



What Is New and Improved With DS8870?

Scott Cardinell
IBM

Wednesday, February 6 2013, 1:30-2:30PM
Golden Gate 7, Lobby Level
Session Number 12770

scardine@us.ibm.com



What Is New and Improved With DS8870?



- Introduction
- Performance
- Hardware
- New Function



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

What Is New and Improved With DS8870?



Introduction:



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



Introducing the DS8870

Smarter Storage for enterprise critical information environments

➤ Built on the DS8800 base

- Exceptionally fast with up to 3x performance increase
- Proven architecture and code base for optimal reliability with non-disruptive microcode updates
- Inherited all functionality of DS8800
- RoHS compliance reduces hazardous material

➤ New dual IBM POWER7 controllers

- Scalable processor configurations with 2, 4, 8 and 16 cores per controller
- Scalable cache from 16GB – 1TB
- Everything scales non-disruptively
- Entry-level Business Class configuration also available


➤ New energy-efficient power supply

- Improved efficiency, power dissipation, reliability
- Up to 20% reduction in energy usage
- Designed to meet evolving energy efficiency standards

➤ Full Disk Encryption drives now standard

- Client decides when to encrypt or not *

* Requires deployment of Tivoli Key Lifecycle Manager or IBM Security Key Lifecycle Manager



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Investment protection for existing DS8000 clients



New DS8870 preserves client investments in prior models

- Equivalent functionality across models
- Existing tools and scripts are compatible
- Remote mirror and copy functions are interoperable
- MES upgrade option will be available to enable clients to convert DS8800 to new DS8870 model (2Q13)

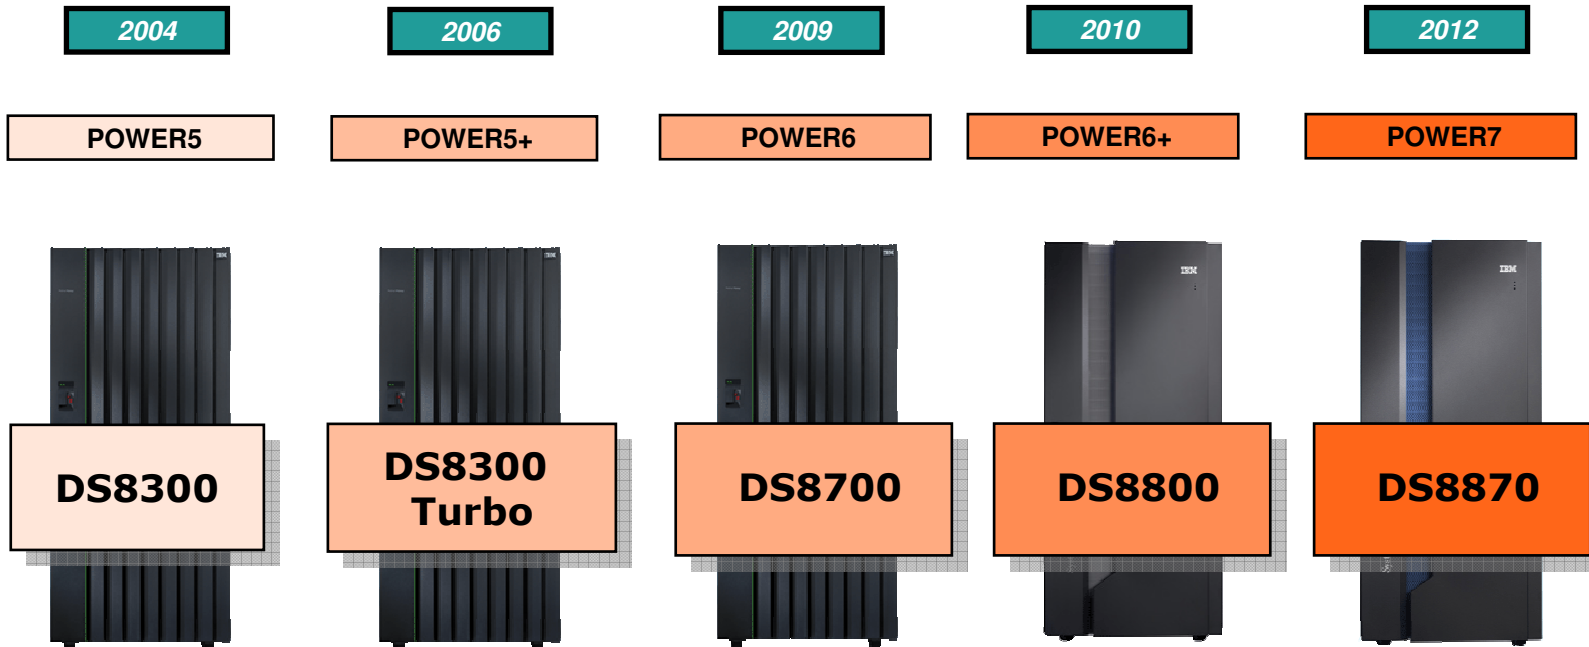


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

5th-generation DS8000 enterprise disk system



- Building on a market-proven, reliable code base!
- 94% of the same proven microcode



- ***Designed for over 5-9's availability natively***
- ***Designed for over 6-9's availability when DS8000 with Metro Mirror is combined with GDPS/PPRC HyperSwap***

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

New hardware delivers up to 3x performance boost



- IBM POWER7-based controllers
- Processor cores that scale from 2-16 per controller
 - Cores aligned with cache configuration
- Cache that scales from 16-1024 GB
- Redesigned power supply for higher energy efficiency
- Dedicated space for new Ultra SSD I/O Drawer (SoD in June)

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Key DS8000 Dates

- **DS8870 R7.0 (242x 961 / 96E)**
 - Announcement date - October 3, 2012
 - General Availability date - October 19, 2012

- **Statement of Direction (SOD)**
 - Field model conversion from DS8800 to DS8870
 - To be available in the first half of 2013
 - Ultra SSD Drawer available 2013

- **Announcing withdrawal for DS8800 Base Frame (242x 951)**
 - Due to new RoHS 2013 compliancy
 - Effective date: 1st February 2013
 - Latest CRAD date: 29th March 2013
 - MES capabilities will continue for existing DS8800s in the field

What Is New and Improved With DS8870?



Performance:




Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



Extraordinary Performance for Enterprise Applications



	DS8300 (R4.3) P5+ 4- core	DS8700 (R5) P6 4- core	DS8800 (R6) P6+ 4- core	DS8870 (R7) P7 16-core	Increase vs. DS8300	Increase vs. DS8700	Increase vs. DS8800
Seq. Read (GB/s)	3.9	9.7	11.8	21.0	5.4x	2.2x	1.8x
Seq. Write (GB/s)	2.2	4.7	6.7	11.0	5.0x	2.3x	1.6x
DB z/OS (K IOPS)	165	201	204	640	3.9x	3.2x	3.1x
DB Open (K IOPS)	165	191	198	550	3.3x	2.9x	2.8x



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Official Storage Performance Council (SPC) results



- **SPC-1 throughput of 451,082 IOPS**
 - #1 result for single, enterprise-class all-HDD system
 - 67% faster than HDS VSP
- **SPC-2 throughput of 15,424 MB/s**
 - 17% faster than HDS VSP
 - 59% faster than DS8800

Source: SPC Benchmark 1 (SPC-1) results page - http://www.storageperformance.org/results/benchmark_results_spc1
SPC Benchmark 2 (SPC-2) results page - http://www.storageperformance.org/results/benchmark_results_spc2


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



SPC-1 comparison vs competition*

Higher throughput with 451K IOPS and lower response time!

SPC-1 Results



DS8300 Turbo: 480 x 73GB 15K hdds, RAID10, 256 GB Cache
HDS USP V: 1024 x 146GB 15K hdds, RAID1, 256 GB Cache
HDS VSP: 1152 x 146GB 15K hdds, RAID1, 512GB Cache+512GB Cache Flash
HP 3PAR V800: 1920 x 300GB 15K hdds, RAID1, 8nodes, 768GB Cache
IBM SVC 6.2 + V7000: 1920 x 146GB 15K hdds, RAID10, 448 GB Cache (8 SVC nodes@24GB + 16 V7000 nodes@16GB)
DS8870: 1536 x 146GB 15K hdds, RAID10, 32x 8Gb FCP, 1024GB Cache

Source: http://www.storageperformance.org/results/benchmark_results_spc1/#spc1

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

What Is New and Improved With DS8870?



Hardware:




Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



Hardware Consistency – DS8800 to DS8870

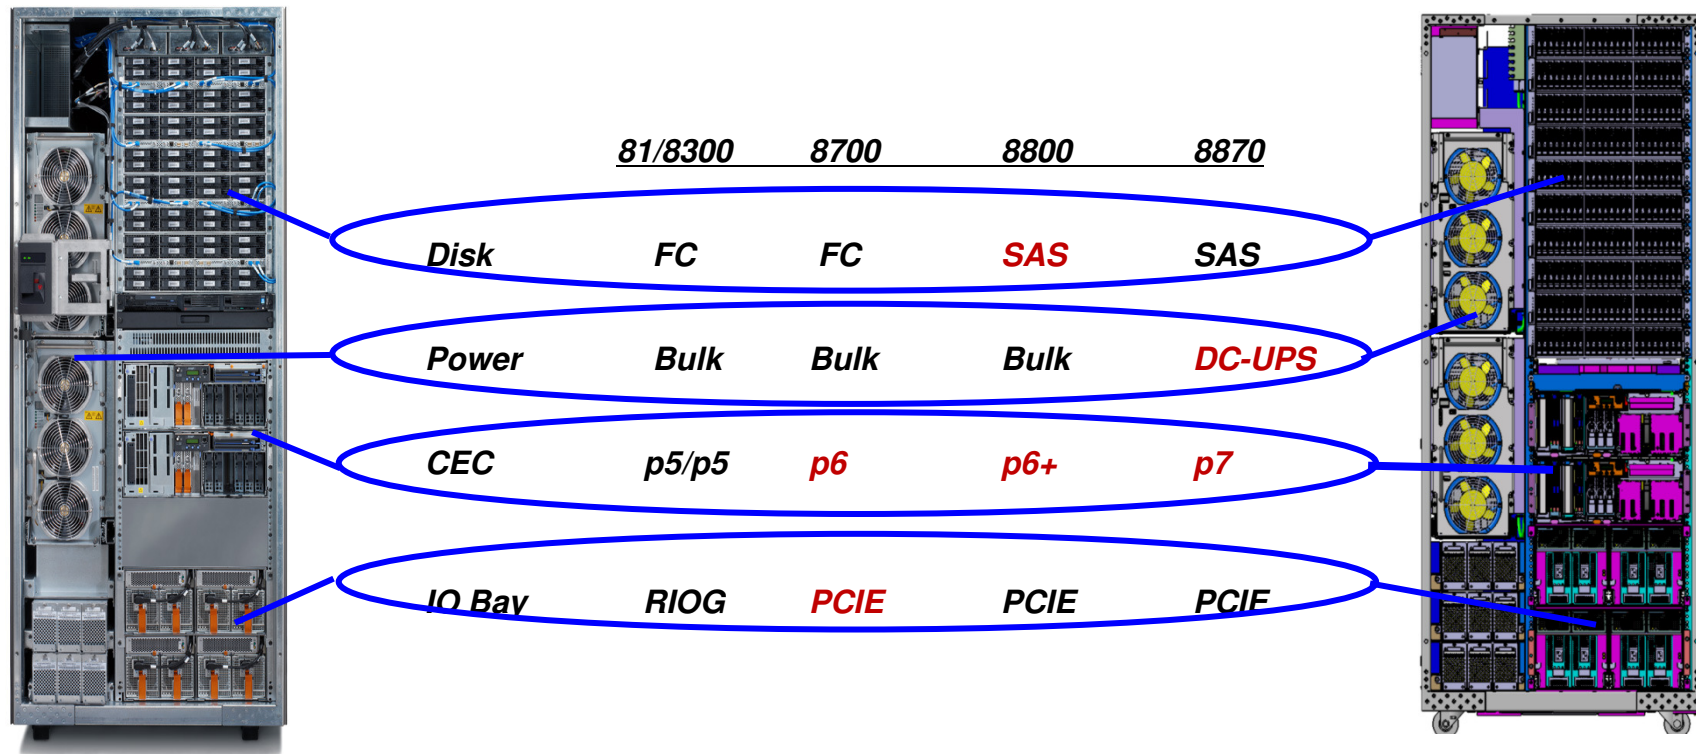


- **DA to DDM mapping / RAID / Intermix / Sparing**
 - Unchanged from DS8800
- **Front to back airflow**
 - Unchanged from DS8800
- **Rack dimensions**
 - Footprint same as DS8800
 - Same number of drives (1536)
- **HA / DA Adapter**
 - Cards unchanged from DS8800
 - Provide equivalent performance per port and per adapter as DS8800
 - Port naming unchanged
 - Plug order unchanged



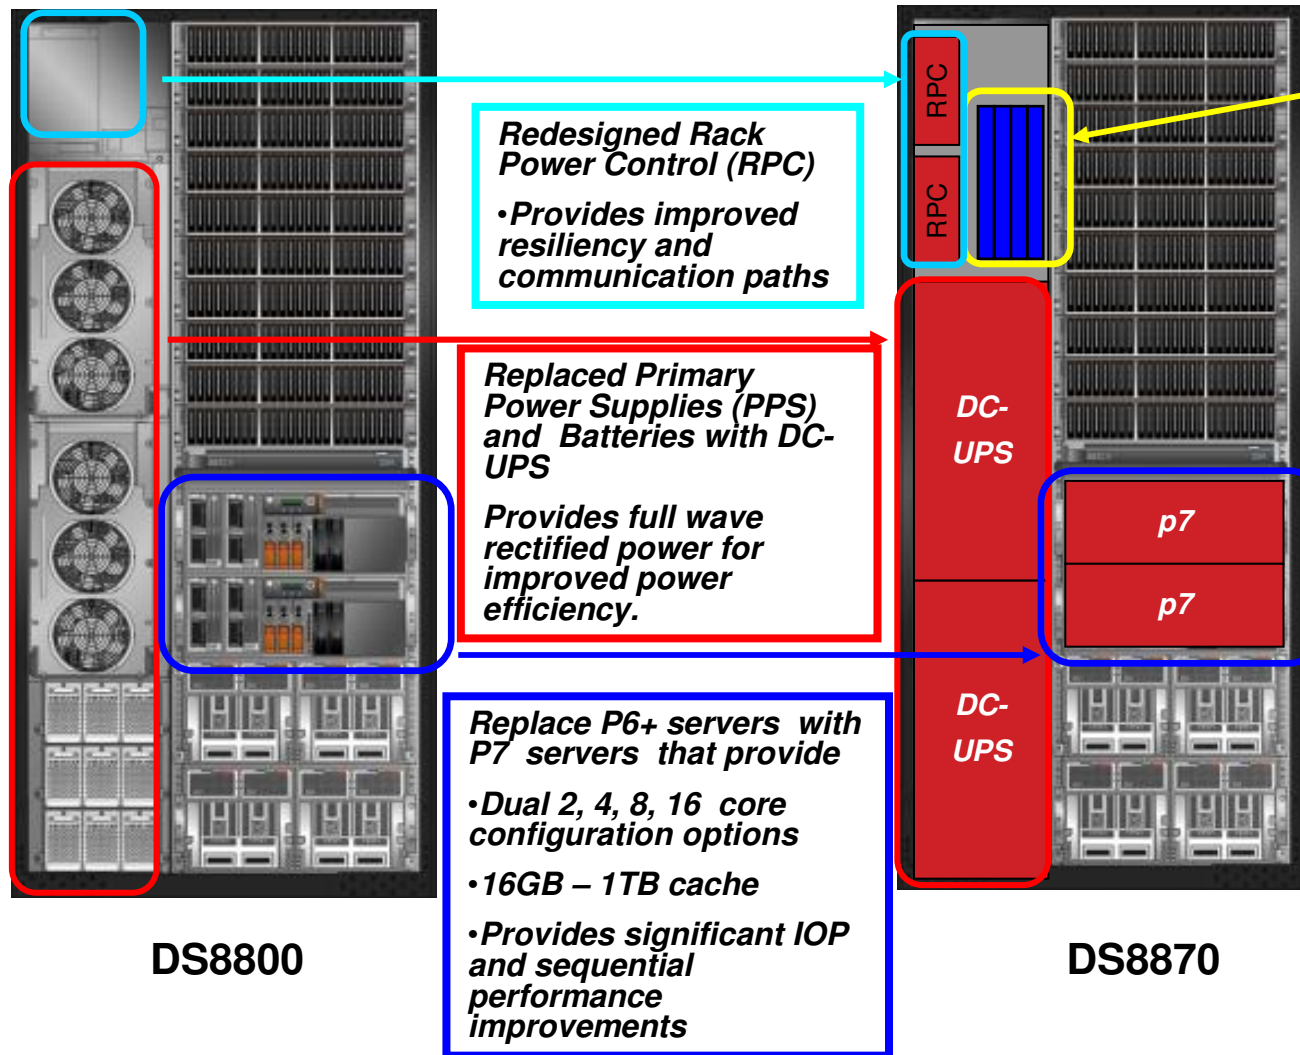
FRONT

DS8000 Enterprise Disk – Hardware Evolution



Incremental changes between versions maximizes quality


Base Frame Hardware Changes – DS8800 to DS8870



June 4 SOD:

- Room for Ultra SSD drawers by rotating 90deg and connecting directly into available PCIe slots
- Will provide substantial additional performance improvement

DS8000 hardware comparison



	DS8300	DS8700	DS8800	DS8870
Processor	P5+ 2.2GHz 4-core	P6 4.7Ghz 2 or 4-core	P6+ 5.0GHz 2 or 4-core	P7 3.55Ghz 2,4,8,16-core
Processor Memory	32 - 256GB	32 - 384GB	16 - 384GB	16 - 1,024GB
Drive Count	16-1,024	16-1,024	16-1,536	16-1,536
Enterprise Drive Options	FC – 73, 146, 300, 450 GB	FC – 300, 450, 600 GB	SAS2 - 146, 300, 450, 600, 900 GB	SAS2 - 146, 300, 600, 900 GB
SSD Drive Options	73, 146 GB	600 GB	300, 400 GB	400 GB
Nearline Drive Options	1 TB	2 TB	3 TB	3 TB
Drive Enclosure	Megapack	Megapack	High-density, high-efficiency Gigapack	High-density, high-efficiency Gigapack
Max Physical Capacity	1,024	2,048 TB	2,304 TB	2,304 TB
Power Supply	Bulk	Bulk	Bulk	DC-UPS
Rack Space for SSD Ultra Drawer	No	No	No	Yes
RAID Options	RAID 5, 6, 10	RAID 5, 6, 10	RAID 5, 6, 10	RAID 5, 6, 10
Internal Fabric	RIO-G	PCI-E	PCI-E	PCI-E
Max Number of LUNs / CKD volumes	64K total	64K total	64K total	64K total
Max LUN Size	2 TB	16 TB	16 TB	16 TB
Host Adapters	ESCON x 2 ports 4 Gb FC x 4 ports	4 Gb FC x 4 ports 8 Gb FC x 4 ports	8 Gb FC x 4 or 8 ports per adapter	8 Gb FC x 4 or 8 ports per adapter
Host Adapter Slots	32	32	16	16
Max Host Adapter Ports	128	128	128	128
Drive Interface	2Gbps FC-AL	2Gbps FC-AL	6Gbps SAS-2	6Gbps SAS-2
Device Adapter Slots	16	16	16	16
Cabinet Design	Top Exhaust	Top Exhaust	Front-to-back	Front-to-back

Information online at SHARE.org/SanFrancisco



DS8000 key feature comparison



	DS8300	DS8700	DS8800	DS8870
Self-encrypting drives	Optional	Optional	Optional	Standard
Point-in-time copies	FlashCopy, FlashCopy SE	Same plus FlashCopy Manager, Remote Pair FlashCopy	Same	Same
Smart Drive Rebuild	No	Yes	Yes	Yes
Remote Mirroring	Advanced mirroring	Same plus - Global Mirror Multi Session - Open HyperSwap for AIX	Same	Same
Automated Drive Tiering	No	Easy Tier Gen. 1,2,3	Easy Tier Gen 1,2,3,4	Same
Thin Provisioning	Yes	Yes	Yes	Yes
Storage Pool Striping	Yes	Yes	Yes	Yes
I/O Priority Manager	No	Yes	Yes	Yes
GUI	DS Storage Manager	Same plus enhancements	New XIV-like GUI	Same
Dynamic Provisioning	Add / Del	Same plus depopulate rank	Same	Same

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Business Class and Enterprise Class configuration options

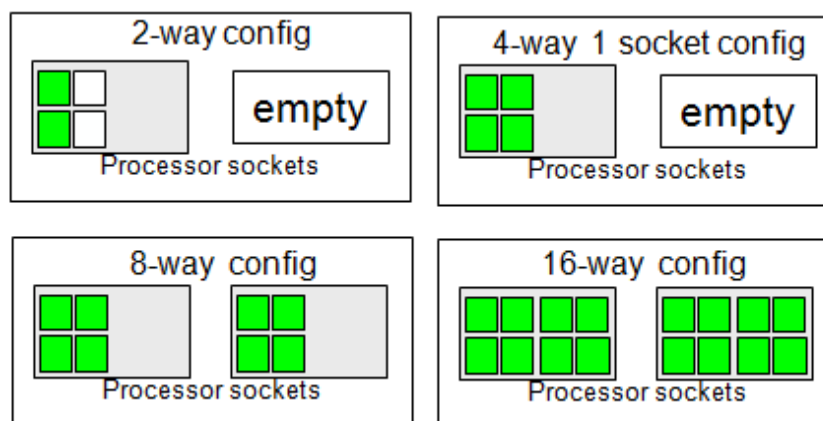


- List price for **Business Class** configuration option is roughly the same as DS8800
 - 30% price premium for Business Class upgrade to Enterprise Class
 - Small servers can be mounted in the Business Class frame to optimize space efficiency
- **Enterprise Class** offers better price/performance than previous models
- **Non-disruptive upgrade** from smallest to largest configuration

Model	Processor	Physical Capacity (max.)	Disk Drives (max.)	Memory	Host Adapters (max.)	9xE Attach
Business Class						
961	2-core	216 TB	144	16/32	4	0
Enterprise Class						
961	4-core	360 TB	240	64	8	0
961	8-core	2,304 TB	1536	128/256	16	0-3
961	16-core	2,304 TB	1536	512/1024	16	0-3
First Expansion Frame						
96E	N/A	504 TB	336	N/A	8	N/A
Second/Third Expansion Frame						
96E	N/A	720 TB	480	N/A	N/A	N/A

DS8870 Processor Configurations (Power 7)

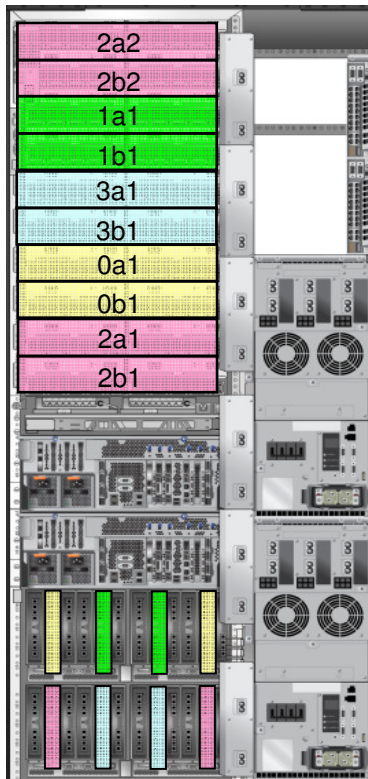
System Class	Active Processor Configuration per CEC	Processor Cards per CEC	Feature Code	Total Active Cores per DS8870	Total System Memory / Persistent	Expansion Frames Supported
Business Class	2 cores	1 x 4 core	#4401	4 cores	16 GB/1 GB	None
					32 GB/1 GB	
Enterprise Class	4 cores	1 x 4 core	#4402	8 cores	64 GB/2 GB	None
	8 cores	2 x 4 core	#4403	16 cores	128 GB/4 GB	0 – 2
					256 GB/8 GB	0 – 3
	16 cores	2 x 8 core	#4404	32 cores	512 GB/16 GB	0 – 3
					1024 GB /32 GB	0 – 3



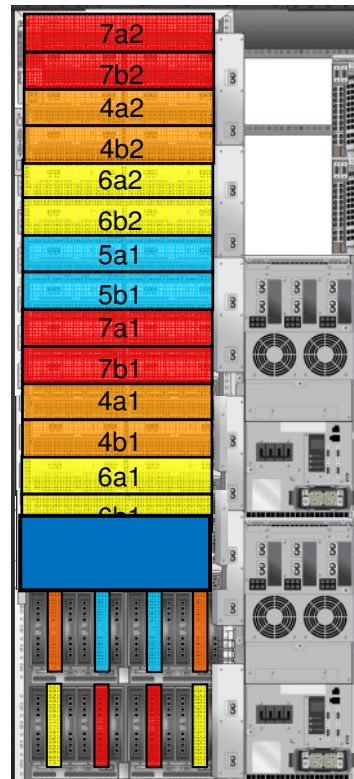
DS8870 Enterprise Class enclosure layout (no change from DS8800)



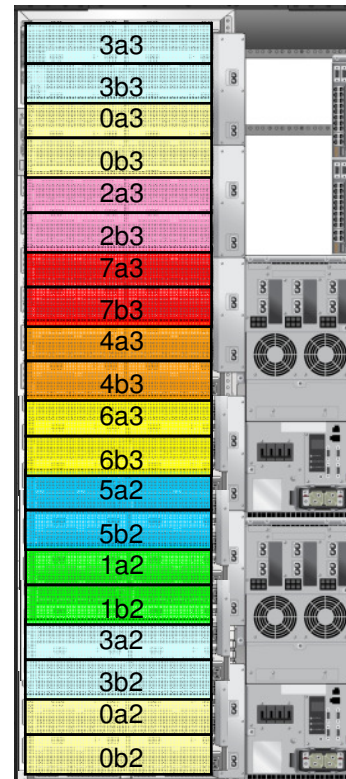
240 drives



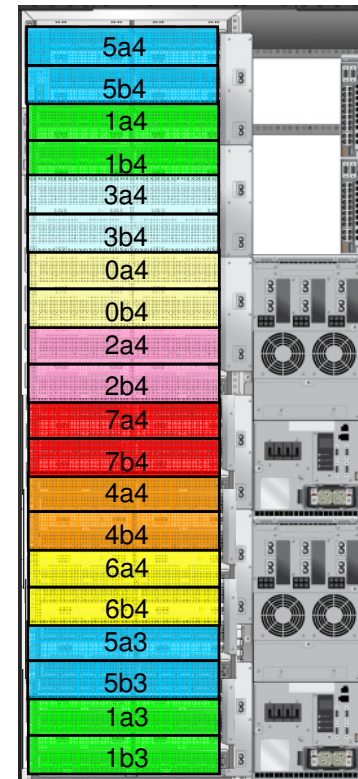
336 drives



480 drives




480 drives



DS8870 Business Class maximum Configuration (reduced drive support from DS8800, upgrades to Enterprise Class)

144 drives



- 2 active processor cores per CEC
- 16GB processor memory
 - No copy services support
 - No I/O priority manager support
- 32GB processor memory
 - support copy services and IOPM
- Concurrent upgrade to Enterprise Class
- SSDs supported on all models
- Easy Tier is supported on all models

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Energy consumption comparison



DS8700 with 1024 drives

- Base frame: 6.8kW
- Exp frame: 7.1kW
- Exp frame: 6.1kW
- Exp frame: 6.1kW
- Exp frame: 3.1kW

TOTAL: 29.2kW



DS8800 with 1536 drives

- Base frame: 7.5kW
- Exp frame: 6.2kW
- Exp frame: 6.3kW
- Exp frame: 6.3kW

TOTAL: 26.3kW



DS8870 with 1536 drives

- Base frame: 6.0kW
- Exp frame: 5.6kW
- Exp frame: 5.8kW
- Exp frame: 5.8kW

TOTAL:

23.2 kW

Base frame is 20% more energy efficient



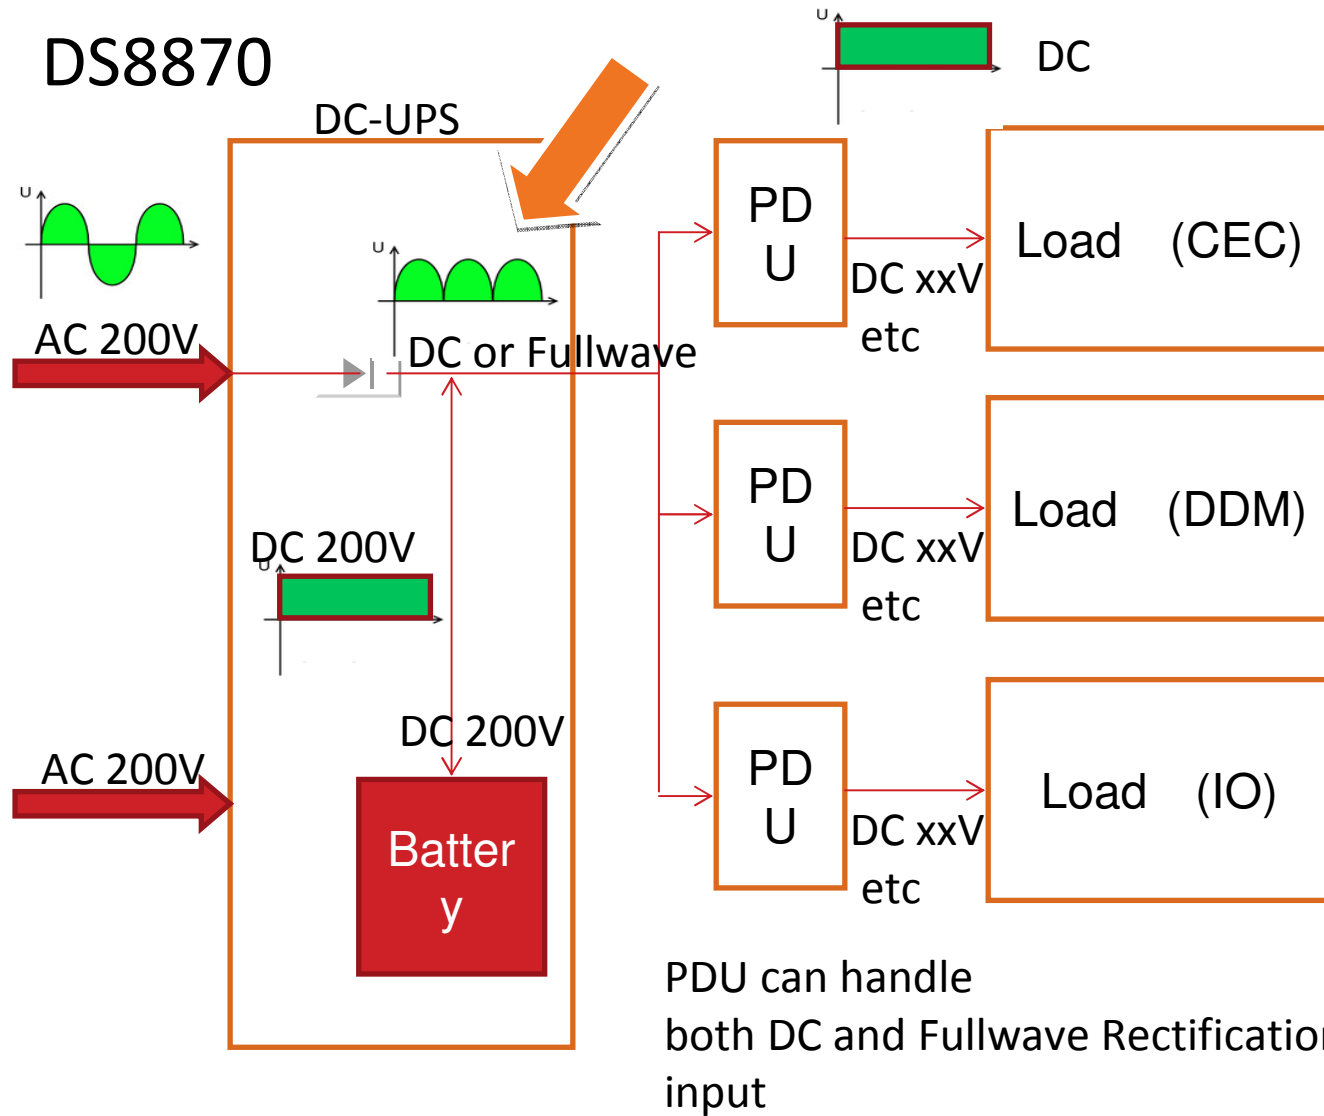
Higher efficiency rating positioned to meet emerging energy efficiency standards

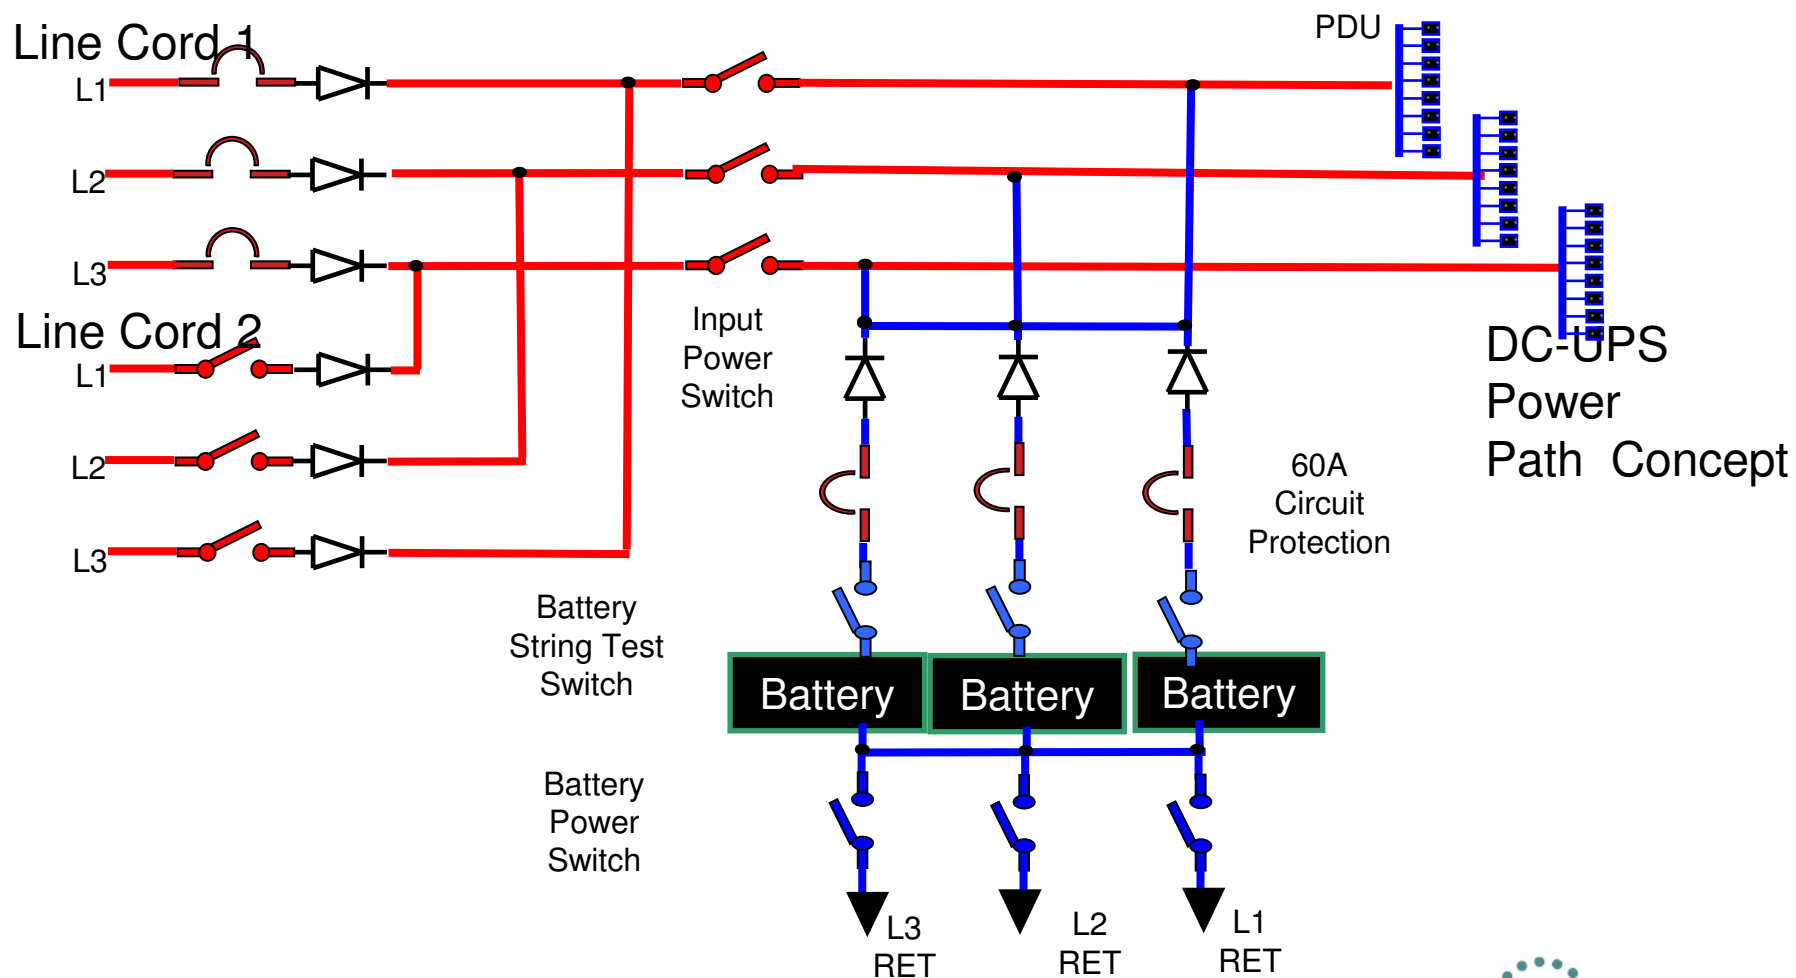
Note: Measurements taken on 100% read miss workload

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DS8870: Why DC-UPS


- Higher efficiency: **From 85% with PPS to 98% with DC-UPS. Higher efficiency needed to compete for Energy Star in 201?.**
- Substantially greater reliability: **Fewer parts, slower parts, cooler parts.**
- Opportunity to simplify/improve code: **Autonomous decision making, individual batteries per DC-UPS, improved code loading, overall requirements simplification.**





RAS Updates / Enhancements

- New DC-UPS vs Old PPS
 - Cabling Topology Different for Load Balancing
 - DC-UPS to PDUs
 - PDUs to IO-OS and CEC-PS
 - PDU to GP-PS
- Multi-Element, Heavy Battery
 - Battery Service Module (BSM)
 - BSM set = 4 BSMs
 - All replaced at once
- Extended Power line disturbance
 - Protects storage unit up to 50 seconds




Best Practice: Hot Aisle / Cold Aisle

Data centers are moving to hot aisle / cold aisle designs to optimize energy efficiency



DS8800 and DS8870 are designed with complete front-to-back airflow.



Cold air pumped beneath raised floor, up through perforated tiles.

Complete your session evaluation online at SHARE.org/SanFranciscoEval

RAS Updates / Enhancements

- New HMC Write-able Media
 - DVD-RAM Media is going End-Of-Life
 - Switch to SDHC Media
 - Used for
 - Data Offload when no remote connection
 - Saving Physical Configuration on Discontinue
 - Not Boot-able
 - Read-only Media & Drive still included



Full Disk Encryption is now standard

- FDE options across all drive tiers
- Same performance as standard drives
- Key manager supports both disk and tape
- New European Union Privacy Directive makes breaches very costly
- Encryption is the least expensive data disposal technique
- **Supports Easy Tier environments**
- Over a thousand IBM encryption disk and tape solutions deployed worldwide

"Do not wait for an event-driven reaction to secure your data. Proactively securing your data will help ensure against a worst-case scenario and a financial impact that is likely to far surpass that of the data security purchase itself."

John Monroe
Gartner Hype Cycle for Storage Technologies,
July 2012



"Within five years, all HDDs and SSDs will be shipped preloaded with some kind of industry-standard FDE technology" – Gartner Hype Cycle for Storage Technologies, July 2012

* Requires deployment of Tivoli Key Lifecycle Manager or IBM Security Key Lifecycle Manager
tion online at SHARE.org/SanFranciscoEval

DS8870 is RoHS compliant



- **Reducing hazardous materials for a greener data center**
 - Meets 2013 European Union requirements for restricting hazardous materials in system and manufacturing
 - Requirements being adopted across geographies

RoHS - Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



DS8000 Integration of EXP30 UltraSSD storage



June 4, 2012 SOD

- 1U drawer
 - Up to 30 encryption capable eMLC SSD drives
 - With 387 GB drives = up to 11.6 TB (raw)
 - Each drawer is independently installable; up to 4 per rack
- Two 6 Gb SAS adapters especially designed for SSD
 - Directly connected to DS8000 internal PCIe fabric via currently open slots
- Enterprise class RAS
 - RAID protected with spares
 - Dual power supplies, sensors, adapters and fans
 - Concurrently replaceable components
- Easy Tier enabled



[IBM June Storage SOD announcement link](#) . [UltraSSD Announcement link](#)

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval



What Is New and Improved With DS8870?



New Function:



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval




Support for T10 Protection Information standard (formerly Data Integrity Field (DIF))



End-to-end data integrity through the SAN

- Provides advanced, end-to-end data integrity (initial support for Linux on System z)
- Checks data integrity to and from the host bus adapter and the disk through the SAN fabric
- Checking done by hardware, so no performance impact

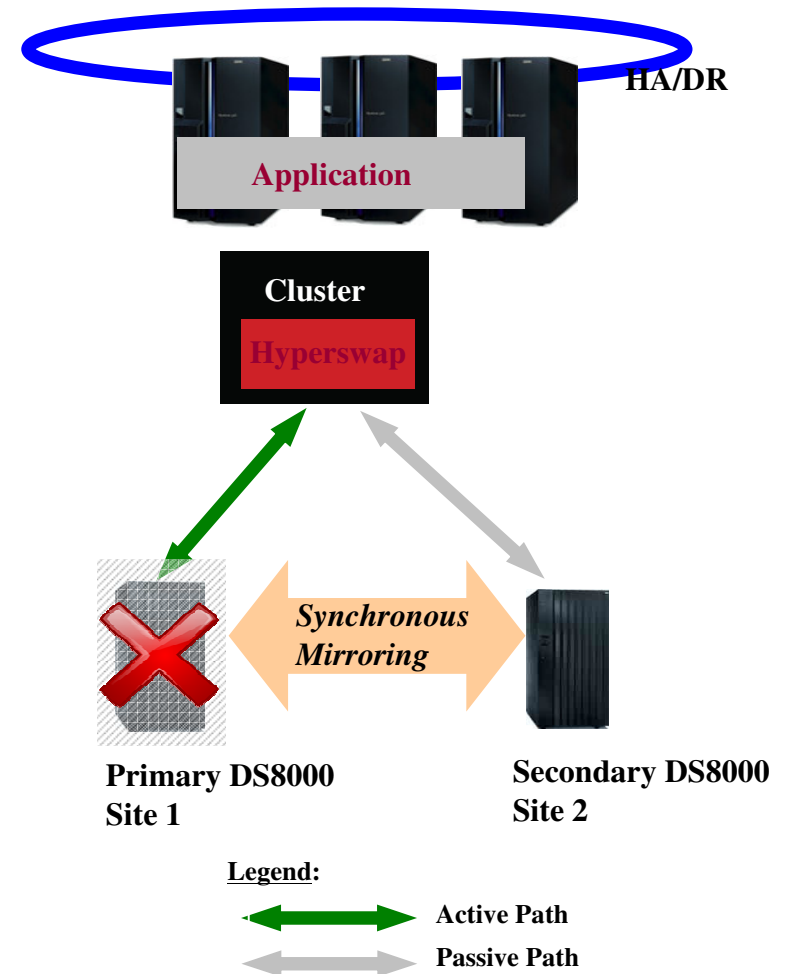


Power HyperSwap Technology

- HyperSwap non-disruptively substitutes secondary for primary devices for planned and unplanned events
 - Transparent to applications - continue to use same the same device
- Customer Benefits
 - Unplanned HyperSwap
 - Continuous availability against storage failures
 - Planned HyperSwap
 - Storage migrations without downtime
 - Storage maintenance without downtime
- Requirements
 - AIX 6.1 TL8 with SP1 or AIX 7.1 TL2 with SP1
 - PowerHA 7.1.2 SP1 with APAR IV2758
 - Available November 9, 2012
 - DS8700/DS8800 R6.3 SP2 (DS8870 in 2013)

Oct 3 PowerHA HyperSwap Announcement

Brings together AIX, PowerHA and DS8000 to provide a comprehensive application and data availability solution



DS8870 and IBM Power AIX and Power i



- **Performance**

- Full Easy Tier support for most server OSes, including Power AIX and Power i
- DB2 End-to-End I/O Priorities and Cooperative Caching for Power AIX
- SOD: Automated tiering to SSD Ultradrawer on Power AIX
- Integrated performance monitoring tools between Power i and DS8000

- **Availability**

- DS8000 copy services support many server platforms, including IBM Power i and Power AIX
- DS8000 copy services integration and automation with PowerHA SystemMirror for AIX and i
- Open HyperSwap for Power AIX (managed through TPC for Replication)
- **New** integration between DS8000 copy services and PowerHA HyperSwap (DS8700/DS8800)
- Full System FlashCopy (FSFC) Toolkit for Power i



Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

New VMware support



- VMware vStorage APIs for Array Integration (VAAI) includes 3 “primitives” for offloading virtual machine and storage management operations to storage systems that support them
 - Atomic Test and Set (ATS) enables hardware-assisted locking of files
 - Full Copy - enables the storage array to make full copies of data within the array
 - Block Zeroing enables the array to zero out large numbers of blocks
- ATS and Full Copy support are available with this DS8870 release; Block Zeroing support is planned for 2013
- DS8870 enhancements for VMware planned for 2013
 - VAAI Block Zeroing
 - vCenter Plug-in
 - SRM 5.0/5.1

What Is New and Improved With DS8870?

- Summary:
 - DS8870 preserves client investments in prior models.
 - Up to 3x performance boost.
 - New power system gives 20% percent better efficiency.
 - 1 TB cache
 - Full Disk Encryption
 - Future Support for UltraSSD Storage
 - T10 Protection Information Standard for End-to-end data integrity.
 - Power Hyper Swap solutions
 - VMware VAAI Support

Many Thanks To:



- Irma Flores-Mendoza
- Chip Jarvis
- Allen Marin
- Hank Sauter
- Vic Peltz
- Alison Pate
- Andy Walls
- Tony Pearson
- Jared Minch
- Brian Sherman
- Paul Spagnolo
- Joanne Brown
- Mark Kremkus
- David Whitworth
- Yan Xu
- Brian Rinaldi

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval





What Is New and Improved With DS8870?

Scott Cardinell
IBM

Wednesday, February 6 2013, 1:30-2:30PM
Golden Gate 7, Lobby Level
Session Number 12770

scardine@us.ibm.com



Trademarks and Disclaimers



IBM Corporation 1994-2012 All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Symmetrix are trademarks of EMC Corporation.

3PAR is a trademark of Hewlett Packard.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

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

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

