



Information On Demand

IBM System Storage News and Trends

Keijo Ekman
Consultant IT Specialist
kekman@fi.ibm.com



ON DEMAND BUSINESS™

Large Systems Update 2006

© 2006 IBM Corporation

IBM System Storage



Agenda

- **Disk Family**
 - ▶ Overview
 - ▶ DS8000 news and trends
 - ▶ New Business Continuity Solutions
 - ▶ Open systems disk trends
- **Tape Family**
 - ▶ Tape Encryption – why and how?
 - ▶ New virtual tape solution
- **Software news**
 - ▶ Selected DFSMS 1.8 news


Disk Family news

IBM 305 RAMAC



The world's first disk storage system, the IBM 350 Disk Storage Unit (in the IBM 305 RAMAC computer), with 5 megabytes capacity, was introduced in 1956.

IBM 350 Disk Storage Unit



If an 18-wheeler “improves” as much as disk drives...



Dimension	18-wheeler (2006)	305 disk unit (1956)	500 GB SATA drive (2006)	Improvement	“Hyper Transport Device” (2056)
Length x width x height	23 x 4 x 2.5 m	1.5 x 1.7 x 0.75 m	14.5 x 10 x 2.5 cm	17.2x smaller	1.3 x 0.25 x 0.15 m
Volume	244 m ³	1.85 m ³	360 cm ³	5,128x smaller	1.68 ft ³
Cargo	30,000 kg	5 MB	500 GB	100,000x	3,000,000 tons
List price	\$150,000	\$35,000	\$1,280	27x less	\$5,485
Speed	100 km/h	8,800 char/sec (transfer rate)	65 MB/sec	61,960x	6.2 million km/h

IBM System Storage

IBM® System Storage™ DS8000 series Featuring IBM System Storage DS8000 Turbo models



DS8000 Turbo

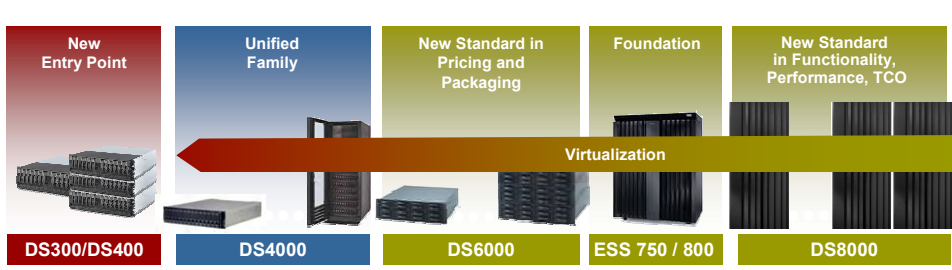
- **Setting a “*New Standard*” in Cost Effectiveness**
 - ✓ **Balanced Performance** – Up to 7X ESS Model 800
 - ✓ **Exceptional Scalability** – Up to 320TB physical capacity
 - ✓ **Virtualization for Simplification** – Storage System LPARs
 - ✓ **Flexibility** – FC and FATA disks to fit access/cost needs
 - ✓ **Extendibility** – Dramatic addressing capability
 - ✓ **Storage Management** – Full complement of interfaces/tools
 - ✓ **Availability** – Designed for 24X7 environments
 - ✓ **Resiliency** – Industry Leading Copy and Mirroring Capability
 - ✓ **Long Term Cost** – Four Year Warranty
Model to Model Upgradeability
- **Delivered through**
 - ✓ Server/Storage Integration – POWER5™ Technology
 - ✓ Exploitation of IBM Virtualization Engine™ Technology
 - ✓ IBM technology leadership and innovation

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

7 | © 2006 IBM Corporation | Large Systems Update 2006

IBM System Storage

Positioning the IBM TotalStorage and System Storage Family



DS300/DS400 **DS4000** **DS6000** **ESS 750 / 800** **DS8000**

Common management platform

Common suite of copy services

Virtualization

Compelling price points

Industry leading service and support

Enterprise Storage Continuum

IBM TotalStorage and System Storage Family innovations are designed to help you:

- Simplify the underlying IT infrastructure of storage and its management to lower help cost and complexity while increasing the ability to respond to changing needs.
- Maintain business continuity, security and data durability.
- Efficiently manage information throughout its lifecycle, relative to its business value.

8 | © 2006 IBM Corporation | Large Systems Update 2006

DS8000 – R2 Announcement Summary

- **Announcing new models**
 - Turbo Models 931, 932, and 9B2
- **Announcing new features for ALL models:**
 - IBM POWER5+ processor – standard on Turbo models
 - Processor memory for POWER 5+ processor – standard on Turbo models
 - 4Gb FCP / FICON adapter
 - 500GB 7,200 rpm FATA drives
 - 3-site Metro / Global Mirror
 - Earthquake resistance kit (feature code 1906!)
 - Ethernet adapter pair (for TPC RM support)
 - Performance Accelerator (Models 932, and 92E only)
- **Announcing new features for 931, 932, and 9B2 models:**
 - FICON/ESCON Attachment Licensed Function
 - Separation of Metro Mirror, and Global Mirror Licensed Functions
- **Withdrawal of current models**
 - 921, 922, and 9A2 models

Introducing: IBM System Storage DS8000 Turbo Models

Faster, more scalable, lower cost

- **Faster than a speeding disk system for quick access to data**
 - ▶ New Power5+ processor yields up to 15% performance improvement for transaction processing workloads over previous models
 - ▶ AIX/DB2 synergy items can enable greater efficiency and higher performance for high priority applications
 - ▶ 4Gbps Fibre Channel/FICON® adapters support high-bandwidth host attachments
- **More scalable to support massive amounts of information, greater consolidation and simplification**
 - ▶ Up to 320 TB physical capacity with 500GB disk drives
- **Able to leap to new levels of cost effectiveness**
 - ▶ Lower base price on new Turbo base models
 - ▶ 500GB FATA disk drives enable lower cost per GB for less frequently accessed near-line application needs
 - ▶ Use fewer host ports and network infrastructure with 4Gbps Fibre Channel/FICON® adapters



POWER5+ Processor

- Compared to current POWER5 processor, the POWER5+ processor may enable up to 15% performance improvement in I/O operations per second in transaction processing workload environments
- POWER5+ enhancements over POWER5
 - Better Performance
 - 15% Higher frequencies
 - Higher bus clock rate
 - Large page size support
 - Memory controller improvements
 - Faster Processors in DS8000 Turbo models
 - **DS8300: 1.9 Ghz, DS8300 Turbo: 2.2 Ghz**
 - **DS8100: 1.5 Ghz, DS8100 Turbo: 2.2 Ghz**
- Upgrade possible from current models

DB2 / AIX Synergy Item – Cooperative Caching (with I/O priority)

- **Host gives cache management hints to storage facility**
 - Up to three conditions can be signaled in a read/write CDB
 - Currently can indicate that data will not be re-accessed soon
- **AIX Support**
 - Trusted Applications can provide a cache hint
 - Supported by MPIO PCM provided with SDD (SDDPCM V2.1.1.1)
 - Cache hint passed from application to storage facility
 - AIX 5.3E+ IY85769 or AIX 5.3H (future) and above
 - Multi-pathing: Devices.fcp.disk.ibm.mpio.rte and devices.sddpcm.53E.rte
 - 64 bit kernel and raw devices only (i.e. not with file system)
- **DB2 Support**
 - DB2 V9.1
 - DB2 is a trusted application and provides cache hints
- **DS8000 Support**
 - Storage facility uses cache hint to manage retention period of cached data
 - Data not to be re-accessed is placed on accelerated LRU queue
 - GA: November 17, 2006 and requires R2

IBM DS8000 Turbo 4Gb/s FC/FICON high-speed connectivity



- Enables potential reduction in DS8000 host ports needed, helping to lower costs*
- Can help enable host adapter/channel and link consolidation to help simplify management and reduce the cost of the SAN infrastructure*
- Designed to offer up to 50% improvement in single port MB/second throughput*
 - up to 300MB/sec per port
 - 540MB/sec per 4-port card
- Flexibility with both 2Gb/s and 4Gb/s FC/FICON® adapters – shortwave and longwave*

IBM has a full range of Disk, SAN, Tape, Software, & Services ready to support 4Gb FC Connectivity

*IBM is first vendor to offer **4Gbps FICON** disk system connectivity*

4Gb FCP/FICON Host Adapters

- **Designed to offer up to 50% throughput performance improvement in a single port MB/second**
 - **Potential cost savings with reduction in the number of host ports needed**
- Overall card throughput remains constant
- **4Gb host adapters available**
 - Longwave
 - Shortwave
 - 4Gb 10km LW FCP/FICON
 - Autonegotiate to either 4Gb or 2Gb link speeds
- **The 4Gb FCP/FICON adapter is an optional feature for 92x, 9Ax, 93x and 9Bx Models**

DS8000 Flexible Range of Disk Options

- **Full range of FATA and Fibre Channel disk drive options**

- ▶ 500GB 7,200 RPM FATA drives – Near-line class
- ▶ 73GB 15k RPM, 146GB 10k or 15k RPM, 300GB 10k RPM FC drives – Enterprise class
- ▶ 300GB 15k FC drives under consideration




- **16-disk drive pack ordering of all drives, FATA and FC**

- ▶ All DS8000 disk drives are installed in disk drive sets that contains 16 identical disk drives (of the same capacity and rpm)
- ▶ FC and FATA drives are both available in Capacity on Demand features

- **Intermix enabled for FATA and FC disk sets**

- ▶ Intermix of different rpm/capacity FC and/or FATA disk sets
- ▶ Supports intermix within a system and within a rack
- ▶ Supports intermix of DDM class, capacity, and RPM on a DA-Pair
- ▶ Homogeneous FATA or FC systems are also supported – there is no need to configure FC drives if all FATA is desired
- ▶ Intermix of FATA and FC drives in FlashCopy and mirroring relationships is supported



Disk Drive Classes

Class	Definition	Usage / Duty Cycle
Fibre Channel <i>Enterprise-class</i>	<ul style="list-style-type: none"> ▪ Drives are intended for servers or multi-user systems ▪ Higher performance, availability, reliability and functionality ▪ Higher costs due to limited manufacturing (comparatively) and more robust and expensive technology 	<ul style="list-style-type: none"> ▪ Designed for very high activity, >70% duty cycle ▪ 1.2M MTBF <ul style="list-style-type: none"> ▶ Operating times are 24 hours/day
FATA / SATA <i>Enterprise-class near-line</i>	<ul style="list-style-type: none"> ▪ Drives are intended for multi-user systems under low workloads ▪ Lower performance, but similar reliability to Enterprise drives under low workloads ▪ Highest capacity 	<ul style="list-style-type: none"> ▪ Designed for moderate activity, around 20-30% duty cycle ▪ 1.0M MTBF <ul style="list-style-type: none"> ▶ Operating times are 24 hours/day
Desktop	<ul style="list-style-type: none"> ▪ Generally intended for single-user systems, designed for lower workloads ▪ Lower performance and availability requirements ▪ Generally lower costs due to mass manufacturing and less expensive technology 	<ul style="list-style-type: none"> ▪ Designed for low activity, around 10% duty cycles ▪ .8M MTBF <ul style="list-style-type: none"> ▶ Operating times assumed to be typically 8 hours/day


Drive Comparison – Mechanism



	SATA-1	SATA-2	FATA	Fibre Channel	Fibre Channel
Capacity	250—500 GB	250—500 GB	500 GB	73, 146, 300 GB	73, 146 GB
Rotational speed	7,200	7,200	7,200	10,000	15,000
Average Read	8.5 ms	8.5 ms	8.5 ms	4.9 ms	3.5 ms
Average Write	9.5 ms	9.5 ms	9.5 ms	5.5 ms	4.0 ms
Max. sustained data transfer rate	34—59 MB/sec	34—59 MB/sec	34—59 MB/sec	39—80 MB/s	58—96 MB/s

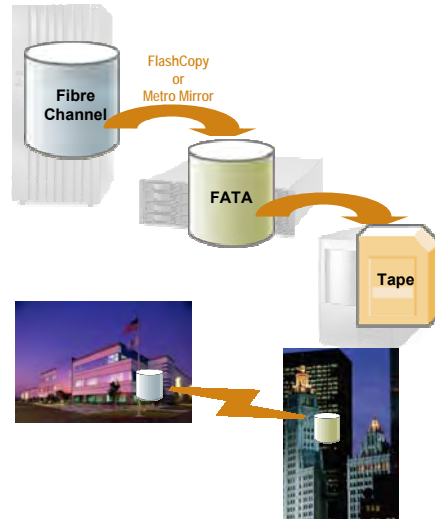
- Drive level specifications of the SATA-1, SATA-2 and FATA are similar
 - Note: Drive level specifications are only one part of a system's overall capabilities
- Focus on the DS storage system level performance and capabilities

DS6800 Performance with FATA Disk Drives with 64 disks



Where to Implement FATA Storage Volumes

- **FATA disks open up new growth opportunities improving backup strategies**
 - Low cost backup helps provide better affordability as compared to using FC drives
- **Shrinking backup windows?**
 - Intermediate storage tier between primary disk and tape
 - FlashCopy or Mirror primary to near-line storage, and then near-line to tape
- **Fast data recovery**
 - Use near-line volumes for quick data restoration
- **Budget too tight for disaster recovery site?**
 - FATA drives offer a lower-cost alternative for implementing a secondary site



Where to Implement FATA disks


- **Disk-based nearline storage**
 - Storage in hierarchy between disk and tape
- **Virtual tape or tape cache**
 - Spinning archive prior to tape enables faster restores
- **Disaster recovery**
 - Target for remote replication
 - Well suited for source of restoration, not a "hot site"
- **Fixed content / Managed Retention data**
 - Medical imaging, object folders, historical reports
- **Temp work space for short-term processing**
 - Minimal IOPS performance requirements



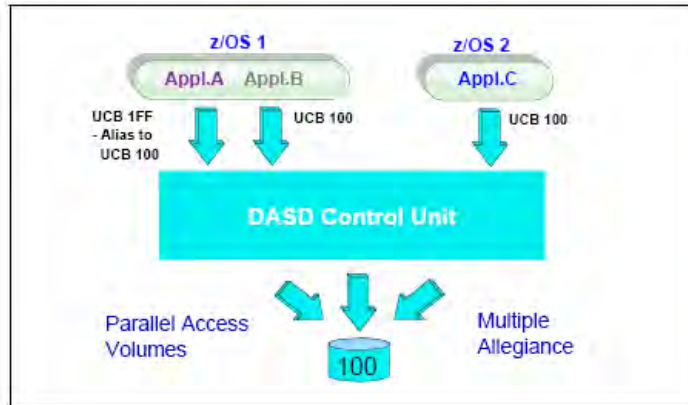
Multiple Allegiance and PAV overview

- Two function that a allows multiple I/Os to be executed concurrently against the same volume in a z/OS environment
 - In the case of Multiple Allegiance, the I/O are coming from different LPAR of Z/OS systems
 - In the case of Parallel Access Volumes, the I/O are coming from the same LPAR of Z/OS systems
 - **Static PAV**: Aliases are always associated with the same Base Address
 - **Dynamic PAV**: Aliases are reassigned to any base address as need dictates: WLM function call Dynamic Alias Management ; reactive alias assignment
 - **“HyperPAV”**: Potential development: On demand/ Proactive alias assignment

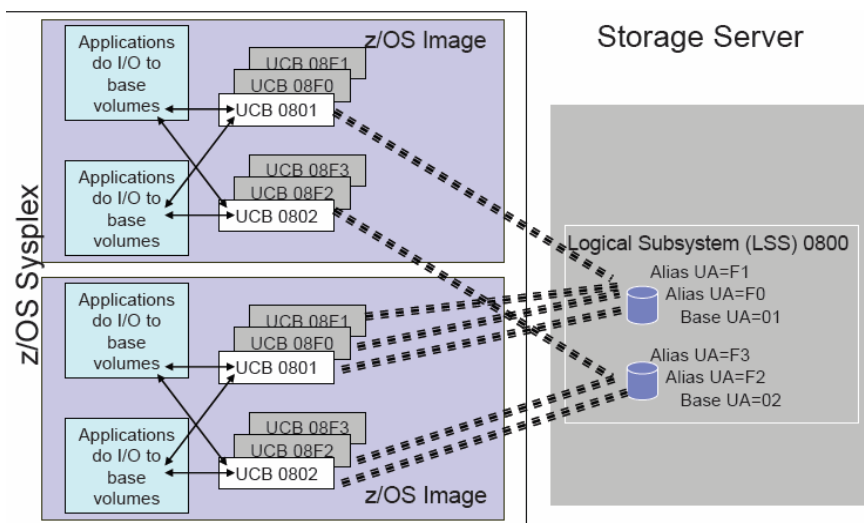
Value of Parallel Access Volumes (PAV)



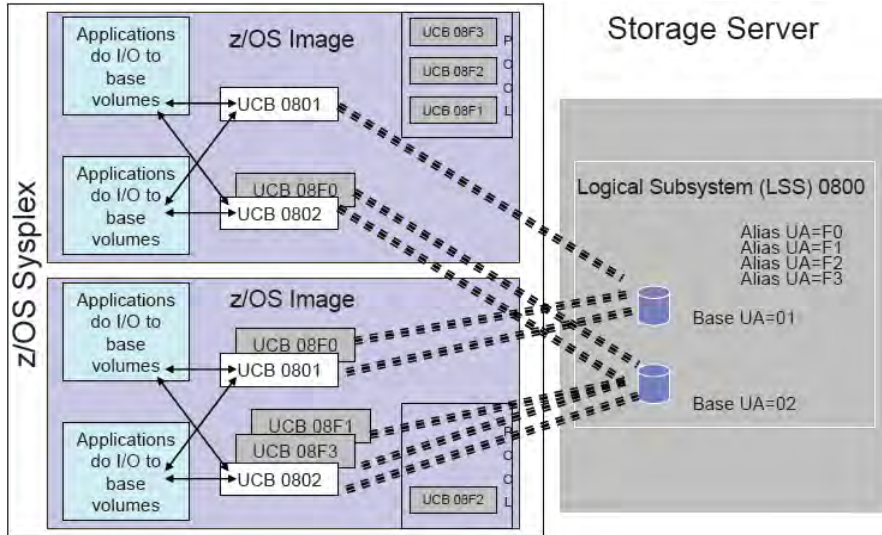
Concurrent I/O with PAV and Multiple Allegiance



Parallel Access Volumes -Today



Potential Implementation



Potential Benefits of new Implementation

- Reduce number of required aliases
 - Give back addressable device numbers
- z/OS can react more quickly to I/O loads
 - Eliminates need for multi-system interlock (DST)
 - React instantaneously to market open conditions
- Overhead of managing alias exposures reduced
 - WLM not involved in measuring and moving aliases
 - Alias moves not coordinated throughout SYSPLEX
- Initialization doesn't require "static" bindings
 - Static bindings not required after swaps
- Increases I/O Parallelism

Possible Migration Scenario

- No HCD changes required
- Deployment can be staged
 - Load code on DS8000
 - Can run without exploiting this feature if necessary using z/OS PARMLIB option
 - Enable new feature on z/OS images that want to utilize it
 - Eventually enable new feature on all z/OS images in the sysplex
 - Reduce the number of aliases defined



GDPS / Global Mirror (GDPS/GM) Overview

Business Continuity Objectives

- Determine your Objectives for Business Continuance (by application)


- Recovery Time Objective (RTO)

- how long can you afford to be without your systems?



- Recovery Point Objective (RPO)

- when it is recovered, how much data can you afford to recreate?



- Network Recovery Objective (NRO)



- how long to switch over network?



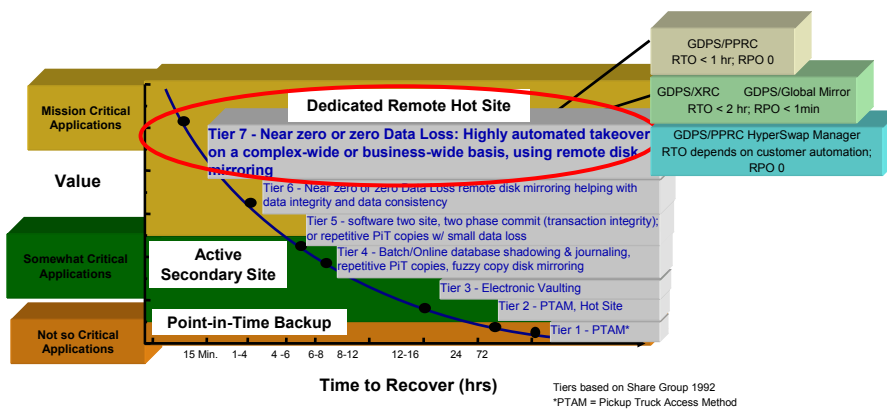
- RTO/RPO is a Business Decision not IT decision

- Balance risk, cost vs recovery curve

- Regulatory compliance

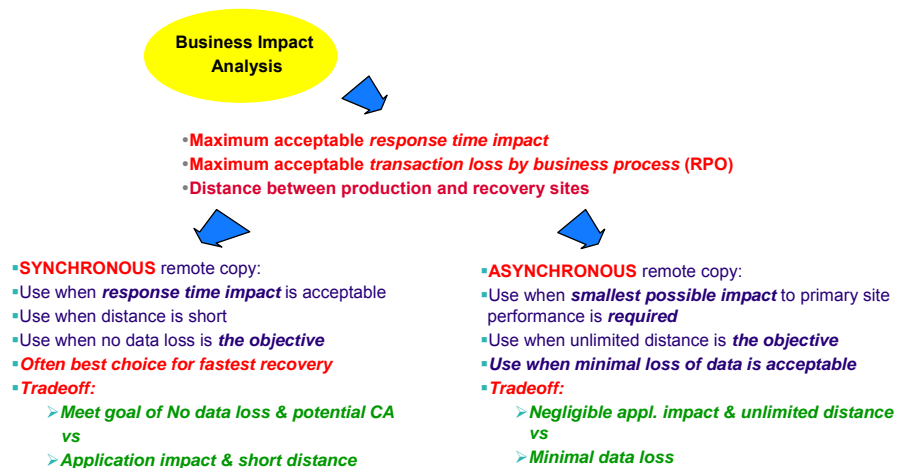
Determining the cost vs. RTO recovery curve is the key to selecting proper solution(s)

Tiers of Disaster Recovery: Level Setting GDPS



Best D/R practice is blend tiers of solutions in order to maximize application coverage at lowest possible cost. One size, one technology, or one methodology does not fit all applications

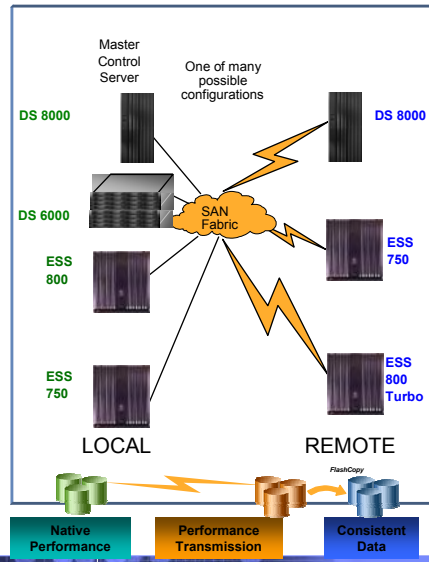
Business Impact Analysis - Synchronous vs. Asynchronous



Business and regulatory pressures

- Various financial services regulatory bodies are requiring long distance disaster recovery solutions
- Many customers especially in financial services already have short distance or campus synchronous solutions providing high/continuous availability and some measure of DR protection
- Business units are not willing to forgo the availability benefits of synchronous mirroring solutions but need to implement longer distance recovery solutions
- In some cases the regulatory solution is seen as having a very low probability of being used and so objective is to satisfy regulator with minimum cost

Global Mirror Copy Technology



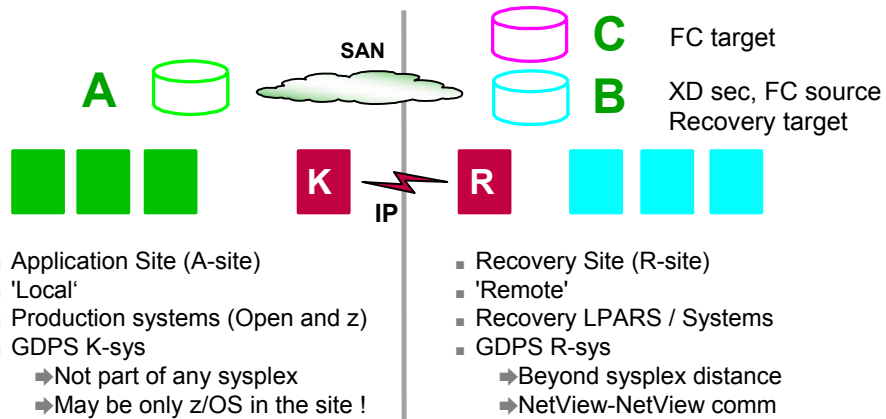
Characteristics:

- ▶ **Global Distance:** Two-site, unlimited distance, data consistency
 - ▶ **Heterogeneous:** System z and/or Open systems data
 - ▶ **Scalable:** Up to 15 total ESSs in Global Mirror session (with RPQ)
 - ▶ **Flexible:** Many possible configurations
 - ▶ **Application Performance:** No impact
 - ▶ **Mirroring Performance:** Two Fibre Channel links per disk subsystem pair sufficient for most workloads
- ### Intended Benefits
- ▶ **Autonomic:** No active external controlling software required to form consistency groups
 - ▶ **Saves cost:** No server cycles required to manage data consistency
 - ▶ **Lowers TCO:** designed to provide improved performance at lower cost point

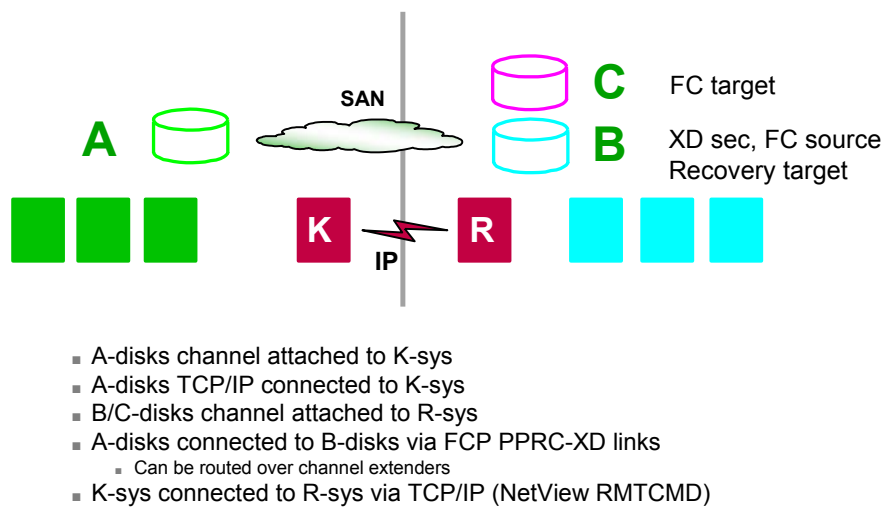
GDPS/GM Solution Overview

- Manage GM copy technology
 - Manage GM configuration
 - Multiple GM sessions
 - Each session can have System z and/or Open systems data
 - Manage day-to-day GM operations from single point of control
 - Manage GM recovery
- Automate restart of production in recovery site
 - System z only
 - Open systems must be manually restarted
- RTO <2 hours after declaration of disaster
- RPO workload/bandwidth dependant


GDPS/GM Terminology / Topology



GDPS/GM Connectivity



GDPS/GM K & Rsys functions

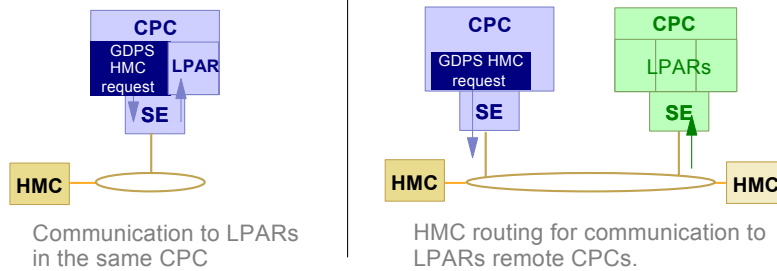


- Ksys
 - Day-to-day GM management
 - Configuration definition/management
 - Monitoring and alerting

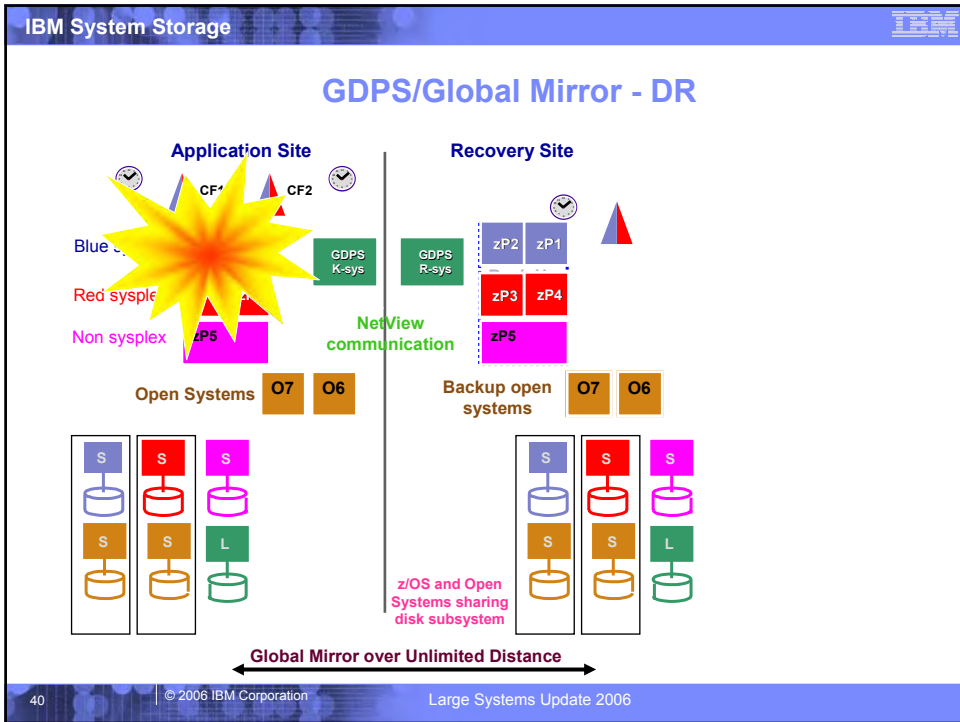
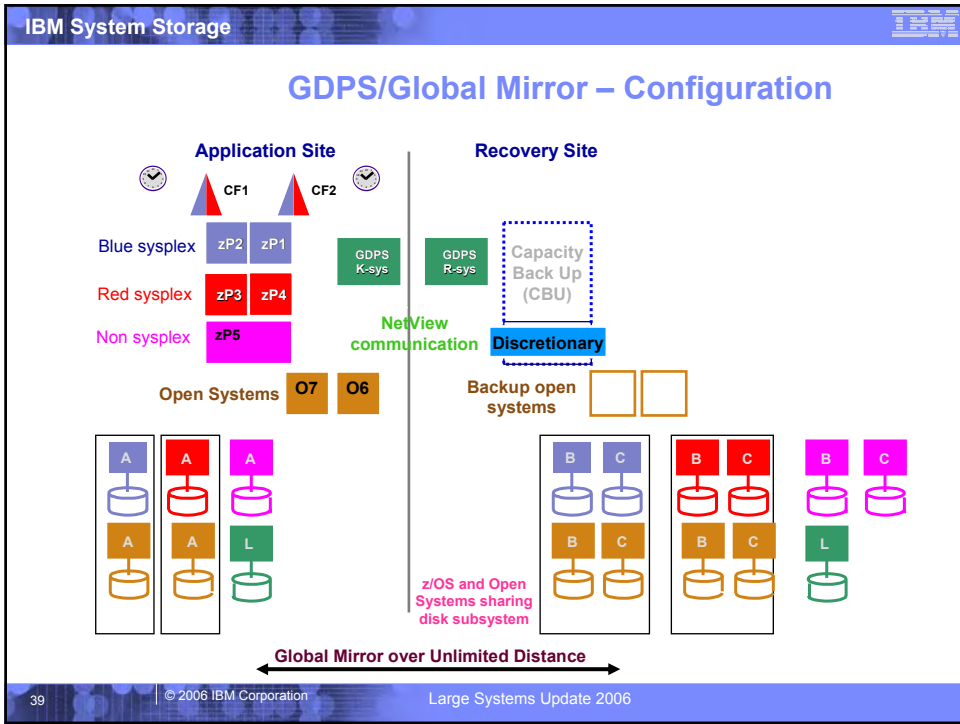
- Rsys
 - Address translation
 - Remote site configuration validation
 - Remote site monitoring / alerting
 - GM recovery management
 - CKD & Open
 - Testing with B-disk & real D/R
 - System restart automation (z only)

Rsys must have the most current configuration to guarantee successful recovery/restart


GDPS HMC automation



- BCP Internal Interface (BCPii)
 - SA for z/OS function exploited by GDPS
 - GDPS/GM R-sys function
 - Reconfigure and restart production in recovery site
 - For testing or for real disaster
- Automation for manually intensive tasks




D/R testing with B-disks



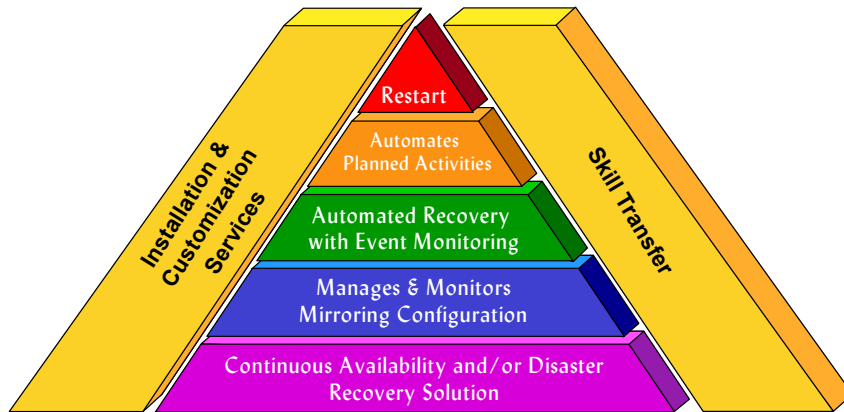
- Suspend mirroring on K-sys
 - Production continues to run on A-disks (changes tracked in bitmap)
- R-sys actions (scripted):
 - Recover B-disks (C-disks will be consistent but will age during test)
 - Reconfigure R-site (activate CBU, activate LPARs)
 - IPL recovery systems into backup LPARs for testing
 - Can refresh B-disks from C-disks for retest
 - Using failover/failback logic R-sys backs out changes from B-disks
- K-sys resynchronizes changes made on production to recovery site
- No need to perform initial copy when testing is complete

GDPS/GM - Future ?



- ➔ Create D disks (FlashCopy) for recovery testing
 - ➔ Maintain DR readiness with B/C disks
- ➔ 'Go home' when Application site is restored following a recovery
- ➔ Note: No SOD for any future function at this time
 - ➔ Future function will depend on customer requirements

GDPS/GM offers a comprehensive D/R solution, not just remote copy technology




IBM Metro Global Mirror functionality and recovery scenarios

Metro/Global Mirror – What is it?

- ▶ Today the DS8000 series supports 2-site hardware-based remote mirror and copy solutions, including IBM System Storage Metro Mirror and IBM System Storage Global Mirror.
- ▶ The DS8000 series has been enhanced to support 3-site Metro/Global Mirror configurations. The Metro and Global Mirror function utilizes synchronous mirroring (Metro Mirror) from a local A-site to a metro distance B-site and asynchronous mirroring (Global Mirror) from an intermediate B-site to a remote C-site.
- ▶ This function, referred to as Metro/Global Mirror, is designed to provide planned and unplanned outage three-site enterprise disk data replication — which can help meet rigorous three-site business resiliency needs of the enterprise data center.

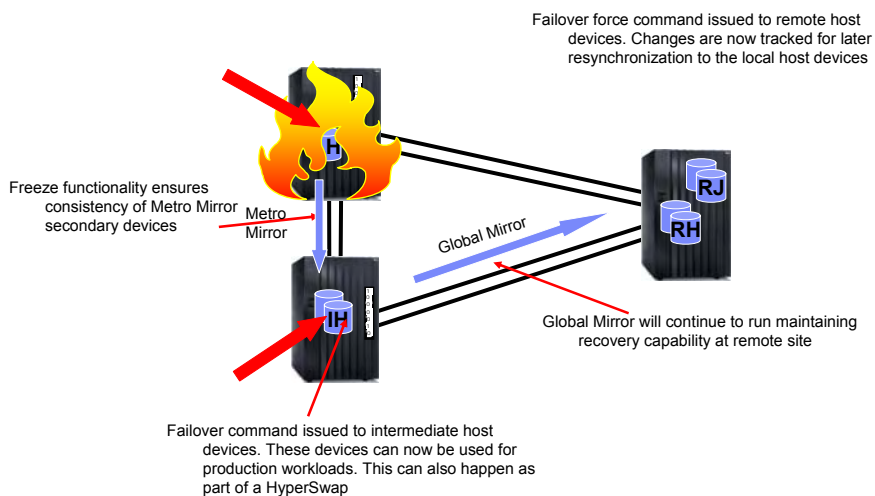
Metro/Global Mirror - Graphically



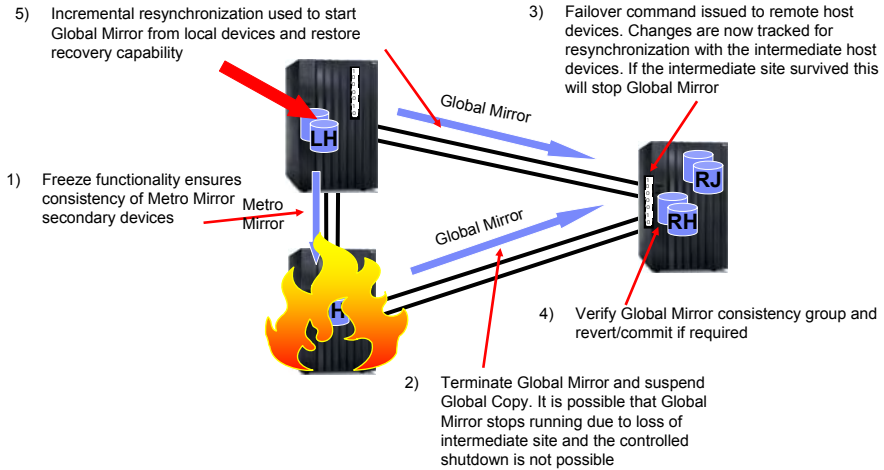
Metro Global Mirror scenarios

- **Loss or planned outage of remote site**
 - Recoverability is maintained through Metro Mirror relationships
 - Remote site is reintroduced by incremental resync from intermediate site
- **Loss or planned outage of local site**
 - Move production and continue Global Mirror from intermediate to remote
 - Intermediate and Local role switch and we are logically in 2-site mode Local to Remote
- **Loss or planned outage of intermediate site**
 - Start Global Mirror from Local to Remote
 - We are running in 2-site mode from Local to Remote
- **Reintroduction of intermediate site when running in 2-site mode Local to Remote**
 - This covers recovery from the above scenarios for both planned and unplanned outages

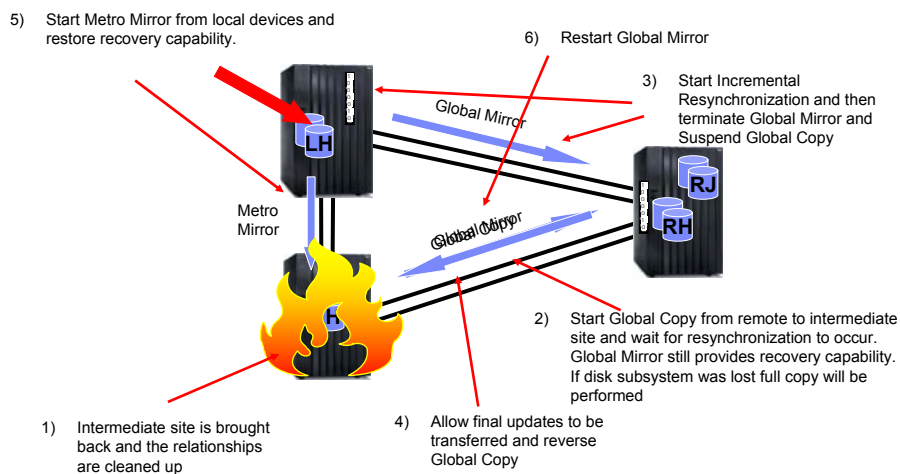
Metro Global Mirror loss of primary site and move to 2-site mode



Metro Global Mirror loss of intermediate site and move to 2-site mode




Metro Global Mirror reintroduction of intermediate site

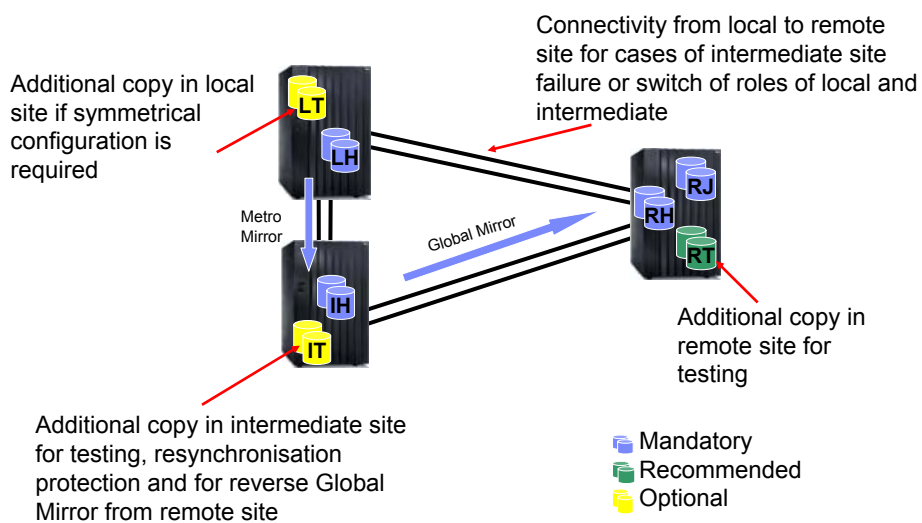


Metro Global Mirror architecture

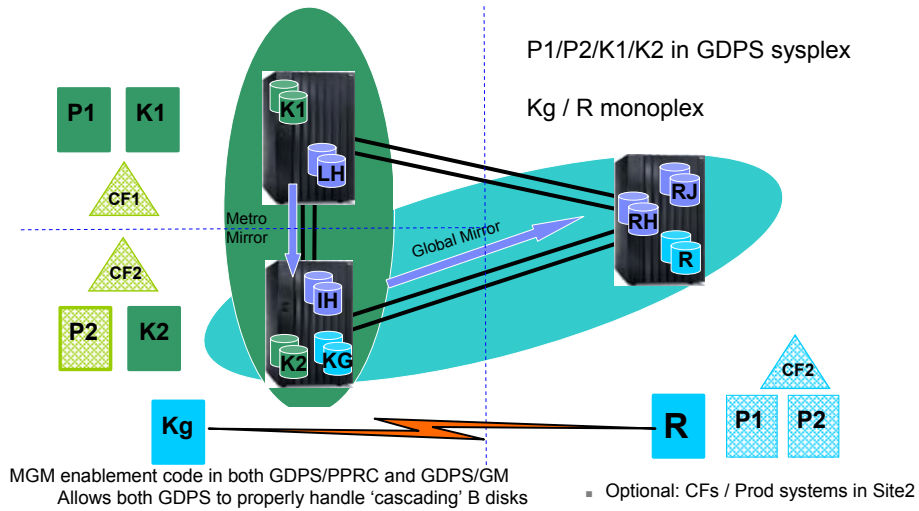
- **Combines Metro Mirror and Global Mirror**
 - Allows for local continuous availability and disaster recovery and out of region disaster recovery
- **RPO of 0 for "A" site failure**
 - Can recover to either intermediate site or remote site
 - Zero RPO implies automation to ensure no production updates if mirroring stops
- **Potential RPO of seconds for local and intermediate twin site failure**
 - Depends on workload and bandwidth between intermediate and remote site



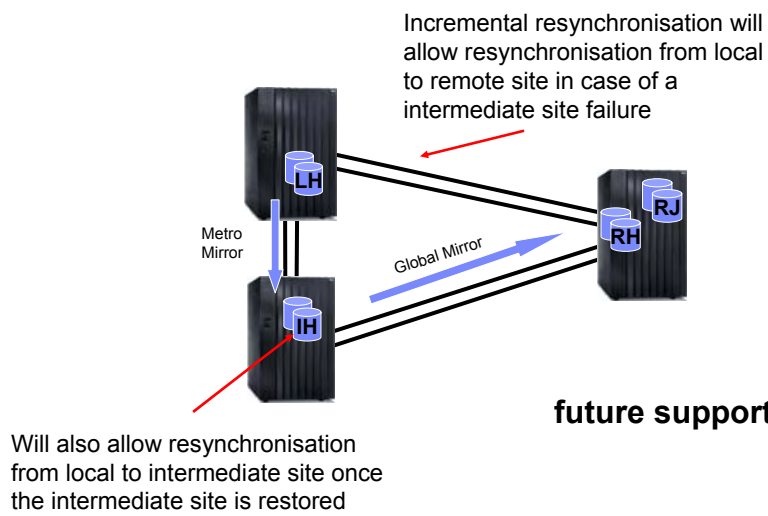
Metro Global Mirror Topologies



GDPS/MGM Configuration (initial release)



Incremental Resynchronisation for Metro Global Mirror



IBM System Storage

SODs for Leadership in Resiliency with DS8000 and IBM System z Servers

- Statement of Direction: Consistency groups that span z/OS® Global Mirror (XRC) and Global Mirror (Asynch. PPRC). This is intended to help customers:**
 - Achieve rapid recovery by providing a common, consistent recovery point for applications and data residing on open and mainframe
 - Enable exceptional scalability with consistency groups that span multiple IBM DS8000s or ESSs
 - Provide flexibility and choice to implement multiple IBM global mirroring solutions across the enterprise with support for data integrity and consistency at a remote site
- Statement of Direction: Continuous data protection function on the DS8000 for data in System z environments. This is intended to:**
 - Offer a new level of protection with a function for data backup and recovery on IBM System z servers using IBM DS8000
 - Record data changes continually, as opposed to halting an application's I/O to create the data copy
 - Store data changes incrementally rather than storing numerous data images
 - Provide rapid recovery of data from any point in time captured in the backup capacity

Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only

Integrating the enterprise

55 © 2006 IBM Corporation Large Systems Update 2006

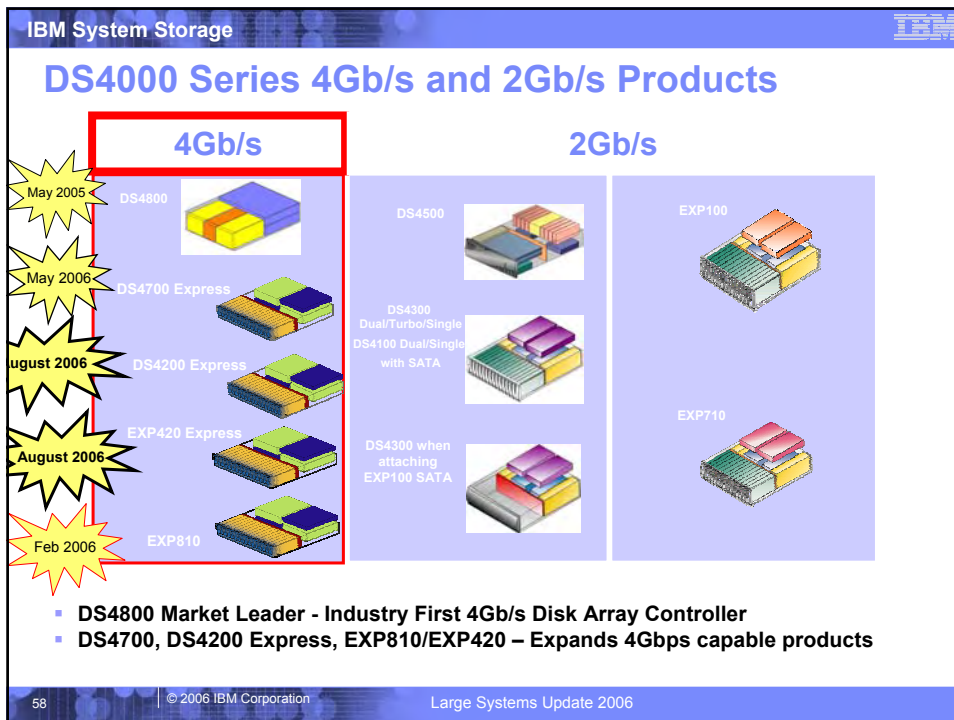
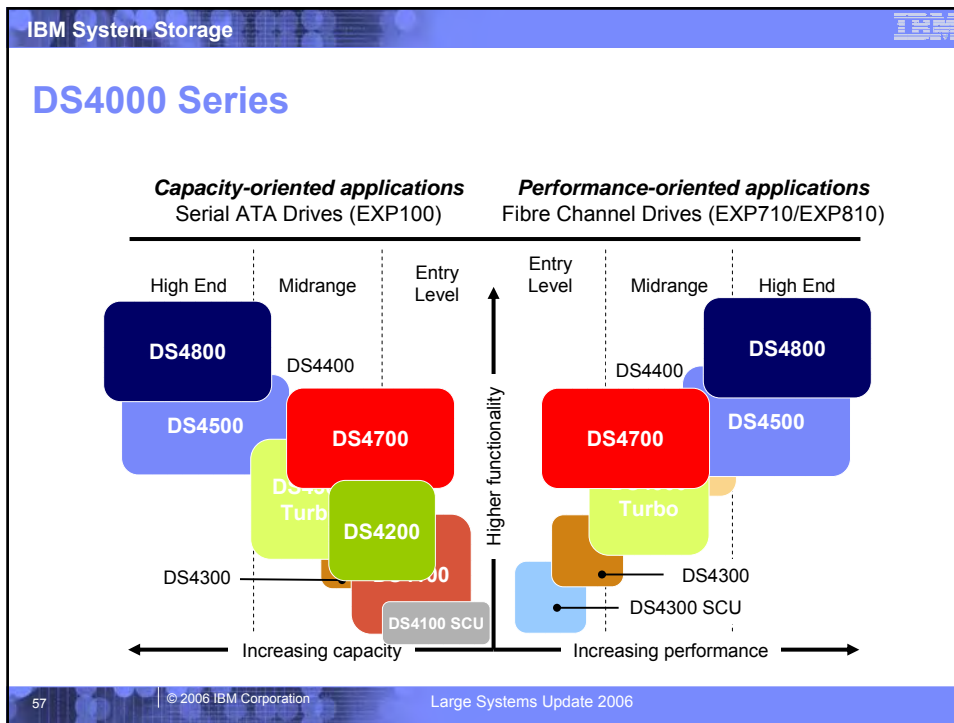
IBM System Storage

DS4000 Series Positioning

This table represents general positioning and not technical capabilities

	Entry-level	Midrange	High-end
Workload	Light workloads	Medium workloads	Heavy workloads
Typical platforms	Windows/Linux direct attach	Homogeneous SAN	Heterogeneous SAN
Applications	operational	business-important	business-critical
Primary drive type	SATA	FC / SATA	FC
Data replication	FlashCopy®	FlashCopy®, VolumeCopy, Enhanced Remote Mirroring	FlashCopy, VolumeCopy, Enhanced Remote Mirroring

56 © 2006 IBM Corporation Large Systems Update 2006



SAN Volume Controller (SVC) Status and News

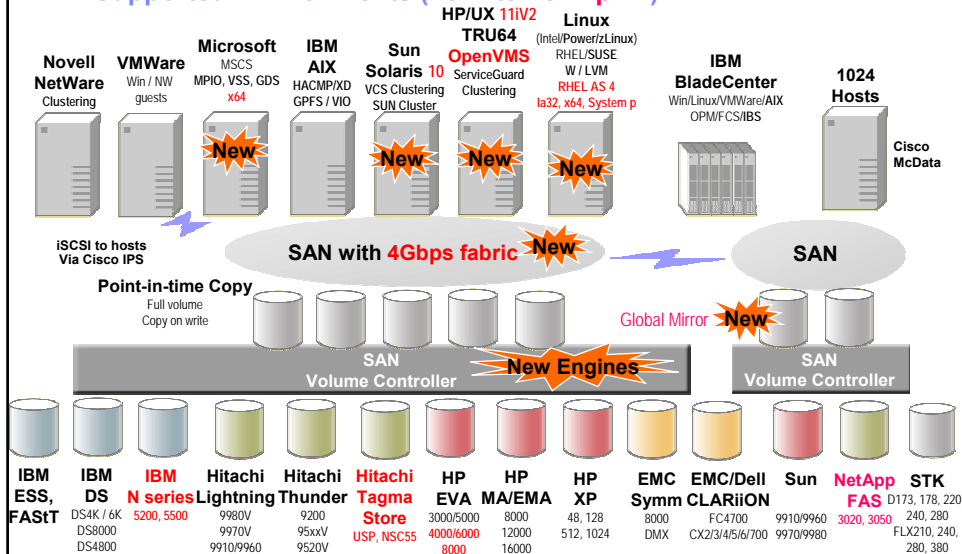
SVC: Did You Know?

- IBM has **40 years experience** in virtualization technologies
- SAN Volume Controller has more than **2,000** clients, over **110** references, and is entering its fourth year of market acceptance
- SAN Volume Controller delivers the **FASTEST** storage performance benchmark ever recorded for ALL controllers
- **15PB** of client data managed today and growing!
- SVC manages 53% of network-based storage virtualization appliance capacity (Source: IDC/IBM)
- SAN Volume Controller can virtualize IBM and **non-IBM** storage (over 75 systems from EMC, HP, HDS, Sun, Dell, STK, NetApp)

What's New with SAN Volume Controller Version 4

- **Global Mirror**
 - Extends business continuity reach to practically unlimited distances
- **4Gbps fabric support**
 - Improves infrastructure simplification through use of latest SAN technologies
- **Improved performance statistics**
 - Enables better monitoring, understanding, and planning of SVC systems
- **Cluster non-disruptive upgrade capability**
 - Foundation for continued growth and technology exploitation
- **Additional server and disk system support**
 - Extends range of environments supported by SVC

System Storage SAN Volume Controller Version 4.1 Supported Environments (new items in pink)



How Much Value?




- **Enterprise Strategy Group reports that early virtualization adopters on average every year save:**
 - 24% on hardware costs
 - 16% on software costs
 - 19% on SAN administration costs
- **With a \$1 million budget spending \$500,000 on hardware, \$200,000 on software, and \$300,000 on administration**

Annual savings would be \$209,000

Tape News
What??? Isn't tape dead?


Over 50 Years of Tape Innovation

- **Starting in 1952**
 - ▶ IBM 726 Tape Unit
 - 7,500 characters per second
 - 100 bits per inch
- **and continuing in 2005**
 - ▶ IBM TS1120 Tape Drive
 - up to 100 MB/sec¹
 - up to 500 GB¹




IBM High Data Density Tape Achievement...

...Projected to Enable up to 8 Terabyte Tape Cartridge of the Future



Tape's Evolving Role


- **Tape is an integral part of the storage hierarchy**
 - ▶ Customers store 4-15X more data on tape than disk
- **Tape is low cost**
- **Tape is intrinsically "On-Demand"**
- **Tape is removable and portable**
- **Tape provides high volumetric efficiency**
- **Tape media has a long life**
- **Tape is ideally suited for:**
 - ▶ Information Lifecycle Management
 - ▶ Infrastructure Simplification
 - ▶ Business Continuance



Sources: Disk IDC, 2003
SATA Disk - IBM estimate
Tape - IBM, based on technology roadmaps

The Linear Tape Open (LTO) Standard

- **Since September 2000 LTO has proliferated**
 - ▶ Collectively the three LTO technology companies have shipped
 - Over 30 million tape cartridges
 - Over 1,000,000 tape drives
 - ▶ LTO Roadmap extended November 2004
 - Added support for LTO Write Once Read Many Cartridges
 - Added Gen 5 and Gen 6



TS1120 Tape Drive Overview

- **2nd Generation of 3592 enterprise tape drive roadmap**
 - ▶ 104 MB/sec performance (up to 260 MB/s at 3:1 compression)
 - ▶ Dual 4Gbps FC ports
 - ▶ 100 / 500 GB native capacity (up to 300 GB / 1.5 TB at 3:1 compression)
 - Re-Writable and Write Once Read Many (WORM) cartridges
 - ▶ **Now with encryption!**
- **Supported by**
 - ▶ All IBM servers (IBM System z™ via TS1120 Controller)
 - ▶ Selected HP and Sun Microsystems servers
 - ▶ Selected versions of Microsoft Windows™
 - ▶ Selected Linux editions
- **Supported in**
 - ▶ IBM 3494 and TS3500 tape libraries
 - ▶ IBM 3592 C20 silo compatible frame
 - ▶ IBM 7014 Rack

