


IBM System Storage™

IBM System Storage DS8000 and tape

Large Systems Update 2007

**A New Standard
in Storage
Leadership**



Henrik Grönberg
Advisory Storage Specialist
henrik.gronberg@se.ibm.com

IBM® System Storage™ DS8000 series

Featuring IBM System Storage DS8000 Turbo models



DS8000 Turbo

**New opportunities to help
increase ROI and decrease
long-term costs**

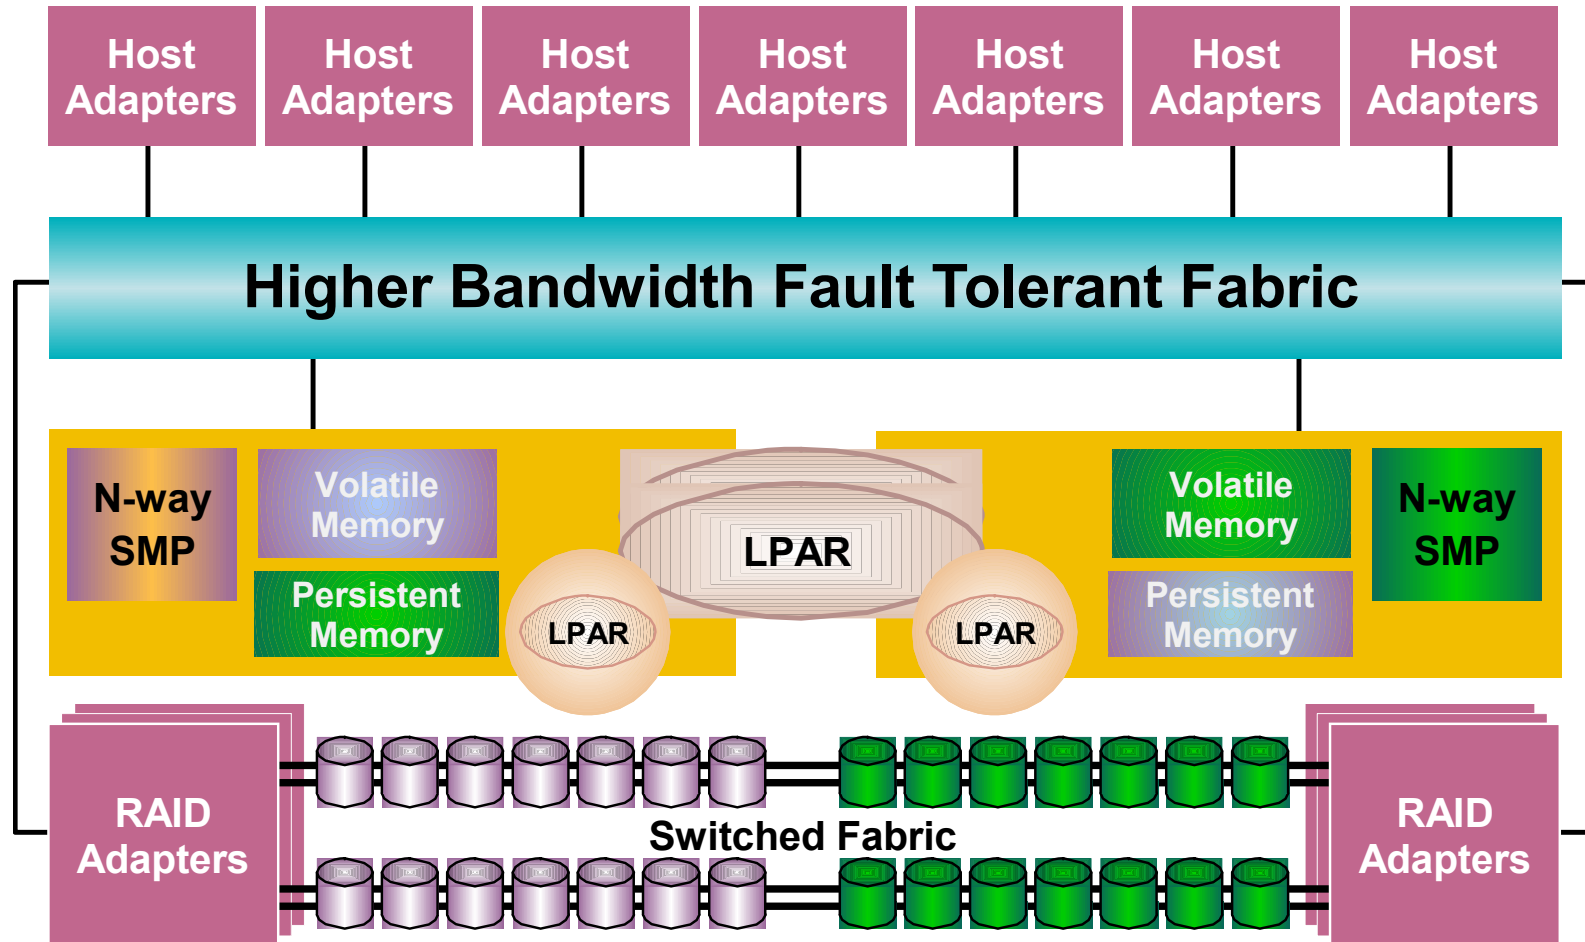
▪ Setting a “New Standard” in Cost Effectiveness

- ▶ **Balanced Performance** – Up to 7X ESS Model 800
- ▶ **Exceptional Scalability** – Up to 512TB physical capacity
- ▶ **Virtualization for Simplification** – Storage System LPARs
- ▶ **Flexibility** – FC and FATA disks to fit access/cost needs
- ▶ **Extendibility** – Dramatic addressing capability
- ▶ **Storage Management** – Full complement of interfaces/tools
- ▶ **Availability** – Designed for 24X7 environments
- ▶ **Resiliency** – Leading Copy and Mirroring Capability
- ▶ **Long Term Cost** – Enterprise Choice Warranty
Model to Model Upgradeability

▪ Delivered through

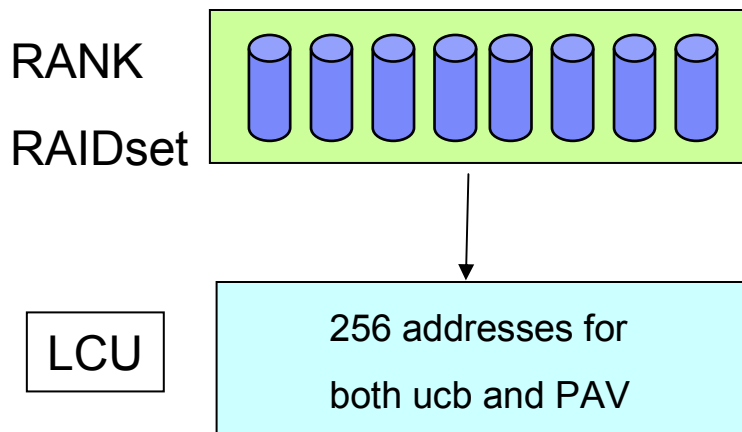
- ▶ Server/Storage Integration – POWER5™ Technology
- ▶ Exploitation of IBM Virtualization Engine™ Technology
- ▶ IBM technology leadership and innovation

Leveraging and Extending IBM's Server-based Architecture



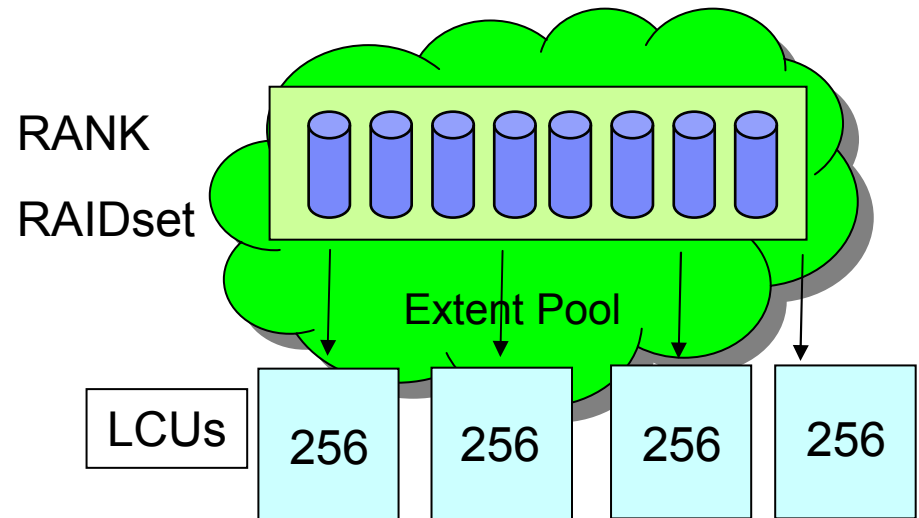
DS8000 virtualization - LCU configuration

ESS



One rank = LSS/LCU = 256 (ucb) addresses
 Serious limitation when using large drives and small CKD volumes.
 145GB drives gave you 840GB usable space
 Configuring 3390-3, not all capacity can be used
 No addresses left for PAV.....

DS8000



One rank/extentpool can be divided in multiple LCUs each with 256 addresses.
 All capacity can be used.
 A lot of addresses available (many PAVs).
 LCU with a "few" GB can be created and assigned to a zSeries LPAR.

DS8000 Turbo Hardware Overview

■ 2-Way (DS8100 Turbo Models)


- ▶ Two dual POWER5+™ based processor servers
- ▶ Up to 128GB Cache
- ▶ 8 to 64 4Gb or 2Gb FC/FICON – 4 to 32 ESCON Ports
- ▶ 16 to 384 HDD
 - Intermixable disk drive packs
 - 73GB, 146GB and 300GB 15k RPM FC
 - 146GB and 300GB 10k RPM FC
 - 500GB 7200 RPM FATA
- ▶ Physical capacity from 1.1TB up to 192TB

■ 4-Way (DS8300 Turbo Models)

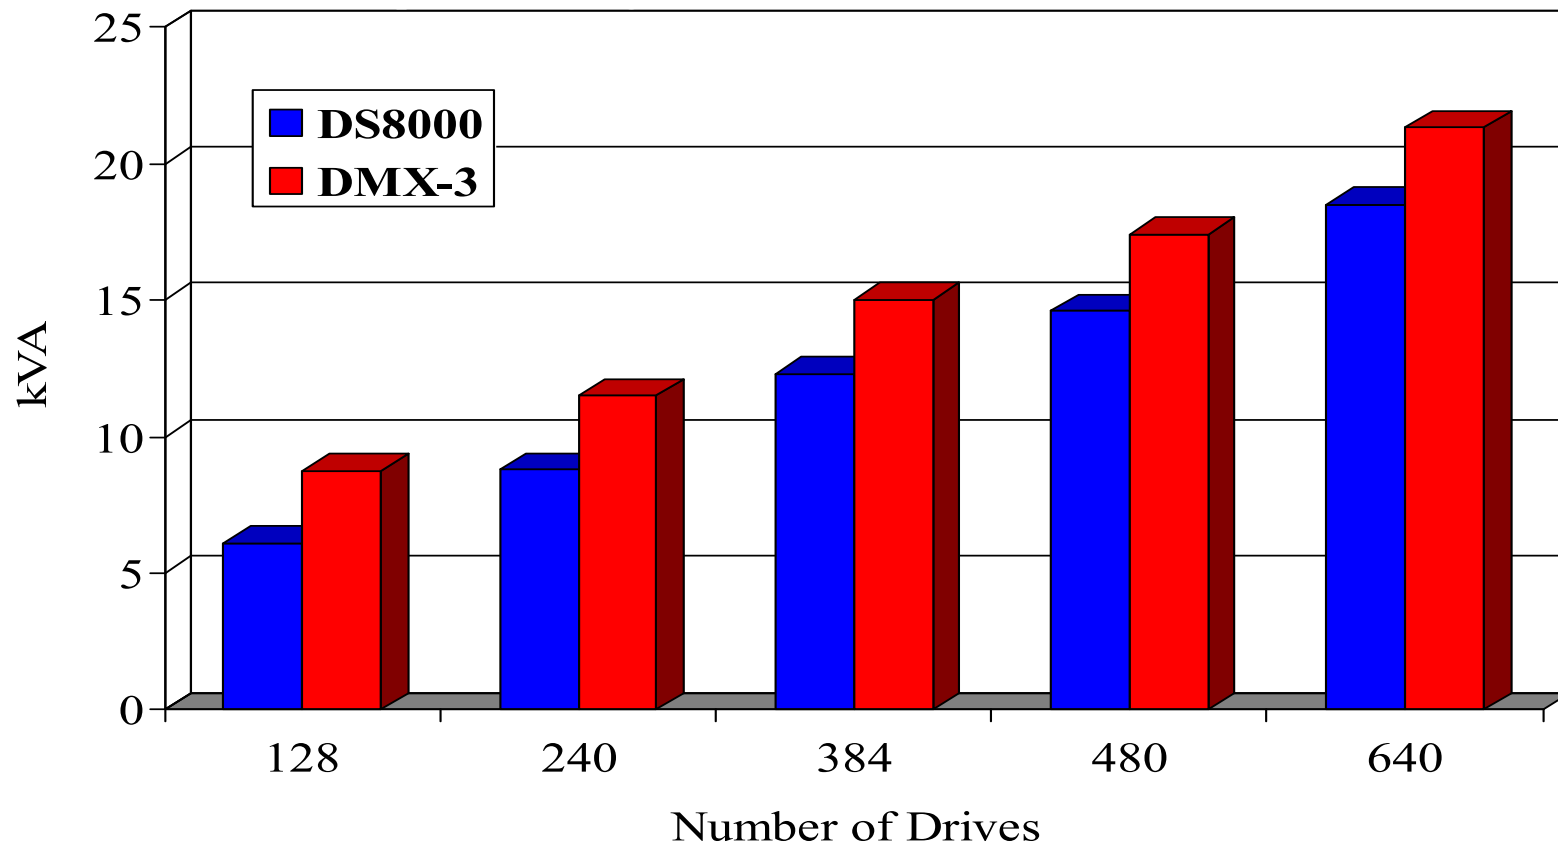
- ▶ Two four processor POWER5+™ based servers
- ▶ Up to 256GB Cache
- ▶ 8 to 128 4Gb or 2Gb FC/FICON – 4 to 64 ESCON Ports
- ▶ 16 to 1024 HDDs
 - Intermixable disk drive packs
 - 73GB, 146GB, and 300GB 15k RPM FC
 - 146GB and 300GB 10k RPM FC
 - 500GB 7200 RPM FATA
- ▶ Physical capacity from 1.1TB up to 512TB



DS8300 Turbo (4-Way with four expansion frames)



Disk System Power Consumption Comparison



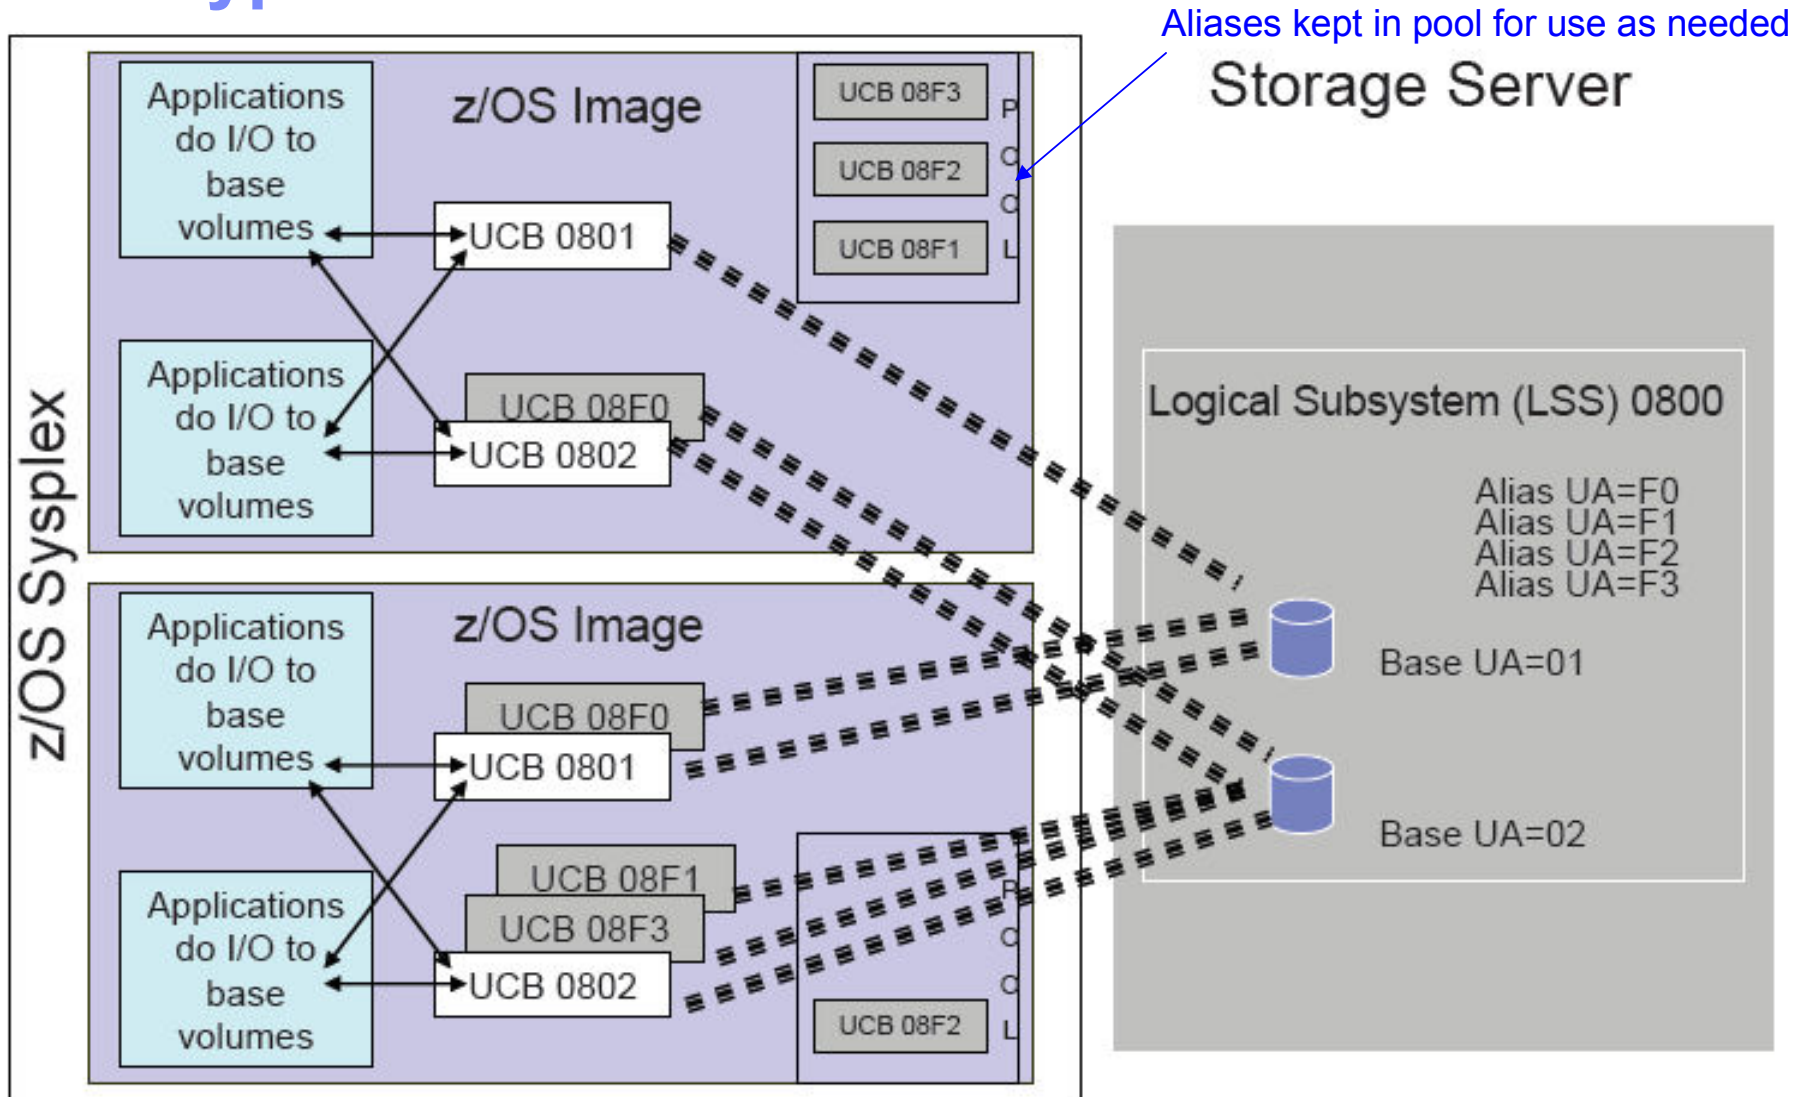
IBM DS8000 can consume up to 31% less power.

Comparisons based on typical power usage ratings using published EMC data and IBM measured power usage.

IBM HyperPAV for Greater Efficiency and Relief from Device Address Constraint for IBM z/OS Environments

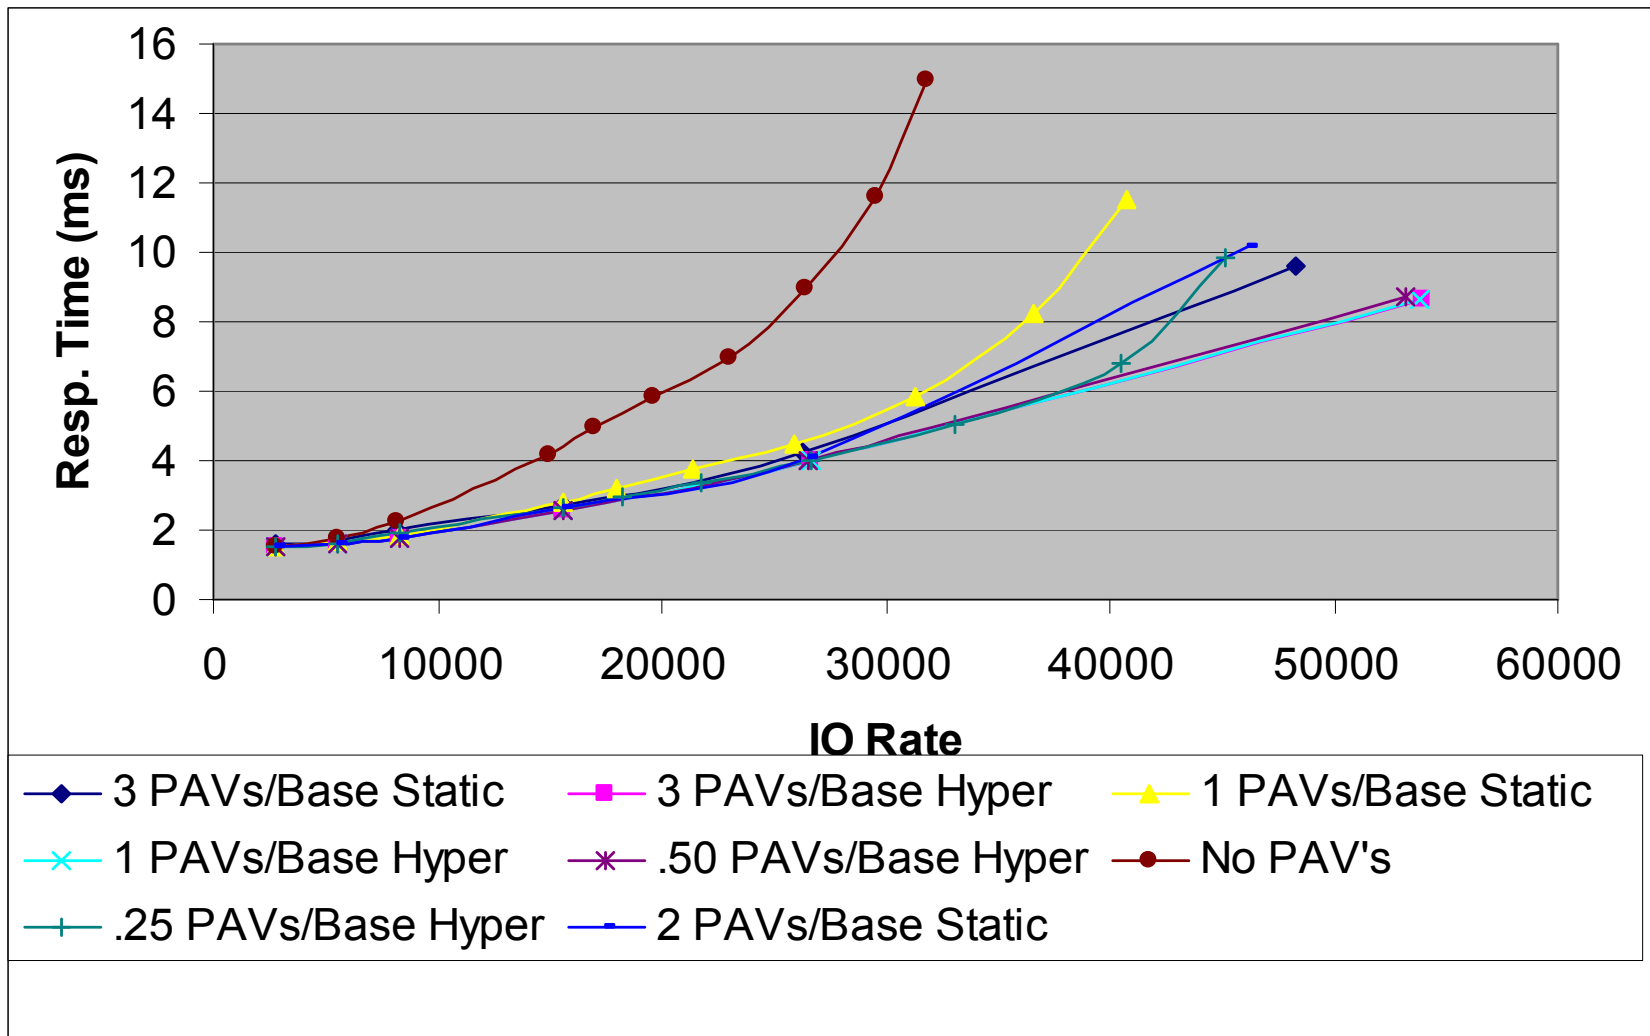
- IBM System Storage™ DS8000 offers unique support for IBM HyperPAV feature designed to:
 - ▶ Provide more efficient Parallel Access Volumes (PAV) function
 - ▶ Help customers who implement larger volumes to scale I/O rates without the need for additional PAV-alias definitions
 - ▶ Can help to reduce overhead, improve addressing efficiencies, and provide storage capacity and performance improvements
 - *More dynamic assignment of PAV-aliases improves efficiency*
 - *Number of PAV-aliases needed may be reduced, taking fewer from the 64k device limitation and leaving more for storage capacity use*
 - ▶ Enable a dynamic response to changing workloads much faster compared to WLM handled PAVs. Aliases assigned on a per IO basis.

IBM HyperPAV



Aliases are assigned on a per IO basis. WLM is not in charge of assigning aliases with HyperPAV

Static PAV vs HyperPAV Cache Friendly



MIDAWs Announcement

The Modified Indirect Address Word (MIDAW) facility is designed to help improve the performance of many applications by reducing FICON channel, director, and control unit overhead. DB2 will benefit the most from MIDAWs, but other examples of datasets that will benefit include VSAM, PDSE, HFS and zFS.

Unlike an IDAW, a MIDAW contains a length field

A MIDAW consists of 8 byte pointer and an 8 byte length

Like IDAWs, MIDAWs are implemented in the SAP


MIDAWs are implemented by Media Manager in z/OS 1.6

What do these datasets have in common? They all use Media Manager and they mostly use 4K blocks. Only Media Manager uses MIDAWs.



Why does DB2 benefit the most? EF datasets benefit the most. DB2 and VSAM require EF for some functions. Not all VSAM applications support those functions, but DB2 does.

Parallel DB2 Table Scan, EF 4K (single channel)

I/O Response Time (sec)



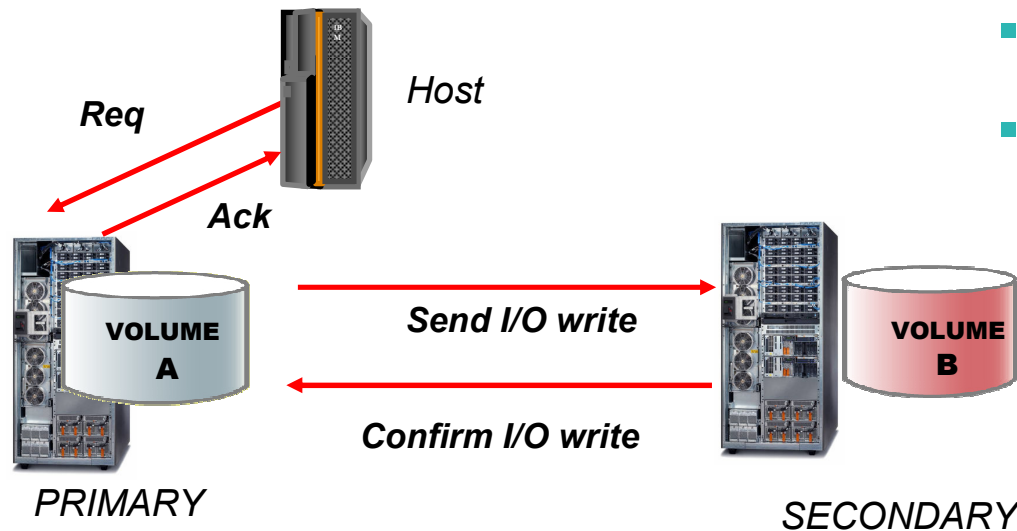
pre-MIDAWs MIDAWs



IBM 2 Site Metro Mirror

Designed to Provide:

- **No Data Loss**
- **Supports:**
 - ▶ Open Systems
 - ▶ iSeries
 - ▶ zSeries
 - ▶ zSeries GDPS Hyperswap
- **Ease of use, Lower cost**
- **Industry Leading Performance**



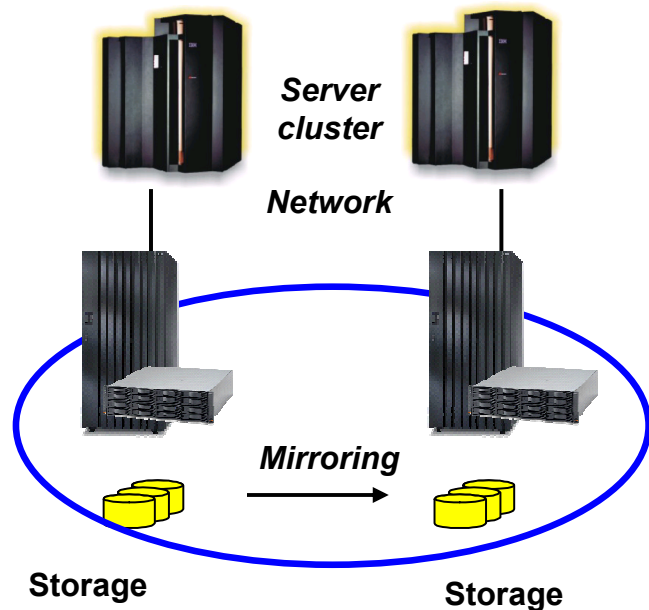
**Approximately
35% of all ESS, DS6000,
DS8000 subsystems have
license for Metro Mirror**

IBM PPRC Migration Manager

z/OS basic disk mirroring Automation Solution



- **zSeries basic disk mirroring management:**
 - *PPRC Migration Manager*
- ▶ **Supports DS6000, DS8000, ESS:**
 - Metro Mirror
 - Global Copy
 - Global Mirror
- ▶ **Considerations:**
 - Supports two site capability only
 - Does not support z/OS (XRC) Global Mirror



Scalable Data Integrity

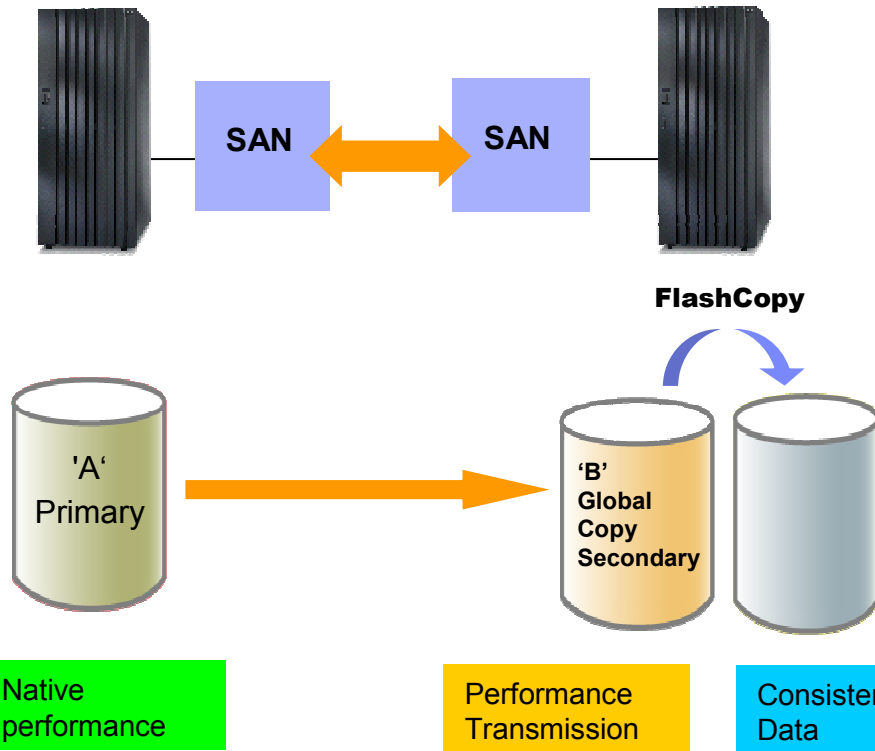
■ **Features:**

- ▶ Provides *basic* management of IBM enterprise disk mirroring for z/OS data
- ▶ Robust configuration auto-discovery and intuitive ease of use
- ▶ Scalability to thousands of volumes
- ▶ Proven in large demanding zSeries customers worldwide
- ▶ When z/OS client strategic plans do not include GDPS, PPRC Migration Manager provides a basic alternative

IBM 2 Site *Global Mirror*

Designed to Provide:

- Unlimited Global Distance
- Reliable Data Currency
 - ▶ 3-5 seconds (bandwidth permitting)
- Scalability
 - ▶ Up to 8 subsystems (17 with IBM RPQ)
- Heterogeneous Consistency Group
 - ▶ zSeries®, iSeries, open systems data
- Native Application Performance
- Ease of use, Lower cost
 - ▶ No active external controlling software or server cycles required to form Con groups




Large European bank using Global Mirror



- Top 50 of Fortune Magazine's list of the world's 500 largest companies**
- 35 TB, 15 ESSs in mirrored config
 - 27,000 IO/sec
 - RPO: 4 to 7 seconds

IBM Metro Mirror and GDPS/PPRC HyperSwap



- **Ability to swap enterprise class zSeries data centers in a matter of seconds**
 - HyperSwap substitutes Metro Mirror **secondary** for **primary** device
 - No operator interaction, GDPS-managed
 - Scales to multi-thousands of z/OS volumes
 - Includes volumes with SYSRES, page data sets, catalogs
 - Non-disruptive - applications keep using same device addresses

Brings different zSeries technologies together to provide a comprehensive application and data availability solution



Benchmark Measurements Planned Disk Reconfiguration: GDPS HyperSwap



Without HyperSwap → **PLANNED ACTION INITIATED** → **1-2 hrs (approx)**
 shutdown systems, remove systems from Sysplex, reverse PPRC (suspend PPRC), restart systems

With HyperSwap No Failover/Failback → **93 Seconds! (2900 vol pairs 4.6 TB)**
 terminate PPRC, swap the primary & secondary PPRC UCBs, systems continue

HyperSwap latest 3.2 release exploits PPRC Failover/Failback → **18 Seconds! (2900 vol pairs 4.6 TB)**
 PPRC Failover, swap the primary & secondary PPRC UCBs, systems continue


Reference Customer	Configuration	Switch Time (without FO/FB)	Switch Time (with FO/FB)
ARZ (Austria)	2300 vol pairs (14 TB) (Note 1)	82-84 secs	
Postbank (Germany)	1800 vol pairs (32 TB)	80-84 secs	
iT Austria (Austria)	4200 vol pairs (24 TB)	75 secs	
iT Austria (Austria)	4500 vol pairs (76 TB)	75 secs	
IBM test facility (MOP)	2900 vol pairs (4.6 TB)	93 secs	18 secs
	Note 1: TB depends on 3390-3, -9, -27 type volumes		



GDPS Solution Examples

High Availability of Data within a Data Center


Single Data Center
Applications remain active
Continuous access to data in the event of a storage subsystem failure



GDPS®/PPRC
HyperSwap™ Manager

High Availability / Disaster Recovery within a Metropolitan Region


Two Data Centers
Systems remain active
Multi-site workloads can withstand site and/or storage failures



GDPS/PPRC

Disaster Recovery at Extended Distance


Two Data Centers
Rapid Systems Disaster Recovery with “seconds” of Data Loss
Disaster recovery for out of region interruptions



GDPS/XRC
GDPS/Global Mirror

High Availability Regionally and Disaster Recovery Extended Distance


Three Data Centers
High availability for site disasters
Disaster recovery for regional disasters



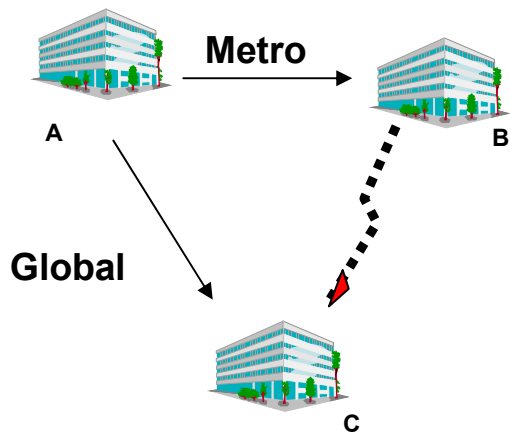
GDPS/Metro/Global Mirror

IBM 3 site Mirroring

Metro/Global Mirror
A->B->C



z/OS Metro/Global Mirror
A->B, A->C



- **Fast Failover / Failback to any site**

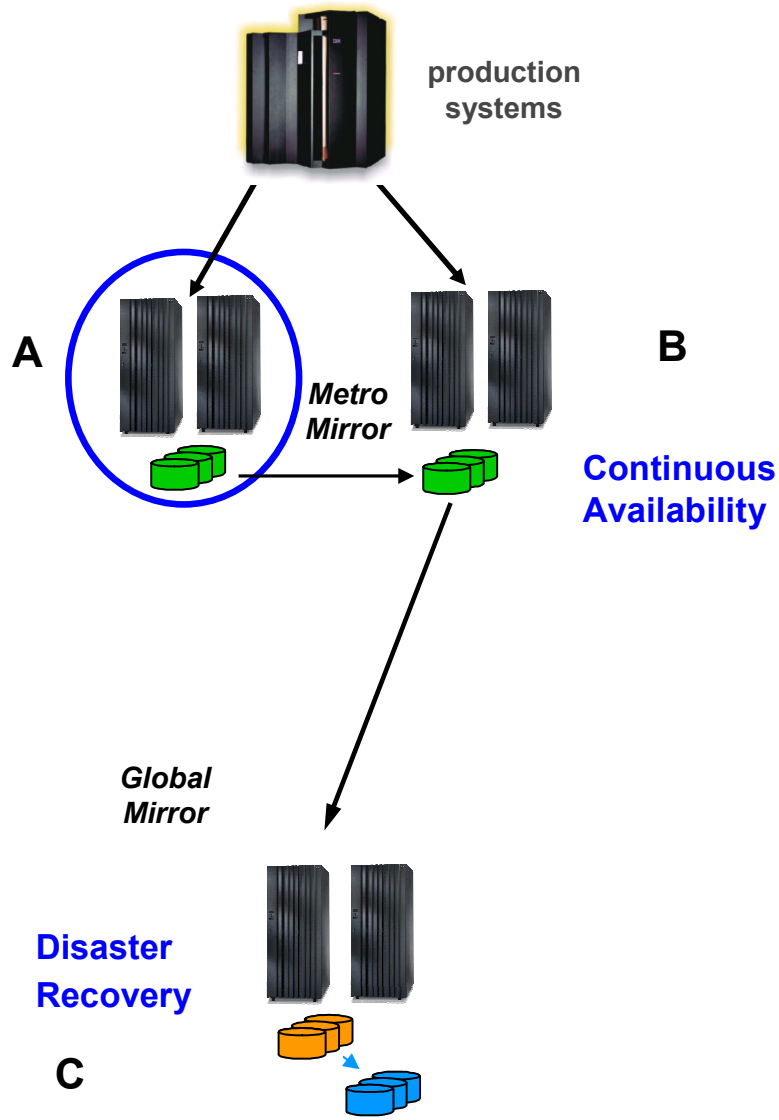
- **Fast re-establishment of 3 site recovery, without production outages**

- **Quickly resynchronize any site with incremental changes only**

- Links and bandwidth assumed between all sites



Three Site IBM *Metro / Global Mirror* Overview



- **Cascaded configuration**

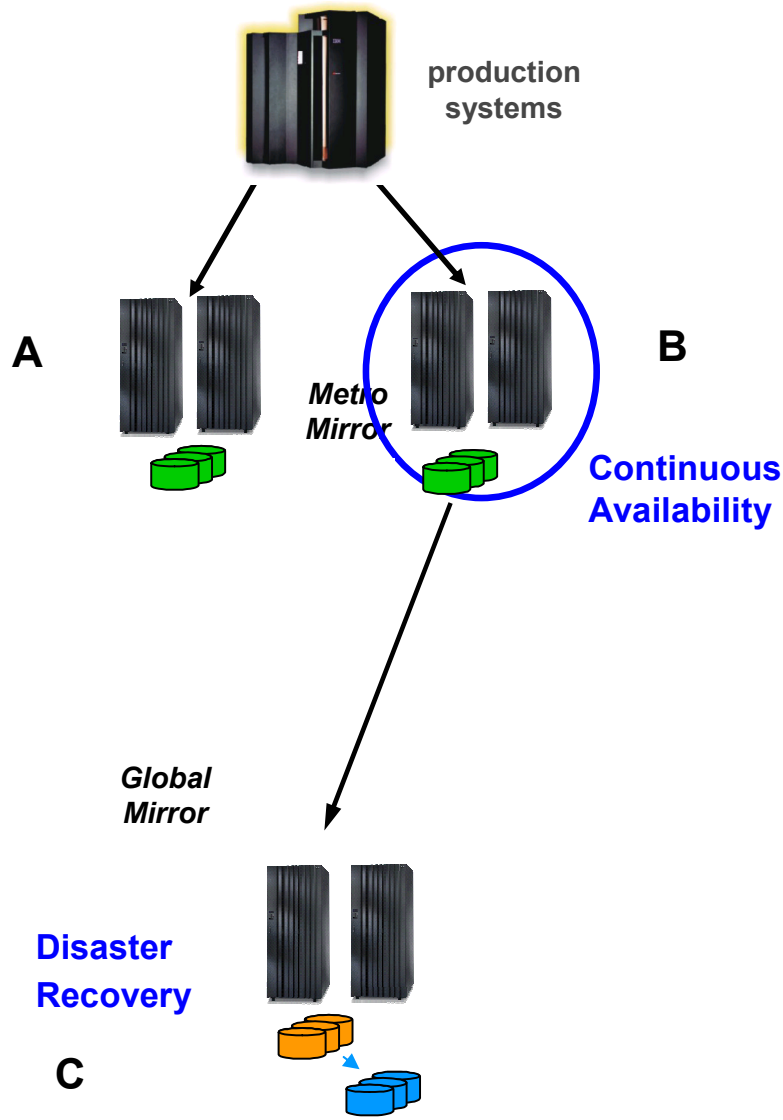
- **Metro Mirror (sync PPRC) and cascaded Global Mirror (async)**

- ▶ **Metro Mirror** provides synchronous mirroring at metro distance from site A to site B

- Load on controller A is reduced compared to Multi-Target

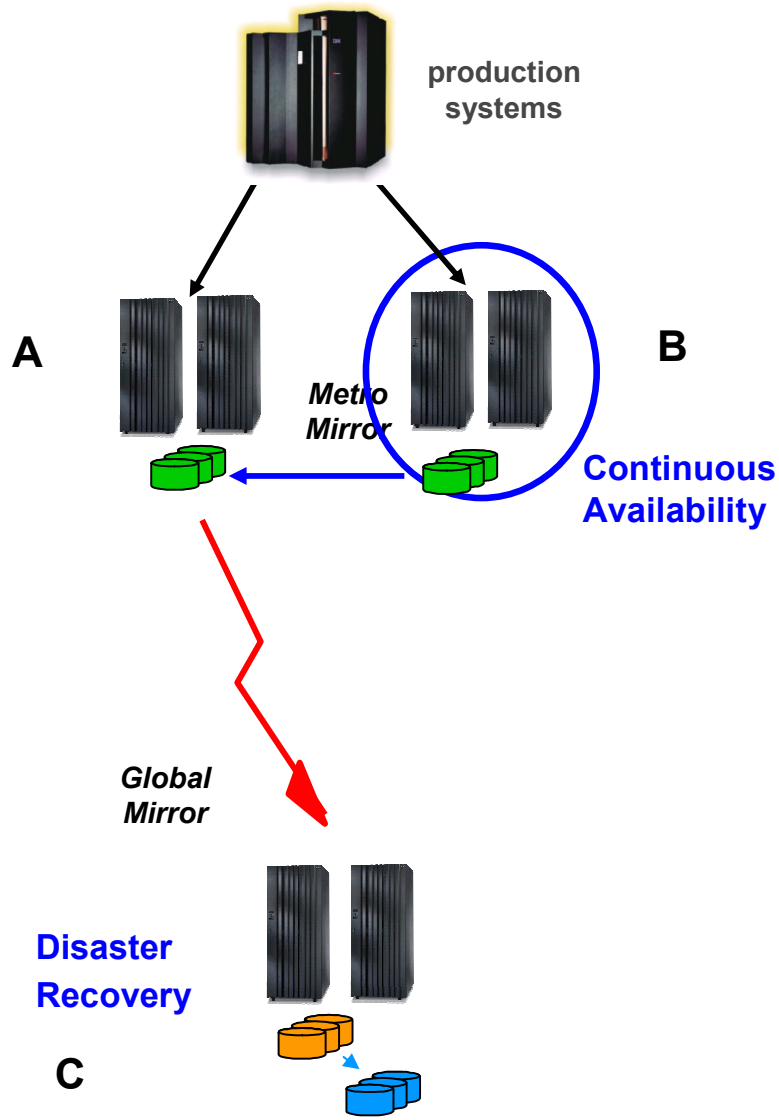
- ▶ **Global Mirror** to provide regional Disaster Recovery at site C

Three Site Metro / Global Mirror Overview



- Metro Mirror (sync PPRC), and cascaded Global Mirror (async)
- Planned or unplanned switch to site B:
 - Continue Disaster Recovery coverage on B-C leg
 - No reconfiguration necessary

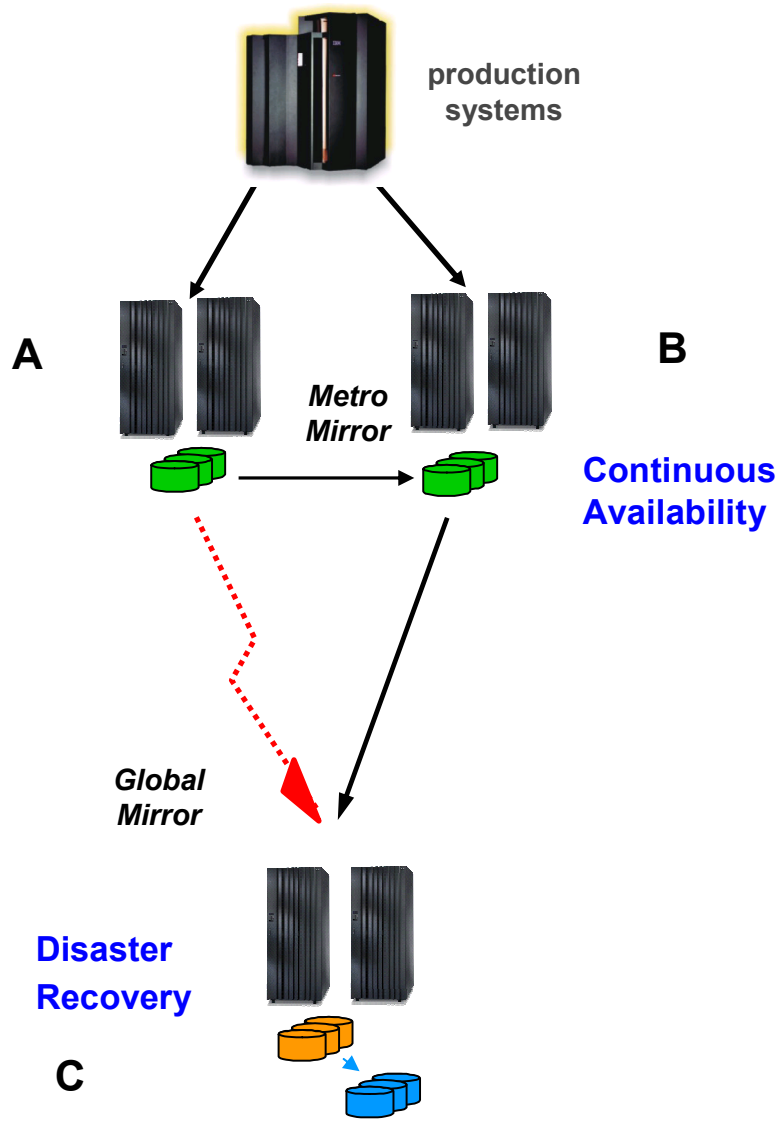
Three Site Metro / Global Mirror Overview



- **Metro Mirror (sync PPRC), and cascaded Global Mirror (async)**

- **Full ability to re-establish B-A-C recoverability while production continues to run at B**
 - ▶ Incremental change resync

Three Site Metro / Global Mirror Overview




- **Metro Mirror (sync PPRC), and cascaded Global Mirror (async)**

- **In event of site B outage, or failure of links A-B or B-C:**

- ▶ Can establish A-C recoverability quickly
 - Incremental change resync
- ▶ Can quickly re-establish A-B-C recoverability as soon as links or site B recover
 - Incremental change resync


TS3500 Tape Library Overview



192 Tape drives and slots – maximum configuration

Configuration Options

- **1 - 16 frames (plus two service frames)**
- **1 - 192 tape drives**
 - ▶ Up to 192 System z attached TS1120 tape drives
 - ▶ Up to 192 open system attached tape drives
 - TS1040/TS1030 or TS1120/3592 J1A tape drives are supported
 - Tape drives types and associated media require unique frames
- **Storage Capacity of over 13.1 Petabytes***
- **1 - 8 TS7700 Virtualization Engines**
 - ▶ Up to 2,048 virtual drives
 - ▶ Up to 2 million virtual volumes
- **1 – 8 Virtual Tape Servers**
 - ▶ Up to 2,048 virtual drives
 - ▶ Up to 4 million virtual volumes



* Using TS1120 tape drive with 3592 JB cartridge at a 3:1 compression

TS1120 Tape Drive Overview

- **2nd Generation of 3592 enterprise tape drive roadmap**
 - ▶ 104MBps performance (up to 260MBps at 3:1 compression)
 - ▶ 100 / 500 / 700GB native capacity cartridges (up to 300 GB / 1.5TB / 2.1TB at 3:1 compression)
 - Re-Writable and Write Once Read Many (WORM)
 - ▶ Supports data encryption and key management

- **Attaches to**
 - ▶ All IBM eServers (IBM System z™ via TS1120 Controller)
 - ▶ Selected HP and Sun Microsystems servers
 - ▶ Selected versions of Microsoft Windows™
 - ▶ Selected Linux editions

- **Supported in**
 - ▶ IBM TS3400, TS3500 and 3494 tape libraries
 - ▶ IBM 3592 C20 silo compatible frame
 - ▶ IBM 7014 Rack



3592 Cartridge Media



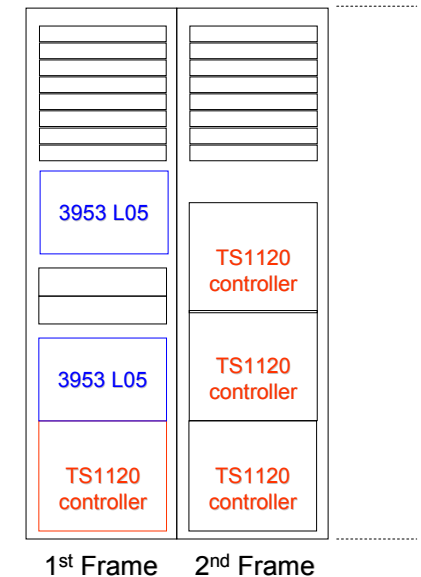
- **Cartridges are in rewritable and WORM**
- **Cartridges are available in three lengths**
 - ▶ JJ and JR cartridges provide rapid access to data
 - ▶ JA and JW cartridges provide fast access to data or high capacity
 - ▶ JB and JX extended data cartridges provide higher capacity
- **Cartridges can be formatted to either Gen 1 or Gen 2 formats¹**
 - ▶ TS1120 tape drives can read or write Gen 1 or Gen 2 formats
 - ▶ 3592 J1A tape drives can read or write the Gen 1 format

3592 Cartridge Media		TS1120 Tape Drive		3592 J1A Tape Drive	
Type	Format	Capacity	Performance	Capacity	Performance
JJ / JR	Gen 1	60GB	50MBps	60GB	40MBps
	Gen 2	100GB	104MBps		
JA / JW	Gen 1	300GB	50MBps	300GB	40MBps
	Gen 2	500GB	104MBps		
JB / JX	Gen 2	700GB	104MBps		

¹ iSeries only supports writing cartridges native 3592 Gen 2 format

System z Support

- **3953 L05 Library Manager**
 - ▶ Functionally equivalent to 3494 Library Manager
 - ▶ Employs identical microcode and server
- **The 3953 L05 is mounted in the first 3953 F05 frame**
 - ▶ Requires dual internal LAN features
 - ▶ Requires TS3000 System Console (TSSC)
- **Optional 3953 L05 provides redundancy**
 - ▶ Mounted in first 3953 F05 frame
 - ▶ Facilitates microcode updates
- **Issues TS3500 mount requests**
 - ▶ Accepts PLF commands via TS7700 VE, VTS or TS1120 controller
 - ▶ Issues SCSI media changer 'move medium' requests




TS7700 Virtualization Engine Overview

- **Includes significant architectural changes**
 - ▶ Architecture re-designed to facilitate future enhancements
 - ▶ Advanced IBM technology to increase performance and capacity
 - ▶ New business continuity option to increase flexibility and reduce cost
- **Provides higher performance and capacity**
 - ▶ Supports performance of up to 560MBps
 - ▶ Provides up to 18TB of native cache capacity (3:1 compression)
- **Supports attachment to IBM TS1120 and /or 3592 J1A tape drives**
 - ▶ Supports faster cache miss recall times and migration datarate
 - ▶ Supports up to 700GB on a 3592 JB cartridge
- **Supports tape drives in an IBM TS3500 or 3494 tape library**
- **Supports TS1120 data encryption**
- **Supports an optional GRID feature to support business continuity**



TS7700 Two Cluster GRID Configuration

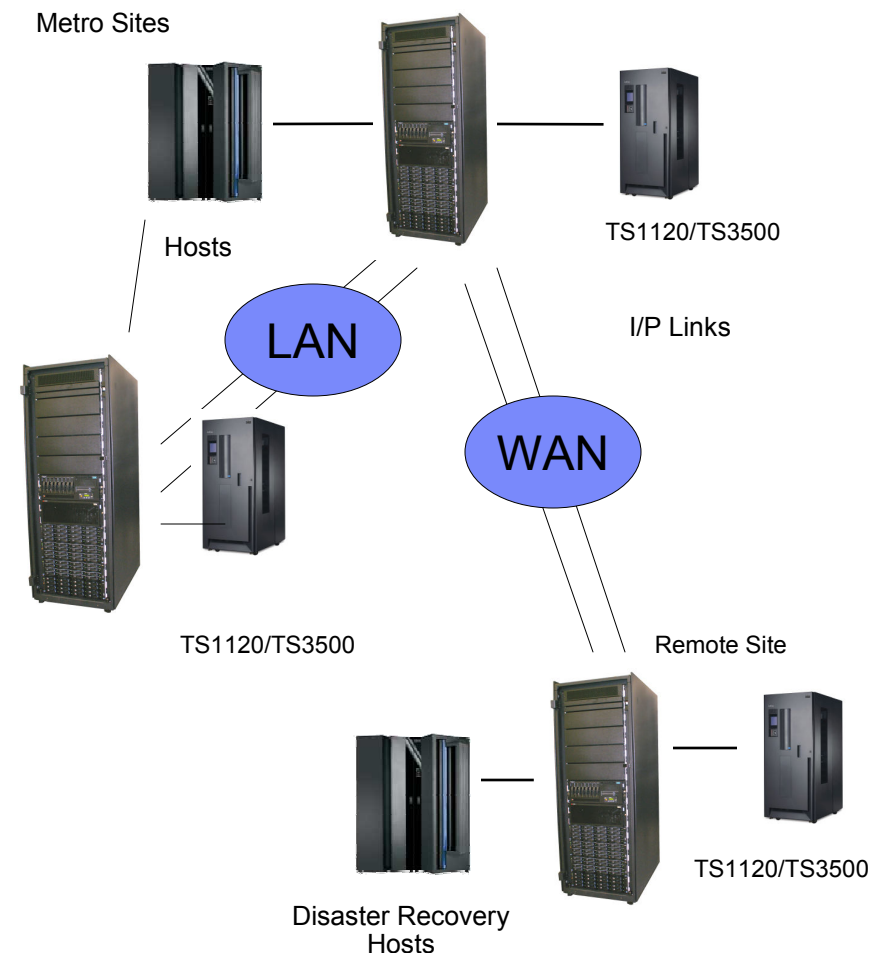
- **Couples two TS7700 Clusters together to form a GRID configuration**
 - ▶ VTCs have been eliminated
 - ▶ Hosts attach directly to the TS7700 Clusters
- **Any volume accessible through either TS7700 cluster in the GRID configuration**
- **I/P based replication**
 - ▶ Two 1Gbps Ethernet links
 - RJ45 Copper (Cat 6) FC1030
 - Optical Fiber FC 1031
 - ▶ Standard TCP/IP
- **Policy-based replication management**
- **Can be configured for disaster recovery or higher availability environments**



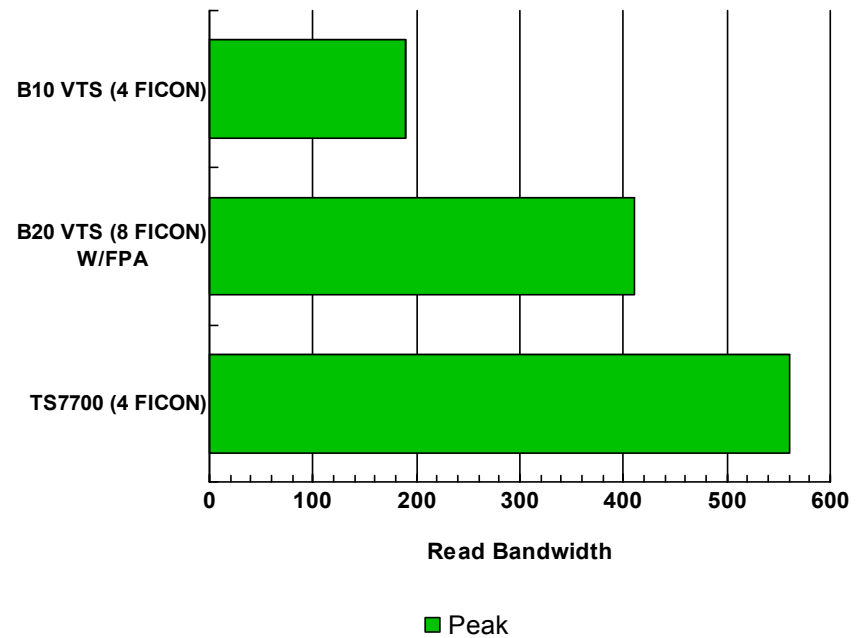
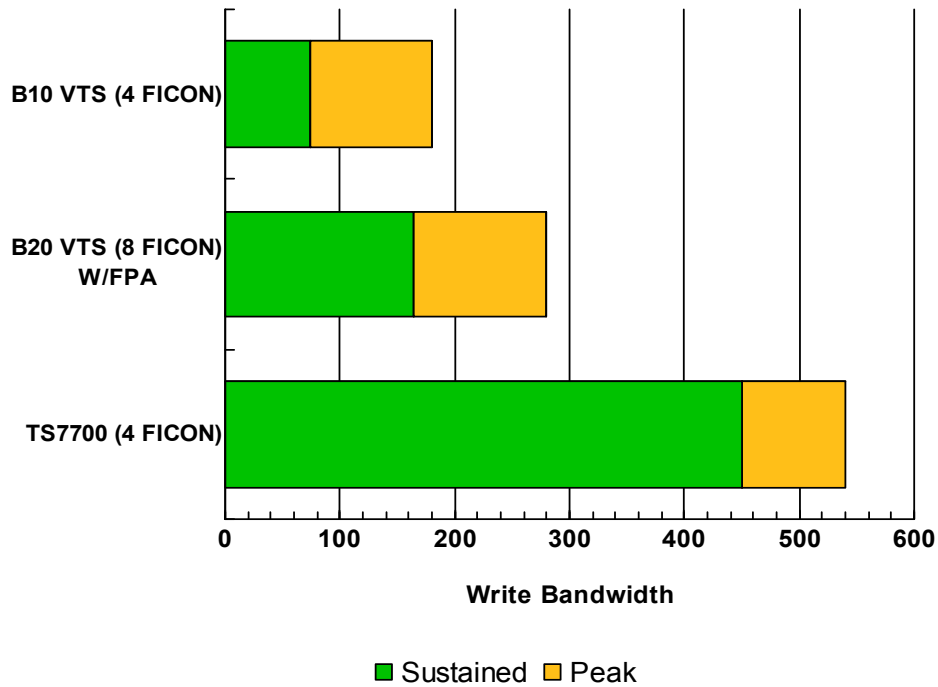
I/P replication may greatly simplify the infrastructure and management needed for a disaster recovery solution as compared to IBM's existing PTP VTS

TS7700 Three Cluster GRID Configuration for High Availability and Disaster Recovery

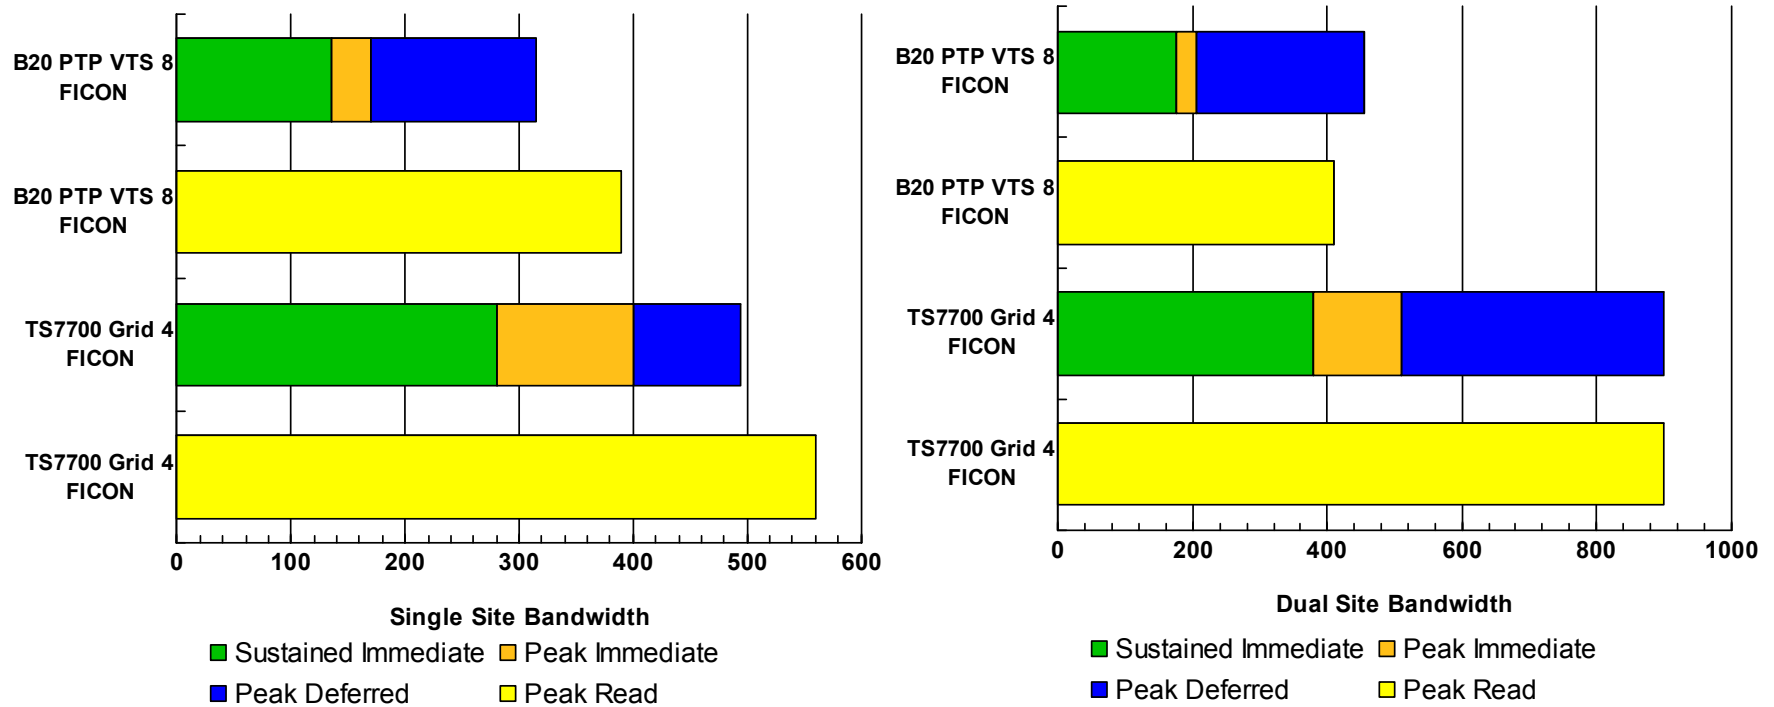
- Two TS7700s are located within 100km
- Interconnected through a Local Area Network
- Hosts are attached to both TS7700s
- One TS7700 located remotely, connected via Wide Area Network
- Immediate copies between the metro sites
- Remote site receives primarily deferred copies
 - ▶ Immediate copies reserved for critical volumes only



TS7700 Virtualization Engine FICON Performance



TS7700 Virtualization Engine GRID Performance¹



¹Local Area Network Connection, 0 distance, two cluster GRID

Deploy More Power Efficient Storage – *Facts* Tape Power & Cooling is Dramatically Better Than Disk


	SATA Disk	Tape
Hardware	\$5,067,652	\$269,771
Prod + DR Carts	\$0	\$168,750
Maintenance	\$0	\$240,036
Power & Cooling	\$805,098	\$40,396
Floorspace	\$493,200	\$227,452
Estimated Cost	\$6,365,950	\$946,405

← **Tape is very "green" @ 20x less energy expense**

Customer Storage Goals:

- Performance
- Compliance
- Data Security
- Disaster Protection
- Reduce TCO and energy costs

10 Year TCO Analysis



Scenario:

- Store 250TB
- 25% Growth Rate
- Over 10 Years
- DS4700 SATA Disk
- LTO 4 Tape Library

Source: IBM MI TCO Analysis

धन्यवाद

Hindi

謝謝

Traditional Chinese

תודה רבה

Hebrew

Tesekkurler

Turkish

Спасибо

Russian

شكراً

Arabic

Thank You

English

Gracias

Spanish

Grazie

Italian

ありがとうございました

Japanese

Obrigado

Brazilian Portuguese

ขอบคุณ

Thai

谢谢

Simplified Chinese

Merci

French

Danke

German

நன்றி

Tamil

감사합니다

Korean