~~Energy-Star Qualified~~

- 1 to 4 Sockets
- 4 or 6 or 8 Cores / Socket
- 3.2 to 3.7 GHz
- Energy-Star Qualified

Power 750 Processor Card Positioning

Number of cores
GHz cores
Memory per core
I/O per core

4-core 3.7 GHz
- Fewest cores lowers server price
- High GHz
- Up to 32 GB memory per core
- Max I/O per core

6-core 3.7 GHz
- High GHz
- Up to 21 GB memory per core
- Good I/O per core

8-core 3.6 GHz
- Good GHz
- Up to 16 GB memory per core
- OK I/O per core

8-core 3.2 GHz
- Lowest cost per core (lowest GHz)
- Up to 16 GB memory per core
- OK I/O per core

Subject to change until announced

Performance on Power 750 vs POWER6 550

Even more performance in 2011

More performance per core

POWER6 550

# Cores	CPW	rPerf
8	37,950	78.6

Power 550 5.0 GHz

POWER7 750

# Cores	CPW*	+%	rPerf*	+%
8	47800	26%	85.29	9%
32	171400	352%	307.03	291%

Power 750 3.2 GHz

8	52700	39%	93.05	18%
32	183200	383%	334.97	326%

Power 750 3.6 GHz

More performance per system

Subject to change until announced

Power 755 Refresh

Faster processor offered

2011 proposals

32 Core @ 3.612 GHz

2010 proposals

32 Core @ 3.33 GHz

Up to 9.24 TF/ Rack (10 nodes per Rack)

Up to 8.4 TF / Rack (10 nodes per Rack)

755 Processor Card

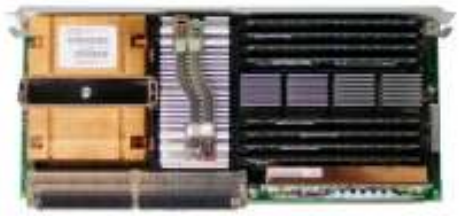

Processor Card

- **8-core 3.3 GHz #8332 – 4 per server (32 core)**
- **8-core 3.6 GHz #EPA1 – 4 per server (32 core)**

Processor activation feature structure

- Chargeable - none
- No-charge #2325 (3.3 GHz)
- No-charge #EPF1 (3.6 GHz)

Pricing insight - 100% activations are no charge

FC #2319 Factory De-configuration for Express Servers

- Allow customers to optimize SW licensing by only licensing the cores required by their workloads.
- Clients pay up-front for all hardware and activations
- Clients do not have to license de-configured cores
  - For each feature code #2319, manufacturing will de-configure one core, preventing its use
  - The default number of AIX licenses will be reduced accordingly for each #2319 feature
  - The default number of PowerVM licenses will also be reduced by one with each #2319
  - The default number of IBM Systems Director licenses will be reduced by one with each #2319
  - IBM i doesn't need #2319 because you can select fewer processor core licenses than the number of activated processor cores on the server
  - If customers need more cores at a later date, they can reconfigure them using the ASMI interface
- Terms and conditions of SW used on the server are met.

Note – the client needs to work with their ISVs to determine if #2319 is accepted by the ISV.

FC #2319 April 2011 Enhancements

#2319 applicability announced in 2010
- For Power 710/720/730/740 with 4, 6, 8-core
- Initially not for PowerVM, but added later in 2010

#2319 applicability announced in April 2011
- Also for Power 730 and 740 with 12-core or 16-core
- For Power 750 Express up to 32 cores

This function effectively takes the place of Capacity Upgrade on Demand for entry servers by controlling the number of active processor cores

LPAR / 7.3 Firmware Enhancements

Maximum Partitions supported / PowerVM

- Power 710/ 720: 80
- Power 730 / 740 160
- Power 750: 320
- Power 770 / 780: 640
- Power 795 1000

HMC: 1024 max partitions/HMC

HMC V7 R7.3.0

FW 7.3 for all systems

- Remote Restart (AIX/Linux)
- HMC Performance and scaling
- EnergyScale partition level power management
  - Per-core frequency control

EnergyScale Enhancements…

Partition Power Management (PPM)

- Allows the customer to set power management modes on a partition basis.
- Provides customer with a new interface in Active Energy Manager (IBM Systems Director plug-in) to put any partition into a supported power management mode.
- Allows dedicating/donating of cores to take place while maintaining the partition power mode.
- TPMD needs to be aware of partition/core groups.
- For dedicated processors only.

Power 795 Concurrent Maintenance

New eFW 7.3 firmware provides the following

- **Hot Node Add**
- Upgrade Memory
- **Hot Node repair**
- **Concurrent and Hot GX Adapter repair**
- **Concurrent System Controller repair**

NOTE: Requires Electronic Service Agent (Call-Home) and compliance with enablement guidelines

Subject to change until announced

SAS Hard Disk Drive (HDD) Options May 2011

3.5"	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
10k	n/a	n/a	n/a
15k	73 GB #3646 wfm	69 GB #3676 wfm	433B
15k	**146 GB #3647**	**139 GB #3677**	**433C**
15k	**300 GB #3648**	**283 GB * #3678**	**433D**
15k	**450 GB #3649**	**428 GB ** #3658**	**198E**

3 3.5" options available

SFF-1	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
10k	73 GB #1881 wfm	n/a	
10k	146 GB #1882 wfm	n/a	
10k	**300 GB #1885**	**283 GB** #1911**	**198D**
10k	**600 GB #1790**	**571 GB** #1916**	**19A3**
15k	73 GB #1883 wfm	69 GB ** #1884 wfm	198B
15k	**146 GB #1886**	**139 GB ** #1888**	**198C**

3 SFF-1 options available

New

wfm = withdrawn from marketing

* not supported as IBM i 5.4 load source
** IBM i 6.1 or later required

Subject to change until announced

SAS Hard Disk Drive (HDD) Options May 2011

3.5"	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
15k	73 GB #3646 wfm	69 GB #3676 wfm	433B
15k	**146 GB #3647**	**139 GB #3677**	**433C**
15k	**300 GB #3648**	**283 GB * #3678**	**433D**
15k	**450 GB #3649**	**428 GB ** #3658**	**198E**

3 3.5" options available

SFF-1	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
10k	73 GB #1881 wfm	n/a	
10k	146 GB #1882 wfm	n/a	
10k	**300 GB #1885**	**283 GB** #1911**	**198D**
10k	**600 GB #1790**	**571 GB** #1916**	**19A3**
15k	73 GB #1883 wfm	69 GB ** #1884 wfm	198B
15k	**146 GB #1886**	**139 GB ** #1888**	**198C**

3 SFF-1 options available

New

SFF-2	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
10k	**300 GB #1925**	**283 GB #1956**	**19B7**
10k	**600 GB #1964**	**571 GB #1962**	**19B3**
15k	**146 GB #1917**	**139 GB #1947**	**19B0**

3 SFF-2 options available

New

SFF-2 or Gen2 SAS drives used in new I/O drawer

Subject to change until announced

SAS Hard Disk Drive (HDD) – Prices 12 April 2011

3.5"	USA List price 520/720 March	USA List price 520/720 April
15k 73/69 GB wfm	$ 498	$ 498
15k 146/139 GB	**$ 498**	**$ 498**
15k 300/283 GB	**$1,150**	**$ 950**
15k 450/428 GB	**$1,599**	**$1,300**

No change to entry drive price

17% price reduction

SFF-1	USA List price 520/720 March	USA List price 520/720 April
10k 73/-- GB wfm	$ 498 wfm	$ 498 wfm
10k 146/-- GB wfm	$ 650 wfm	$ 650 wfm
10k 300/283 GB	**$1,050**	**$ 798**
10k 600/571 GB	n/a	**$1,200**
15k 73/69 GB wfm	$ 498 wfm	$ 498 wfm
15k 146/139 GB	**$ 798**	**$ 498**

18% price reduction

24% price reduction
100% more GB for $798

37% price reduction
100% more GB for $498

SFF-2	USA List price 520/720 March	USA List price 520/720 April
10k 300/283 GB	n/a	**$ 798**
10k 600/571 GB	n/a	**$1,200**
15k 146/139 GB	n/a	**$ 498**

Lowest $/GB
Only $2.00/$2.10 per GB

wfm = withdrawn from marketing

IBM USA suggested list prices as of April 12 2011. Reseller prices may vary. Prices subject to change without notice.

Subject to change until announced

#5887 EXP24S SFF Gen2-bay Drawer

2 U drawer with 24 SAS bays for HDD (contains no PCIe slots)
Supports SAS SFF HDD on POWER6 & POWER7
Gen2 (SFF-2) bays …. Different carrier/tray vs SFF-1

Attached to SAS PCI Adapter(s) or integrated SAS controller
AIX 5.3 or later
IBM i 6.1 or later (with or without VIOS)
Linux REHL 5.6 or later, SUSE 11 or later
VIOS 2.2.0.12 or later

New SFF Disk Only I/O Drawer (has no PCIe slots)

#5887 EXP24S SFF Gen2-bay Drawer
2 U drawer with 24 SAS bays for HDD

Compared to existing #5886 #EXP12S
List price is 45% less per drive bay
- EXP24S at $5400 = $225/bay
- EXP12S at $4950 = $412/bay

Maintenance price is about the same per drive bay
- EXP24S drawer is about 2X maintenance of EXP12S per drawer

Prices are USA suggested list prices as of April 2011 when ordered with the 720 server. Prices and are subject to change without notice. Reseller prices may vary.

New SFF Disk Only I/O Drawer (has no PCIe slots)

#5887 EXP24S SFF Gen2-bay Drawer
2 U drawer with 24 SAS bays for HDD

Greener

- Smaller foot print – saves floor space
  - Twice the number of drives (24 vs 12) in 2U rack space
  - Four times number of AIX boot drive partitions (4 vs 1) vs EXP12S
  - Same number of AIX boot drive partitions in ½ rack space vs #5802
- Energy efficient vs #5886 EXP12S
  - SFF HDD (~ 9W) vs 3.5-inch HDD (~17W) - almost 50% less per drive
  - EXP24S = 300 W max, EXP12S = 340 W max
    - Up to 128% more efficient per bay (340/12=28.3 W/bay 300/24=12.5 W/bay)
    - Up to 13% more efficient per drawer

Prices are USA suggested list prices as of April 2011 when ordered with the 720 server. Prices and are subject to change without notice. Reseller prices may vary.

Subject to change until announced

SAS I/O Drawers (no PCI slots in drawer)

	3.5-inch #5886 EXP12S	SFF #5887 EXP24S
Max HDD per drawer	12	24
Max SSD per drawer (April 2011)	8	0
Size in 19-inch rack	2U (2 EIA)	2U (2 EIA)
PCI SAS adapters which support	PCI-X & PCIe	PCI-X & PCIe
Capable of “partitioning” bays	no	Yes (AIX/Linux/VIOS)
Max SAS interface speed	3 Gb	6 Gb
Power Systems supporting	POWER5/6/7	POWER6/7
Minimum IBM i support	5.4	6.1
Minimum AIX support	5.3	5.3
USA List price on Power 720	$4950	$5400

Prices are USA suggested list prices as of April 2011 when ordered with the 720 server. Prices and are subject to change without notice. Reseller prices may vary.

Subject to change until announced

#5886 EXP12S Compare: HDD for Data

Want 24 HDD for data

Config 1 two #5886 EXP12S Drawers
- 4U of rack space
- 4 power cords (uses more PDU outlets)
- PCI adapter(s) and 2 Y cables
- $9,900 list price for 2 drawers
- plus cost of 3.5-inch drives
- 3.5-inch HDD about 2X energy of SFF HDD

Config 0 One #5887 EXP24S Drawer
- 2U of rack space
- 2 power cords
- PCI adapter(s) and 1 or 2 Y cables
- $5,400 list price for 1 drawer
- plus cost of SFF drives
- SFF HDD about ½ energy of 3.5-inch HDD
- Maint for one #5887 ~same for two #5886

EXP24S
50% rack savings
45% lower price
50% energy savings
Same maint cost

Prices are USA suggested list prices as of April 2011 when ordered with the 720 server. Prices and are subject to change without notice. Reseller prices may vary.

Subject to change until announced

#5887 EXP24S SAS Adapters/Controllers

POWER7 or POWER6* PCI adapters

- PCIe 380 MB Cache SAS RAID Adapter (#5805, #5903)
- PCIe SAS Adapter (#5901, #5278) (not with IBM i)
- PCI-X 1.5 GB Cache SAS RAID Adapter (#5904, #5906, #5908)
- PCI-X SAS adapters not supported: #5900, #5902, #5912

POWER7 or POWER6* integrated SAS controllers via integrated SAS port in rear of CEC

- Power 710**/730, Power 720**/740, Power 750/755, Power 770/780
- Power 520* **, Power 550*,

* POWER6 unified product structure (8203-E4A, 8204-E8A, 8233-EMA, 9117-MMA, 9119-FHA), NOT 9408-M25, 9409-M50, 9406-MMA
** Not 4-core Power 710/720, Not 1-core Power 520

Subject to change until announced

#5887 EXP24S Modes

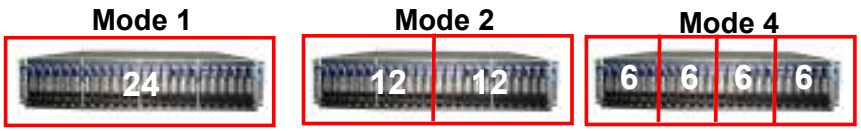

Mode 1	Mode 2	Mode 4
1 set 24 bays AIX, IBM i, Linux, VIOS	2 set 12 bays AIX, Linux, VIOS Not IBM i	4 set 6 bays AIX, Linux, VIOS Not IBM i
All adapter/controllers	All adapter/controllers	Only #5901/5278

IMPORTANT:
#5887 Modes are set by IBM Manufacturing
No option to reset after IBM ships is announced
Order this correctly!!!!

Note: if ordering multiple EXP24S in an MES order, avoid mixing modes within that order. There is no externally visible indicator or switch as to the drawer's mode.

EXP24S Connectors

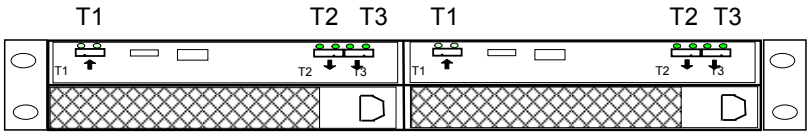

Rear of #5887 EXP24s

- In mode 1 with unpaired/non-dual adapters
  - two of the six ports are used.
    - Two T1s are used with a YI cable
    - Two T3s are used with an a YO cable
- In mode 1 with paired adapters,
  - four of the six ports are used – two T2s and two T3s
- In mode 2, four ports used to access both sets of bays
  - Two T2s and two T3s
- In mode 4, four ports used to access all four sets of bays
  - Two T2s and two T3s

EXP24S Cabling to CEC's Integrated SAS port

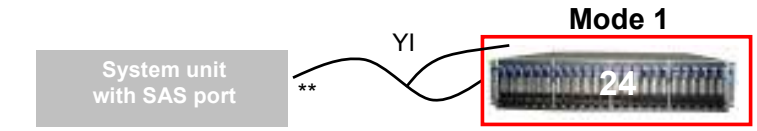

710/720/730/740/750/755/770/780
520*/550*

For CEC SAS port attach:
- Mode 2 or 4 not supported
- HDD only, no SSD in EXP24S

Uses one SAS YI cable attached to the two T1 ports on EXP24S
- Lengths: 1.5 m (#3686), 3 m (#3687)

Mfg code 9384

Supported by AIX/IBM i/Linux/VIOS

* 8203-E4A, 8204-E8A supported --- not 9408-M25, not 9409-M50

** For integrated SAS port to be enabled, use the write cache options of the integrated SAS controllers in the CEC. Split backplane options of CEC do not enable the integrated SAS port

Subject to change until announced

Cabling w/ SAS #5805/5903 Adapter – Mode 1

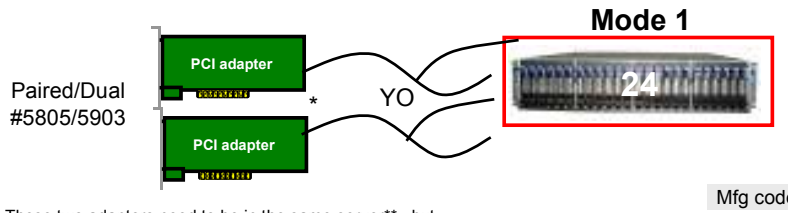

Mfg code 9367

These two adapters need to be in the same server**, but can be in different I/O drawers and/or processor enclosures

Supported by AIX/IBM i/Linux/VIOS

* One of adapter's 2 SAS connectors used in diagram above. The 2nd connector could be used to attach another SAS device such as another EXP24S
** Special exception for AIX PowerHA configurations

Subject to change until announced

Cabling w/ SAS #5904/6/8 Adapter – Mode 1

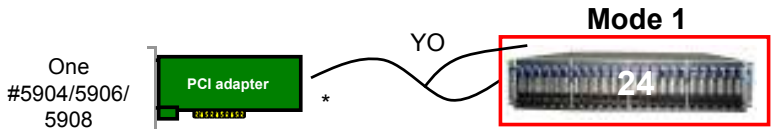

Mfg code 9382

Supported by AIX/IBM i/Linux/VIOS

* One of adapter's 3 SAS connectors used in diagram above. The 2nd and 3rd connectors could be used to attach additional SAS device such as another EXP24S or EXP12S or
** Special exception for AIX PowerHA configurations

Ordering EXP24S - Mfg Specify Codes

To help communication to IBM Manufacturing exactly what should be built, the following no-charge specify codes must be included with EXP24S orders ….
One per EXP24S

Specify	Mode	Adapter/Controller	Cable to Drw	Environment
#9359	Mode 1	One #5901/5278	1 YO cable	AIX/Linux/VIOS
#9360	Mode 1	Pair #5901/5278	2 YO cables	AIX/Linux/VIOS
#9361	Mode 2	Two #5901/5278	2 YO cables	AIX/Linux/VIOS
#9365	Mode 4	Four #5901/5278	2 X cables	AIX/Linux/VIOS
#9366	Mode 2	Two pair #5901/5278	2 X cables	AIX/Linux/VIOS
#9367	Mode 1	Pair #5903/5805	2 YO cables	AIX/IBM i/Linux/VIOS
#9368	Mode 2	Two pair 5903/5805	2 X cables	AIX/Linux/VIOS
#9382	Mode 1	One #5904/06/08	1 YO cable	AIX/IBM i/Linux/VIOS
#9383	Mode 1	Pair #5904/06/08	2 YO cables	AIX/IBM i/Linux/VIOS
#9384	Mode 1	CEC SAS port	1 YI cable	AIX/IBM i/Linux/VIOS

After the order has been shipped, the specify is interesting and perhaps helpful, but does not prevent the client from adjusting the configuration (except mode) to better match their needs and no longer match the mfg specify code.

#5887 EXP24S Config Maximums

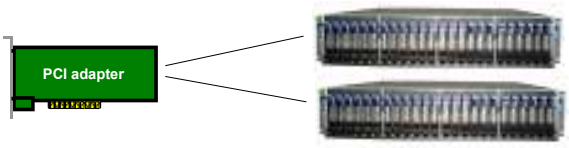

By POWER server
- 710/730 = 4; 720/730 = 14; 750 = 21; 755 = 7; 770/780 = 56; 795 = 168
- 520 = 12; 550 = 24; 560 = 56; 570 = 56; 595 = 93

168 = 4000 HDD

By Controller – Max HDD
- #5901/5278 PCIe Adapter with two SAS connections (max 48 HDD)
  - Max 2 mode1 #5887 … one per connector
  - Or max 1 #5887 & another SAS device on the other connector
- #5805/5903 PCIe 380MB RAID Adapter with two SAS connectors (max 48 HDD)
  - Max 2 mode1 #5887 … one per connector
  - Or max 1 #5887 & another SAS device on the other connector
- #5904/06/08 PCI-X 1.5GB RAID Adapter with three SAS connectors (max 60 HDD)
  - Max 2 mode1 #5887 … one per connector … 3rd connector can attach a #5886
  - Or max 1 #5887 plus 3 #5886 (2 of #5886 cascaded)
- Integrated SAS controller
  - Max 1 #5887

By Controller – Max SSD
- SSD not supported in EXP24S … no mixing SSD & HDD on #5805/5903 or on #5904/6/8

No Cascading #5887s Max one #5887 per SAS connector

Subject to change until announced

#5887 EXP24S Physical Planning

- Physical Specifications
  - Width: 448.6 mm (17.7 in) Depth: 530.0 mm (20.9 in)
  - Height: 87.4 mm (3.4 in) with support rails (2U)
  - Weight: 17.7 kg (39 lb) w/ no drives; 25.4kg (56 lb) w/ 24 drives
- Electrical/Thermal
  - Operating voltage: 110* - 240 V AC at 50 - 60 Hz single phase
  - Max power usage with 24 HDD -- 320 VA or 300 W (1024 Btu/hr)
- Acoustic
  - Only 6.0 bels (Quieter than EXP12S #5886's 6.5 bels)
- Maximum altitude: 2134 m (7000 ft) above sea level

- Rails are FIXED LENGTH and designed to fit Power Systems provided racks with 28 inch depth (715 mm) such as 7014-Txx
  - Other racks may have different depths and these rails won't adjust
  - No adjustable depth rails are orderable at this time

* Power Systems assumes racking and rack mounted 200-240V PDUs. Thus no 100V or wall-attached power cords are orderable, only drawer to PDU power cords.

Subject to change until announced

EXP24S Miscellaneous

Boot / load source support in EXP24S … POWER6 or POWER7
- Yes, boot drive with a supported SAS adapter / controller
- Yes, Load Source drives supported

EXP24S Pre-reqs
- AIX 5.3 with 5300-11 Technology Level, or later
- AIX 5.3 with 5300-10 Technology Level and Service Pack 2, or later
- AIX 5.3 with 5300-09 Technology Level and Service Pack 5, or later
- AIX 5.3 with 5300-08 Technology Level and Service Pack 8, or later
- AIX 6.1 with 6100-04 Technology Level, or later
- AIX 6.1 with 6100-03 Technology Level and Service Pack 3, or later
- AIX 6.1 with 6100-02 Technology Level and Service Pack 6, or later
- AIX 6.1 with 6100-01 Technology Level and Service Pack 7, or later
- IBM i 6.1 with 6.1.1 machine code or later
- SUSE Linux Enterprise Server 10 Service Pack 1 or later
- SUSE Linux Enterprise Server 11 or later
- Red Hat Enterprise Linux 4.6 or later
- Red Hat Enterprise Linux 5.1 or later

EXP24S CCIN = 2BE2

EXP24S SFF-2 or SFF Gen2 Hard Disk Drive

- Different carriers/trays which hold the HDD fit into the EXP24S SAS bays. These do NOT fit into #5802/5803 12X PCIe I/O drawers of into SFF bays of the Power System units.
- There is no conversion offered between Gen1 (SFF-1) and Gen2 (SFF-2) drives

SFF-2	512-byte sectors AIX/Linux formatted	528-byte sectors IBM i formatted	CCIN
10k	300 GB #1925	283 GB #1956	19B7
10k	600 GB #1964	571 GB #1962	19B3
15k	146 GB #1917	139 GB #1947	19B0

Qty 150 #1925 = #1869 Qty 150 #1956 = #1844
Qty 150 #1964 = #1818; Qty 150 #1962 = #1817,
Qty 150 #1917 = #1866; Qty 150 #1947 = #1868

Power Systems SSD Configuration Options

SAS-bay-based

- Option introduced 2009 – Enhanced 2011
  - 69GB SSD
  - 177GB SSD

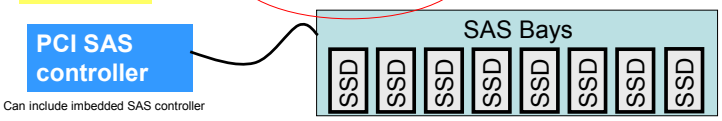

PCIe-based

- Introduction September 2010
  - 177GB SSD

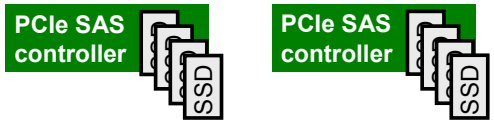

177GB SSD for #5803 SAS Bays

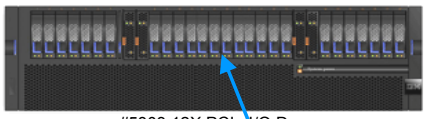
#5803 12X PCIe I/O Drawer

- For POWER7 795
  - Did not have 177GB SSD before
  - No other servers announced April 2011 for this SSD
  - No other I/O drawers announced April 2011 for this SSD
- Placed in #5803 SFF SAS bay
- Hot swap capability – just like HDD
- Controlled by #5805/5903 PCIe 380MB RAID Adapters in #5803
- Support
  - AIX 5.3 or later
  - IBM i 6.1 or later (VIOS optional)
  - Linux: REHL 5.6 or later, SUSE 10 or later
  - VIOS 2.2.0.X or later

177GB SSD for SAS Bay

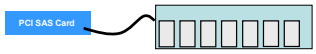

69GB SFF SSD

177GB SFF SSD

Greener

- 177GB SSD is 2.5x more GB per SAS bay vs older 69GB

Better Price

- 30% lower list price per SAS-bay-based drive
  - $4700 per drive vs existing $6882 per drive
- Nearly 75% lower list price per GB
  - $26.6/GB vs existing $100/GB

Prices are USA suggested list prices as of April 2011 when ordered with the 795 server. Prices and are subject to change without notice. Reseller prices may vary.

Subject to change until announced

SSD SAS-Bay-Based Options

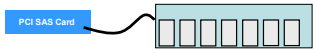

SAS-Bay-Based SSD	Announced 2009	Announced 2011
GB capacity	69 GB	177 GB
Flash Technology	SLC	eMLC
Performance	Very similar, but answer varies by workload and randomness of data	
Hot plug	Y	Y
3.5-inch SAS bay	Y	N
SFF SAS bay	Y	Y
On POWER6 configurations	Y	N
On POWER7 710/720/730/740 750/755/770/780 configurations	Y	N
On POWER7 795	Y	Y
IBM i minimum support	5.4 or later	6.1 or later
AIX minimum support	5.3 or later	5.3 or later

2.5x more per drive

Subject to change until announced

New 177GB SSD: Configuration Details

177GB SFF-1 SSD w/ eMLC
- #1775 for AIX/Linux/VIOS
- #1787 for IBM i
- #1775/1787 physically identical, but 2 features for IBM configurator usage
- CCIN = 58B3
- Fits Gen1 SAS bay of #5803 12X PCIe I/O Drawer

Placement
- Only in #5803 12X PCIe I/O Drawer on Power 795

Controlled by
- #5805/5903 380MB RAID SAS Adapter in the #5803

Configuration rules
- Protection: same as HDD and other SSD. Must protect for IBM i. Highly recommended AIX/Linux/VIOS
- Max 9 SSD per #5805/5903 pair
- Max 18 SSD per #5803 in mode 2 using two pairs of #5805/5903 adapters
- Max zero SSD in #5803 in mode 1
- Max zero SSD in #5803 in mode 4
- Can not mix SSD and HDD on one pair of #5805/5903 adapters, but can mix SSD and HDD in the same #5803 in mode 2 under different adapter pairs
- Can not mix 69GB SSD and 177GB SSD in the same array, but can mix 69GB SSD and 177GB SSD on the same pair of SAS adapters
- A 177GB SSD can be the RAID hot spare for a 69GB SSD, but not vice versa

Performance / usable capacity insight
- For best performance the #5805/5903 pair should run more than one array to take advantage of Active/Active capability. This means that assuming RAID5 and 8 drives attached in two arrays, you get 6 drives capacity for the 8 physical drives. Or if RAID5 and 9 drives in two RAID5 arrays, you get seven drives of capacity.

Existing PCIe-based SSD Update

1- IBM i 6.1 support of PCIe-based SSD without VIOS
- Available May 2011 with latest IBM i fixes

2- The #2054 available for Power 720, 740, 750, 755

3- In March 2011 limited the #2053 in the 710/730 to a max of 2 SSD per #2053 adapter. Submit RPQ request if this is a problem.

4- Just 3 #1995/1996 SSD modules in one PCIe SSD adapter (#2053/2054/2055) – OS mirroring two adapters still recommended
- eConfig support to configure and order planned 2011 May 10

Existing SAS-bay-based SSD Update

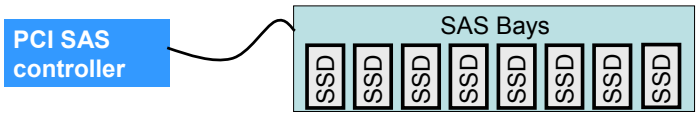

1- If you put a 69GB SSD in a 710/720/730/740, then we need to speed up the fans. Is noisier, especially the 720/740..
- If first added to server as MES order, client needs to turn off "acoustic mode" using ASMI

2- 69GB SSD are being withdrawn from marketing effective 29 July 2011
- Note, IBM has not been able to obtain/build additional 69GB SSD. Once existing inventory is sold, we are out. It is possible this could happen before July.

IBM i Technology Refresh 2 & More

IBM i suspend / resume –
- Ability to suspend a partition and resume it from the suspend point through PowerVM (Enterprise or Standard Edition)
- Customer Value: Saves IPLs. Suspend before CEC maintenance, do the repair or CEC IPL and then resume.
- One of the underlying technologies required prior to future potential Live Partition Mobility

I/O performance – enhanced multi-path –
- Smarter load balancing multi-path algorithm that tries to push more I/O faster
- Example: When using NPIV, running both disk and tape traffic on the same port more feasible since the disk traffic will move off the port if the tape traffic starts slowing down the disk performance.

Support of new hardware
- 571GB disk, EXP24S I/O drawer, Power 795 177GB SSD, and more

IBM i Technology Refresh 2 & More

IBM i to IBM i virtual tape –

- In an environment where an IBM i partition is hosting other partitions
- Allow multiple IBM i partitions to share the tape drive owned by a another IBM i partition
- No longer need to reassign the tape drive (actually assign its connecting adapter) to a different partition
- Alternative to VIOS/NPIV sharing
  - Note: no tape library robotics controls. That needs VIOS/NPIV

Enhanced support of existing hardware

- 6.1 support of existing PCIe-based 177GB SSD
- 7.1 support of PCIe Gen2 Riser Card

DS5000 NPIV support –

- Allow better utilization of FC adapter/switch resources by using NPIV with the DS5300/DS5100
- For POWER6/POWER7 servers (excepting Blades)

7216-1U2 Storage Enclosure Enhancements April 2011

1U 19-inch rack mount drawer
First shipped 3Q 2010
Supported on POWER7 or POWER6 servers
Max 2 HH drives or max 4 DVD or max 1 HH & 2 DVD drives
7216 can attach to more than one server

April 2011 announcements:

- Add SAS DAT160
- Announce marketing withdrawal DAT320 effective July 2011
- Add capability to support multiple drives with one SAS adapter

- SODs:
  - Add USB DAT160 drive & add 1TB RDX drive

7216-1U2 Cabling Enhancement April 2011

One 7216 can physically attach to more than one server or partition

April 2011, 7216-1U2 announces a new quad cable which allows one SAS adapter to attach to up to 4 SAS/SATA drives in the 7216.
SAVE PCI adapters and PCI slots

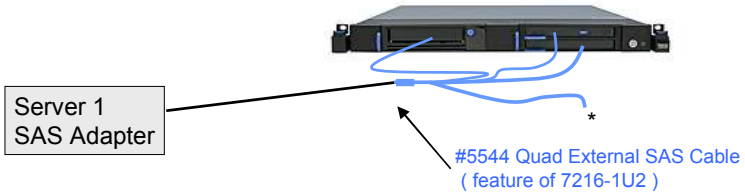

* Multiple 7216-1U2s can be connected to one #5544 cable

#5724 Storage/Media Drawer – Power 795

#5724 DVD/Tape SAS External Storage Unit w/SATA Media

- For Power 795 CEC rack – in either 12U or 34U location
- Max one drawer per 795
- Holds up to 1 HH tape drive and up to 2 DVDs
  - HH DAT160 (#5619)
  - DVD SATA (#5762)
  - HH LTO-5 (#5638)

Replaces #5720 media drawer in new Power 795 proposals

April SODs

- Systems Director Management Console
  - IBM intends to enhance the interoperability of Systems Director Management Console with IBM Systems Director management servers.
- PCIe Riser Card (Gen2) support by IBM i 6.1
  - IBM plans to support the PCIe Riser Card (Gen2)(#5685) without VIOS in an IBM i 6.1 environment with machine code 6.1.1 in fourth quarter 2011.
- PCIe SAS adapter
  - IBM plans to introduce a large-cache PCIe SAS adapter in 3Q2011 for clients with large numbers of HDD and/or SSD per adapter. It is planned to be supported on POWER7 configurations (Power 720, 740, 770, 780, 795), POWER7 750 servers with 12X PCIe I/O drawers, or on POWER6 servers with 12X PCIe I/O drawers. It is planned to support SSD or HDD located in #5802/5803 12X PCIe, #5886 EXP12S and #5887 EXP24S I/O drawers. It is planned to be supported by AIX 5.3 or later, IBM i 6.1 with Virtual IO Server (VIOS), or IBM i 7.1 with or without VIOS. Support by IBM i 6.1 without VIOS is planned 4Q2011. The adapter is planned to be configured as pairs of single-wide PCIe cards, much like the existing #5805/5903 PCIe adapters.

PCIe SOD Positioning

PCIe SAS Adapters
- Zero cache (existing)
- 380MB cache (existing)
- Large cache (SOD)
  - high performance, supports lots of internal HDD and/or SSD

High interest to IBM i clients with internal Power HDD

High interest to AIX/IBM i clients with larger amounts of Power SSD

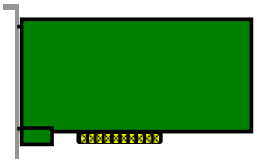

April SOD – PCIe SAS Adapter

- IBM plans a large-cache PCIe SAS adapter in 3Q2011
- For clients with large numbers of HDD and/or SSD per adapter
- Planned to be supported on
  - POWER7 configurations (Power 720, 740, 770, 780, 795)
  - POWER7 750 servers with 12X PCIe I/O drawers
  - POWER6 servers with 12X PCIe I/O drawers
- Planned to support SSD or HDD located in
  - #5802/5803 12X PCIe drawers
  - #5886 EXP12S drawers
  - #5887 EXP24S I/O drawers.
- Planned to be supported by
  - AIX 5.3 or later
  - IBM i 6.1 with VIOS
  - IBM i 7.1 with or without VIOS
  - IBM i 6.1 without VIOS is planned 4Q2011
- Planned to be configured as pairs of single-wide PCIe cards, much like the existing #5805/5903 PCIe adapters.

Existing 2010 SSD SOD

IBM plans to enhance its Power Systems Enterprise Class SSD solutions with technology designed to continue to provide significant improvements in performance and storage density over time. IBM plans for these IBM Power Systems enhancements to include both SAS-bay-based and PCIe-based SSD product offerings that will leverage IBM's investments in its SSD optimized Enterprise Class RAID Storage Controllers.

For Web listing of hardware SODs:
www.ibm.com/systems/power/hardware/sod.html

IBM® Systems Director – Platform Management

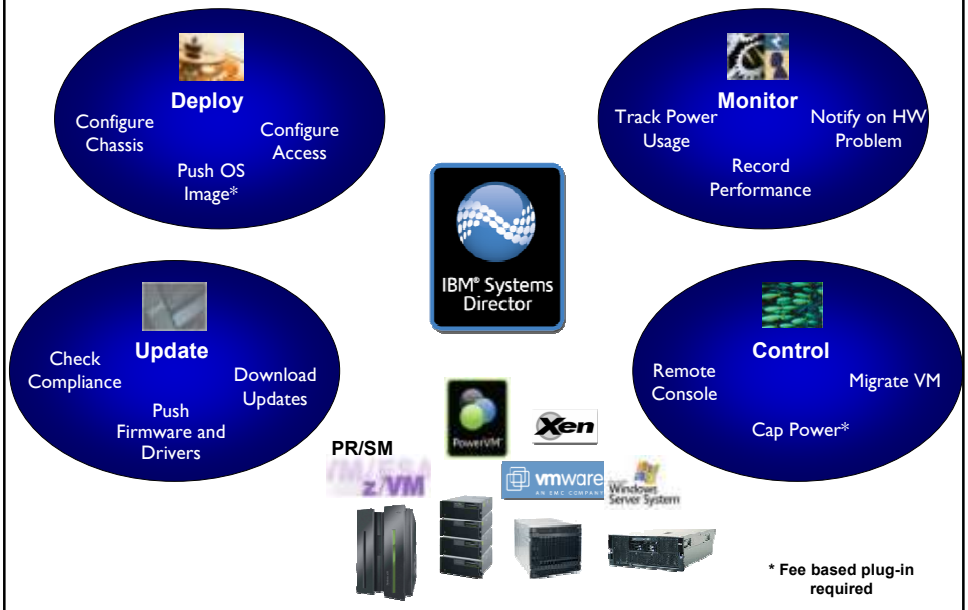

* Fee based plug-in required

Subject to change until announced

Systems Director Editions

Express Edition

- ✓ View the relationships between systems
- ✓ Manage virtual machines across multiple hypervisors
- ✓ Remotely access system status and updates

IBM® Systems Director

Standard Edition

- ✓ Reduce time to deploy virtual AIX and Linux workloads
- ✓ Monitor and control energy use
- ✓ Monitor and configure networking systems

Enterprise Edition

- ✓ Automate workload deployment in system pools of AIX and Linux virtual machines
- ✓ Analyze real-time and historical status
- ✓ Analyze predictive resource capacity data

Subject to change until announced

Systems Director Management Console

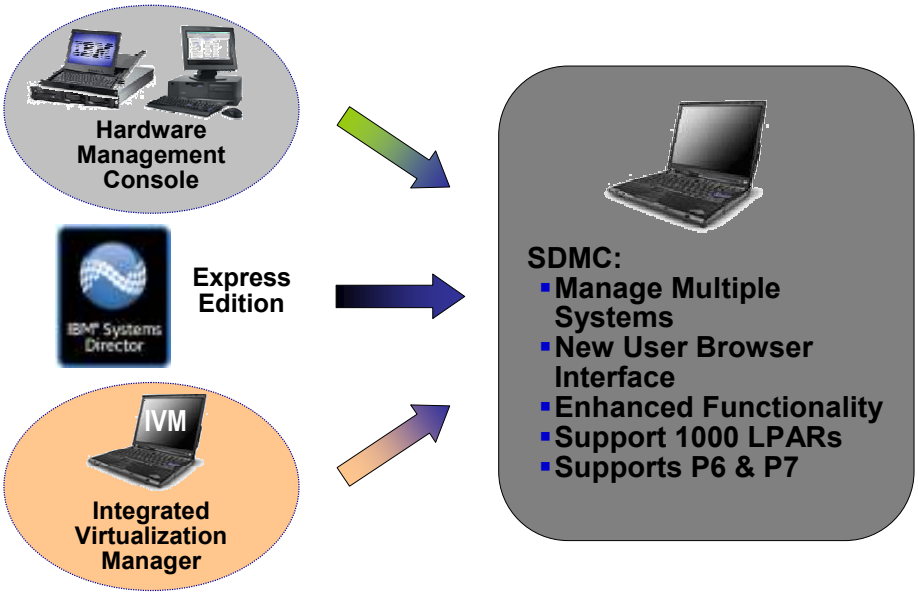

Subject to change until announced

Power Systems Management Options

Power Systems clients have three different management options

Product	Offering / Packaging	Function
HMC	Physical appliance (uses Linux)	Complete virtualization and service management for Power
IVM	Lightweight utility (built into VIOS)	Basic configuration and service for Power
SDMC	Software (installed on management server)	Comprehensive cross-platform mgmt

Strategy: Converge on IBM Systems Director Management

- Standardize the user interfaces. Easier for clients to transition and scale up
- Converge branding: IBM System Director is the strategic management tool

Subject to change until announced

Systems Director Management Options

Hardware Appliance

IBM Provided Hardware

Plus Software Appliance:
- SDMC Management Base
- Systems Director Express Edition

Turnkey Solution

Software Appliance

Software Appliance:
- SDMC Management Base
- Systems Director Express Edition

Non Turnkey Solution

Utilize existing IBM x86 virtualization infrastructure

Customer provided hardware
- Finite set of server options

SDMC Feature Adds Vs HMC

Capability	SDMC	HMC
Simplified Virtualization	✓	
Blades Management	✓	
Mobility between Blades and Servers	✓	
Dual VIOS for Blades	✓	
Active Memory Expansion for Blades	✓	
Virtual Appliance Option	✓	
Consistent Look and Feel	✓	
OS Management and Monitoring	✓	

- SDMC provides simplified, easy-to-use, IVM-like virtualization management
- SDMC can manage Blades
- SDMC is offered as a virtual as well as a turnkey hardware appliance
- SDMC integrates Power hardware, service and virtualization management into a common look and feel (Director)

SDMC Function - Differences

Enhanced Virtualization Management:

- IVM-like Ease-of-Use:
  - Enhanced VIOS / Virtual Adapter management – SDMC will manage your virtual slots automatically.
  - More intuitive DLPAR – ability to modify resource assignments whether the partition is On or Off.

Terminology:

- Fundamental terminology changes.
  - Example: an 'LPAR' is now a 'virtual server'.
  - Managed systems are now 'servers' or 'hosts'.

Additional Function:

- Director provides a lot of additional function including things such as AEM, Image Manager, etc.

Terminology Differences

HMC Terminology	SDMC Terminology
Managed System	Server / Host
Hardware Management Console	Platform Manager
Frame / BPA	Power Unit
LPAR	Virtual Server
VIOS	Utility Virtual Server
Users: hscpe / hscroot	PE sysadmin
Partition Mobility	Relocation
Remove connection	Remove MEP (Managed End-Point)
hmcuser	SMUser
hmcoperator	SMManager

SDMC Function Mapping

Function on HMC	Function on Director / PSM
CEC and Frame Management	Ported from HMC. GUI is new
PowerVM Partition Virtualization	Ported from HMC. GUI is new.
Guided Repair	Ported from HMC. GUI is new
Concurrent Maintenance	Ported from HMC. GUI is new (POWER7 only)
Problem Analysis	Ported from HMC. GUI is new
Serviceable Event Management	Provided by Director: Service & Support Manager
Call Home	Provided by Director: Service & Support Manager
Updates for Firmware / Device microcode	Provided by Director: Update Manager, and ported from the HMC
HMC Appliance Management	Provided by Director: Console Management (new)
Function on IVM	**Function on Director / PSM**
Management of Power Blades	PSM functionality expanded
Command Line Interface	PSM functionality expanded

ion

HMC / IVM / SDMC Comparison – General

General Characteristics	IVM	HMC	SDMC
Delivery Vehicle	**Integrated into the server**	**A desktop or rack-mounted appliance**	**Virtual appliance: Customer provided hardware & hypervisor Physical appliance: A desktop or rack-mounted appliance**
Footprint	**No overhead beyond VIOS. Runs in 60MB memory / minimal CPU**	**2-Core x86, 4GB RAM, 1x500GB HD (latest – used to run with less)**	**4-Core Nahalem x86, 8GB RAM, 2x500GB HD**
Installation	**Installed with the VIOS (optical or network). Preinstall option available on some systems.**	**Appliance is preinstalled. Reinstall via optical media or network is supported.**	**Software appliance: Virtual appliance tooling – apply and go. Relies on third-party / customer tooling. Hardware appliance: Preinstalled.**
Servers supported	**Blades: JS21 & beyond P5/5+: 560Q Express and below P6/6+: All HV P7: All HV**	**Blades: None P5/5+: All P6/6+: All P7: All**	**Blades: P6 & P7 P5/5+: None P6/6+: All P7: All**

HMC / IVM / SDMC Comparison – General

General Characteristics	IVM	HMC	SDMC
Multiple system support	One IVM per server	One HMC can manage multiple servers (48 cecs w/ 1000 lpars spread amongst)	SDMC can manage multiple servers (48 cecs w/ 1000 lpars Director can manage additional non-Power server entities (OS's, etc)
User Interface	Web browser (no local graphical display)	Web browser (local or remote)	Web browser (local or remote)
Scripting and Automation	VIOS command line interface (CLI) and HMC compatible CLI	HMC command line interface	Director command line interface (compatible with HMC & IVM)

Transition from HMC to SDMC

Side-by-side Management
- Servers will support 1 HMC + 1 SDMC
  - Some P7 servers may support more than 2 management consoles in the future
- SDMC & HMC must have common code levels – same requirement as redundant HMCs today (e.g., 730 SDMC would match up with 730 HMC).

Transition of configuration data
- All partition/profile related info is stored on the managed system itself.
- Transition tool to help bring over static IP managed servers
- Future enhancement for bringing over custom users, groups and roles

2011 April – Selected Hardware Highlights

- New POWER7 Blades: PS703 & PS704
- Refreshed Power 750 & 755
- Higher max LPARs per server (750/770/780/795)
- Expanded use of #2319 Factory Deconfig for Entry Systems
- New Firmware 7.3 functionality
- HDD update
- New SFF I/O drawer
- New SAS-bay-based SSD
- Existing SSD update
- Enhanced 7016-1U2 Multi-Media Enclosure
- New 795 DVD/Tape Storage/Media drawer
- SODs
- Systems Director Management Console

Announce: April 12
Planned availability:
- May 20

eConfig/ordering support:
- April 12 for most things
- May 10 IBM/BNT Switches

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.)

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

Subject to change until announced

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/hmmerAltivecGen 2mod

Revised January 15, 2008

Subject to change until announced

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