Session G17:
IBM System Storage Overview for the System z Environment

IBM System z Expo
September 17-21, 2007
San Antonio, TX

Curtis Neal
IBM San Jose EBC Outbound Team
curtisne@us.ibm.com

2007 Storage Top Ten Trends

Tiered data classification, prioritization, and movement
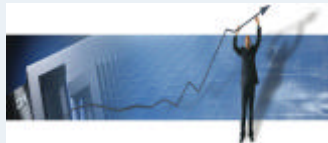
De-duplication (single instance storage) ramping up

Continued compliance, regulatory, and governance roll-out

Secure "at rest" disk encryption

Multimedia , RFID sensors data growth
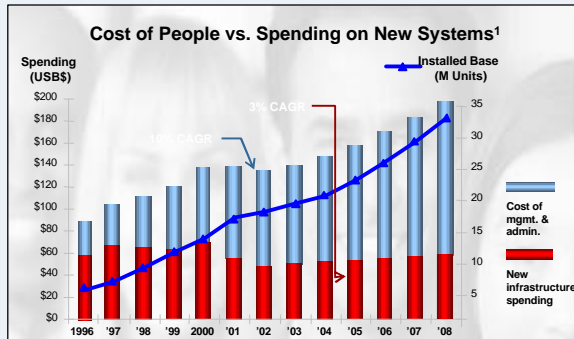
Trickle-down SMB function and focus

Virtualization adoption acceleration

Disk Technologies → SFF (2.5"), Perpendicular capacity (TB), Hybrids (NAND)

SNIA SMI-S and Aperi standards strengthen

Shared to dedicated DR solutions

Information a Strategic Asset

Cost of People vs. Spending on New Systems[1]

Spending (USB$)

Installed Base (M Units)

3% CAGR

10% CAGR

Cost of mgmt. & admin.

New infrastructure spending

$200 / $180 / $160 / $140 / $120 / $100 / $80 / $60 / $40 / $20 / $0

35 / 30 / 25 / 20 / 15 / 10 / 5

1996 '97 '98 '99 2000 '01 '02 '03 '04 '05 '06 '07 '08

[1] IDC 2004 Source: IBM Customer Surveys/Feedback

"During 2002 the world created 5 exabytes of information. An exabyte is the digital equivalent of a trillion novels. The yearly total amounts to somewhat more than all the words ever spoken by human beings – and it was up 68% from two years earlier" –

Forbes: Data of Reckoning, May 10, 2004

IBM System Storage - Strategy
Intelligent Management. Protected Information. Smarter Insights.

Leverage Information
✓ Capitalize on data sharing for collaboration
✓ Align storage investment, information value
  ‣ Tivoli Storage Manager
  ‣ TSM Space Manager

Mitigate Risk
✓ Comply with regulatory, security requirements
✓ Keep your business running continuously
  ‣ TotalStorage Productivity Center for Replication
  ‣ XRC, GDPS

Optimize IT
✓ Automate and Simplify IT operations
✓ Optimize Performance, Functionality
  ‣ DS family
  ‣ FlashCopy, Metro Mirror, Global Mirror
  ‣ TS family
  ‣ Virtual Tape Server

Enable Business Flexibility
✓ Flexible, On Demand IT infrastructure
✓ Protect your IT investment
  ‣ SAN Volume Controller
  ‣ TotalStorage Productivity Center
  ‣ Tivoli Provisioning Manager

The IBM System Storage DS Family

Industry's broadest range of disk storage systems

DS6000

DS8000

- **Enhanced scalability, performance and affordability with IBM TotalStorage DS8000 and DS6000 Series disk systems**
- **Support for enhanced data protection and application availability with leadership products for replication, mirroring, and disaster recovery**
- **Support for greater productivity through infrastructure consolidation / simplification**
- **Support for lower costs through better management of data throughout its lifecycle**
- **Solutions designed for the On Demand enterprise – IBM SAN, Disk, Tape, Software, and Services**

On September 13, 1956, IBM introduced the industry's first disk system, the **IBM 350** disk storage unit for the IBM RAMAC 305.

– 5 million 7-bit characters (approximately 5MB)
– Across 50 disk platters (61cm diameter)
– Was leased for $3,320 USD per month, or purchased outright for just under $50,000 USD.

IBM System Storage DS8000 series

DS8000 Turbo

- **Setting a New Standard in Cost Effectiveness**
  - *Performance* **– Up to 7X ESS Model 800**
  - *Scalability* **– Up to 512TB physical capacity**
  - *Virtualization* **– Storage System LPARs**
  - *Flexibility* **– FC and FATA disks fit cost/access**
  - *Extendibility* **– Extended addressing capability**
  - *Storage Management* **–Full set of interfaces/tools**
  - *Availability* **– Designed for 24x7 environments**
  - *Resiliency* **– Industry leading advanced functions**
  - *Long Term Cost* **– 1,2,3,4 year warranty, model upgrades**
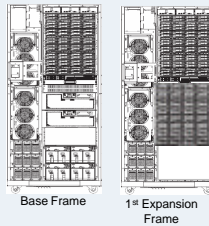 - Over 5000 Shipped

BENEFITS
- Provides most robust remote copy options to safeguard critical business data and operations
- Consolidates hundreds of terabytes to reduce infrastructure complexity and cost
- Integrates FC and SATA disk tiers to lower storage residency cost
- World class performance enhances business transaction speeds and competitiveness
- Reduces operational costs for spiraling data growth with flexible warranty options and no charge software
- Synergistic with System p and System i; provides reference support for System z

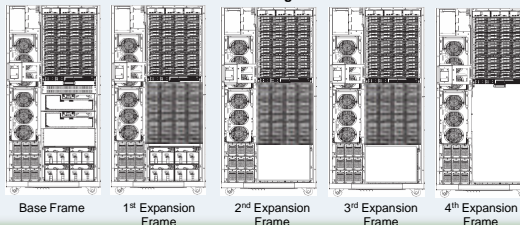# DS8000 Storage Unit - New Greater Scalability

Enables up to 1024 drives, 512TB of physical capacity in a single DS8000 system

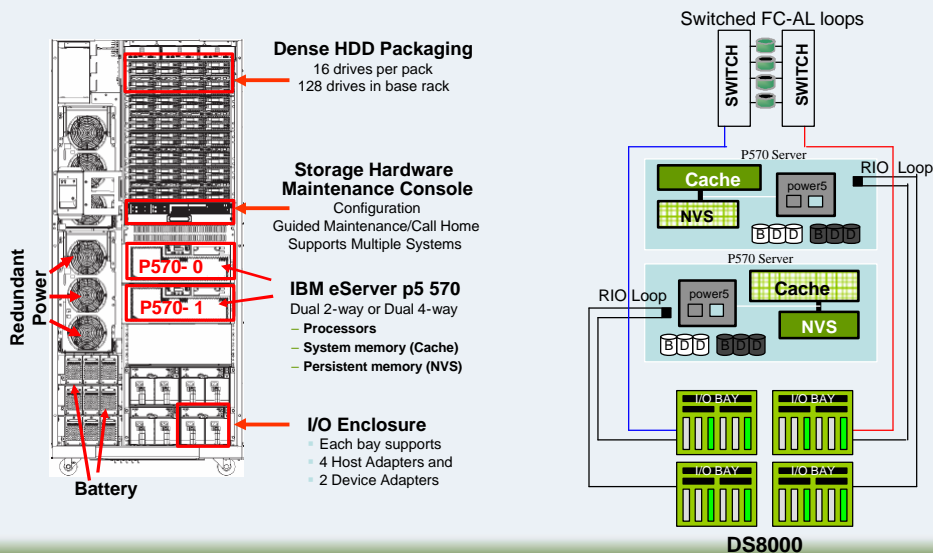| Server Processors	2-way POWER5	4-way POWER5
Cache	16 to 128 GB	32 to 256 GB
Host Ports		
FCP/FICON® (4Gb/s) (4 ports per adapter)	8 to 64	8 to 128
ESCON® (2 ports per adapter)	8 to 32	8 to 64
Drives 73 GB (15K RPM) 146 GB (10K RPM) 300 GB (10K RPM) 300 GB (15K RPM) 500 GB (7.2K RPM)	16 to 384	up to 1024
Physical Capacity	1.2 to 192TB	up to 512 TB
Number of Frames	1 to 2	1 to 5

DS8100 Storage Unit

Base Frame / 1st Expansion Frame

DS8300 Storage Unit

Base Frame / 1st Expansion Frame / 2nd Expansion Frame / 3rd Expansion Frame / 4th Expansion Frame

---

IBM DS8000 High Availability Design

Dense HDD Packaging
16 drives per pack
128 drives in base rack

Storage Hardware Maintenance Console
Configuration
Guided Maintenance/Call Home
Supports Multiple Systems

Redundant Power

P570- 0
P570- 1

IBM eServer p5 570
Dual 2-way or Dual 4-way
– **Processors**
– **System memory (Cache)**
– **Persistent memory (NVS)**

I/O Enclosure
▪ Each bay supports
▪ 4 Host Adapters and
▪ 2 Device Adapters

Battery

Switched FC-AL loops

SWITCH   SWITCH

P570 Server
Cache   power5
NVS
DDD   DDD
RIO Loop

P570 Server
power5   Cache
NVS
DDD   DDD
RIO Loop

I/O BAY   I/O BAY
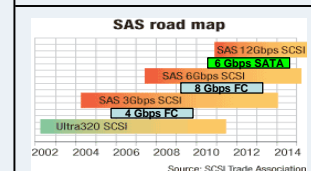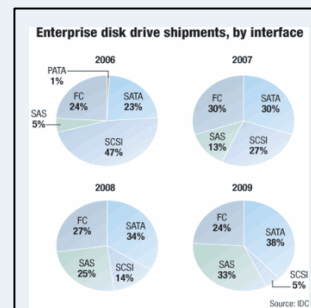I/O BAY   I/O BAY

DS8000

4

Hard Drive Recommendations – DS8000

- **73 GB 15K RPM Disk**
  - Provides significant service time, throughput improvements
  - *Recommended for best performance and cache hostile OLTP*
- **146 GB 10K Disk**
  - Best choice for most customers running workloads with average access density
  - *Best price/performance/throughput balance*
- **146 GB 15K Disk**
  - Significant (30% to 50%) array throughput benefit over 146 GB 10 KRPM disks
  - *Recommended for above average access densities requiring above average performance*
- **300 GB 10K and 300 GB 15K Disk**
  - Use with caution for high access density or with cache unfriendly workloads
  - Consider using as Remote Copy secondary for Disaster Recovery
  - *May provide good price/performance for RAID 10 configurations*
- **500 GB 7.2K Disk**
  - *Use with caution for low access density or archival unfriendly workloads*
  - *Unsuited for OLTP or other fast transactions OK for tier 2 or tier 3 storage*


Disk Drive Evolution

- **SAS (Serial Attached SCSI) is a new serial, point-to-point topology connecting SAS and SATA disks – will replace FC-AL, SCSI**

- **Dominant growth: SATA and SAS disks**

- **FC, SAS, and SCSI disks use dual port, 10K/15K RPM, fast actuator hardware with different interfaces (Fiber, Serial, Parallel)**

- **FATA and SATA disks use 7.2K RPM, slower actuator hardware with different interfaces (Fiber, Serial)**

- **Emergence of 2.5" (SFF) small form factor disk (36GB and 73GB, 10K)**
  - With better performance, packaging, lower power and heat
  - About 50% of shipped disks by 2009
  - Some 15K drives, 10K dominant

Enterprise disk drive shipments, by interface

2006	2007
PATA 1%	
SAS 5%	
FC 24%	FC 30%
SATA 23%	SATA 30%
SCSI 47%	SAS 13% / SCSI 27%

2008	2009
FC 27%	FC 24%
SATA 34%	SATA 38%
SAS 25% / SCSI 14%	SAS 33% / SCSI 5%

Source: IDC

SAS road map

SAS 12Gbps SCSI
6 Gbps SATA
SAS 6Gbps SCSI
8 Gbps FC
SAS 3Gbps SCSI
4 Gbps FC
Ultra320 SCSI

2002 2004 2006 2008 2010 2012 2014

Source: SCSI Trade Association

Patent Number: US20040098541A1

**System and method for implementing an adaptive replacement cache policy**

- **Best caching algorithms in industry**

- **Over 20 years experience**

- **Simplified Adaptive Replacement Cache (SARC)**
  - *Self-Learning algorithms*
    - ✓ **Adaptively and dynamically learn what data should be stored in Cache based upon the recent access and frequency needs of the Hosts**
  - *Adaptive Replacement Cache*
    - ✓ **Most advanced and sophisticated algorithms to determine what data in Cache is removed to accommodate newer data**
  - *Pre-fetching*
    - ✓ **Predictive algorithm to anticipate data prior to a host request and loads it into Cache**

- **Benefits**
  - *Leading performance*
    - ✓ **Been proven to improve Cache hit by up to 100% over previous IBM caching algorithms and improve I/O response time by 25%**
  - *More efficient use of Cache*
    - ✓ **Intelligent caching algorithm profiles Host access patterns to determine what data is stored**
    - ✓ **Need less Cache than competitors**
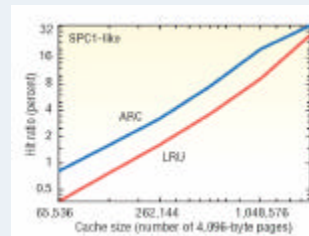
Nimrod Megiddo and Dharmendra S. Modha, "Outperforming LRU with an Adaptive Replacement Cache Algorithm," *IEEE Computer*, pp. 4-11, April 2004.

What is AMP?

- **A breakthrough caching technology from IBM Research called AMP (Adaptive Multi-stream Prefetching)**
  - – Can dramatically improve performance for common sequential and batch processing workloads on the IBM DS8000 with the 2.4G LIC release (July 2007).

- **AMP optimizes cache efficiency by incorporating an autonomic, workload-responsive, self-optimizing prefetching technology .**
  - – The algorithm dynamically decides what to prefetch and when to prefetch
  - – Delivers up to a two-fold increase in the sequential read capacity of RAID-5 arrays.
  - – The bandwidth for a fully configured DS8100 or DS8300 remains unchanged,
  - – May improve sequential read performance for smaller configurations and single arrays
  - – Reduces the potential for array hot spots due to extreme sequential workload demands.
  - – May significantly reduce elapsed time for sequential read applications constrained by array bandwidth such as BI and critical batch processing workloads.
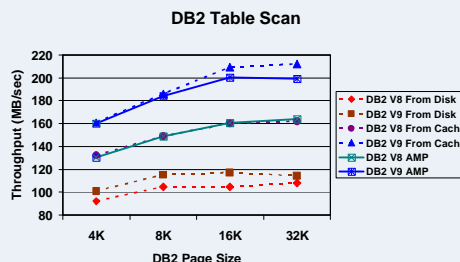
Reference Paper: *AMP: Adaptive Multi-stream Prefetching in a Shared Cache by Binny Gill and Luis Angel D. Bathen*, USENIX File and Storage Technologies (FAST), February 13-16, 2007, San Jose, CA.

6

DS8000 and System z9 Enable Performance

- **IBM introduces a combined server/software/storage configuration designed to offer fast access and low processing time for:**
  - Applications that use Media Manager to access data, such as DB2® tables or other VSAM[1] data sets. No application changes are required to benefit utilities that support these data, including DB2 Utilities for Image Copy and Recovery
- **IBM solution for this high-performance data repository includes:**
  - IBM System z9 EC running z/OS®, enabled for MIDAW feature
  - FICON Express2 or FICON Express4 host adapter connectivity
  - IBM System Storage DS8000[2] support for System z9 MIDAW
- **MIDAW (Modified Indirect Data Address Word) is a system architecture and software exploitation for IBM System z9 with z/OS 1.6 and above**
- **Designed to offer improved performance[3]**
  - Example Results: IBM DS8000 + z9 EC - DB2 table scan test with MIDAW + FICON Express4 - compared to FICON Express2 configs without MIDAW
  - Up to 220% improvement in I/O throughput (read MB/s)

Note: 1. EF VSAM, PS-E, PDSE, HFS and zFS for z/OS 1.6 with APAR. All VSAM, PS-E, PDSE, HFS and zFS for z/OS 1.7 release. Most improvement for EF data with small block size
2. IBM ESS 800 & DS6000 also support System z9 MIDAW feature and may be used as alternatives, but may have less benefit.
3. 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 numbers stated here.

DS8000 - Benefits of HyperPAV

Reduce number of required aliases

– Give back addressable device numbers

– Use additional addresses to

- support more base addresses
- larger capacity devices.

z/OS can react more quickly to I/O loads

– React instantaneously to market open conditions

Overhead of managing alias exposures reduced

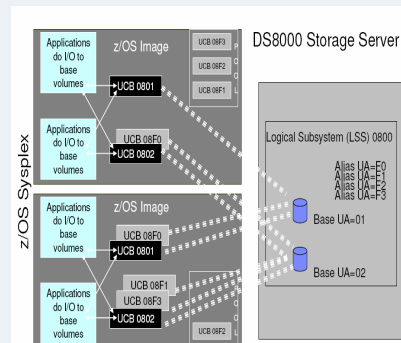– WLM not involved in measuring and moving aliases

– Alias moves not coordinated throughout Sysplex

Initialization doesn't require "static" bindings

– Static bindings not required after swaps

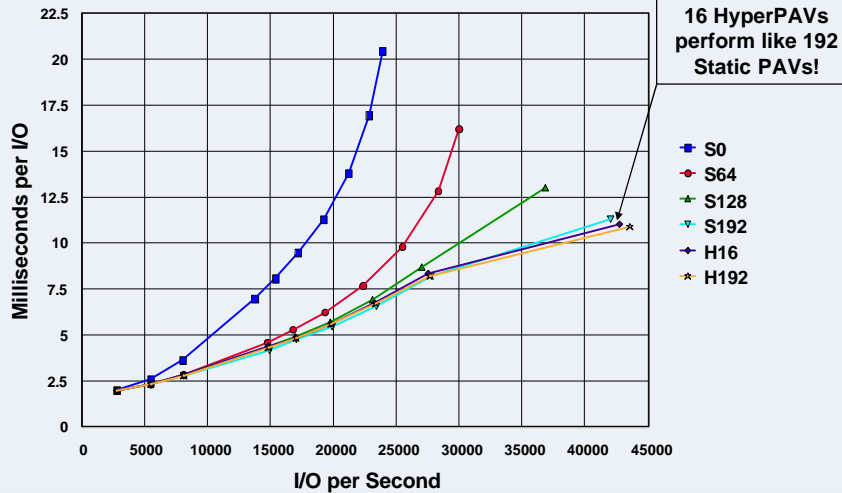**IO reduction, no longer need to BIND/UNBIND to manage HyperPAV aliases**

Increases I/O Parallelism

Static PAV vs HyperPAV – 64 base volumes
DB z/OS Workload (Cache Standard)

16 HyperPAVs perform like 192 Static PAVs!

Y-axis: Milliseconds per I/O (0, 2.5, 5, 7.5, 10, 12.5, 15, 17.5, 20, 22.5)

X-axis: I/O per Second (0, 5000, 10000, 15000, 20000, 25000, 30000, 35000, 40000, 45000)

Legend:
- ■ S0
- ● S64
- ▲ S128
- ▼ S192
- ◆ H16
- ★ H192

DS8000 Exploitation of Storage System LPARs

The Storage System LPAR Advantage!

- **Reduction in**
  - **Floor space, power requirements and cooling requirements, through consolidation of multiple stand-alone storage functions**
  - **Systems management effort**
  - **Implementation effort through LPAR integration, test, pre- installation and   self-configuration**
  - **Storage infrastructure complexity through integration**

- **Hardware based implementation ensures data integrity**

- **Heterogeneous workload support**

- **Efficient use and sharing of system resources through exploitation of dynamic LPAR**

- **Ability to leverage Storage Subsystem LPARs for the purpose of testing**

- **Reduction in physical asset management through consolidation into DS8000 LPARs**

DS8300 with LPARs

8
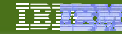
Low-end Mainframe Attach with the DS6800

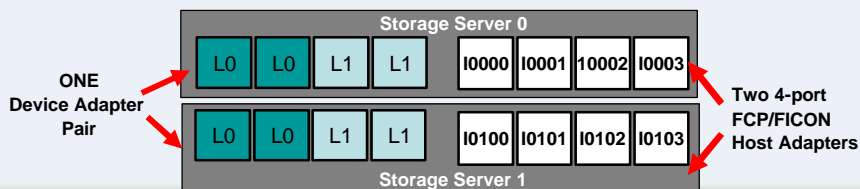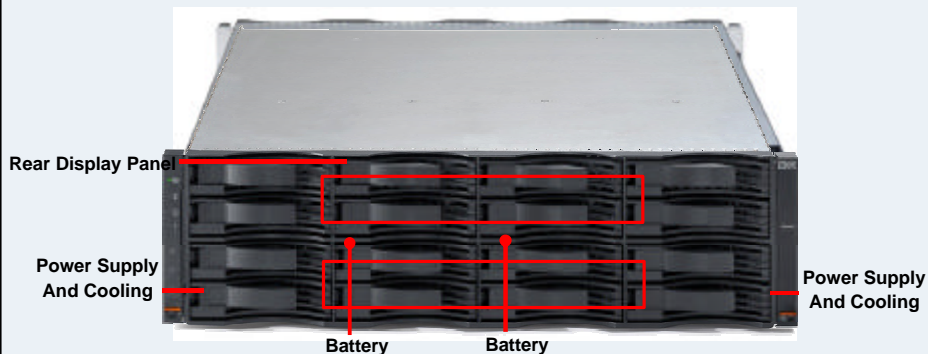- **Value Proposition:**
  - **Affordable zSeries storage solution**
  - **Save customers money in scenarios where the DS6000 fits**
  - **Customers can now afford new technology instead of purchasing older used equipment**
    - Leverage new technology instead of purchasing or retaining older, used hardware
  - **Affordable zSeries solution that fits in standard 19" rack**
  - **Outstanding TCO for customers' whose specific capacity and performance needs are met**

- **DS6000 Models - 522/EX2**
  - **3U package for controllers and 7 disk expansion units**
  - **High storage density footprint -16 drives per 3U package**
  - **Maximum of 1024 host and 8192 LUNs (2TB max)**
  - One year, same-day, 24x7 IBM onsite repair (IOR) warranty
  - 500 GB 7,200 rpm FATA disk drives
    - **Provide additional price, and capacity flexibility to help address specific application and business requirements**
    - **Intermix is at the drawer level**
  - TPC Replication Manager Support
  - CALL HOME SUPPORT VIA TELEPHONE LINE
   - **Provides the ability to E-Mail a customer at Call Home**

DS6800 High Availability Design

Rear Display Panel

Power Supply And Cooling

Power Supply And Cooling

Battery   **Battery**

Storage Server 0

| L0 | L0 | L1 | L1 | | I0000 | I0001 | I0002 | I0003 |

ONE Device Adapter Pair

| L0 | L0 | L1 | L1 | | I0100 | I0101 | I0102 | I0103 |

**Two 4-port FCP/FICON Host Adapters**

Storage Server 1

Preferred Path Data Operations

- Chooses the most efficient and optimum path to store and retrieve data from the storage systems
  - **Beyond load balancing**
  - **Our system selects the best path at each stage of the request**
- *Competitors – can do preferred path, but path is selected at time of initial request*
- Preferred Path
  - **Dynamically selects the most efficient and optimum path to use at each data interchange during read and write operations**
  - **A level above load balancing used by competition**

HBA HBA
(SDD)
Path determined by LUN
Path determined dynamically
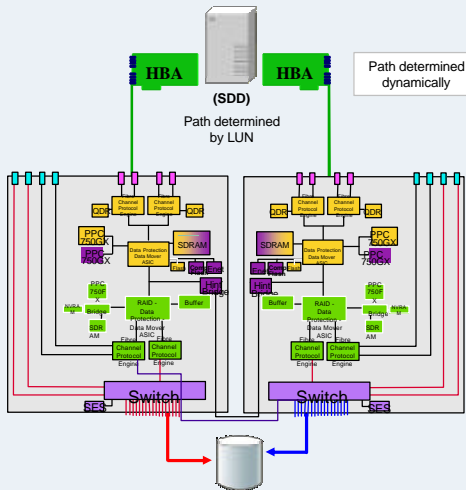
DS6800 Specifications

DS6800	
Controllers	**Dual Active**
Max Cache	**4 GB**
Max Host Ports	**8-Ports; 2Gb FC/FICON**
Max Hosts	**1024**
Max Storage / Disks	**128**
Disk Types	**FC 10K: 146GB, 300 GB** **FC 15K: 73GB, 146GB, 300GB, FATA 7.2:** **500 GB**
Max Expansion Mod	**7**
Max Disk Loops	**4 (2 dual redundant)**
Max LUNs	**8192** **(up to 2 TB LUN size)**
RAID Levels	**5, 10**
RAID Array Sizes	**4 or 8 drives**
Operating Systems	**z/OS, i5/OS, OS/400, AIX, SUN Solaris, HP UX, VMWare, Microsoft Windows, Linux**
Packaging	**3U – Controller & Expansion Drawers**
Power consumption	**Controller: 0.69 kVA** **Expansion drawer: 0.48 kVA**
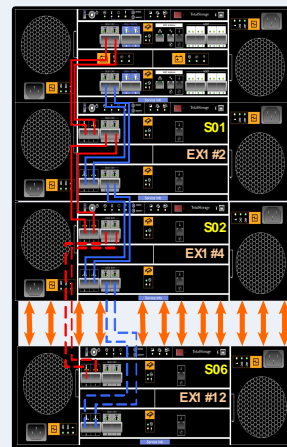
IBM DS8000/DS6000 and Business Continuity

FlashCopy

Can help reduce or eliminate planned outages for critical applications. FlashCopy is designed to provide the same point-in-time copy capability for logical volumes on the DS8000 and DS6000 series.

- **Volume level FlashCopy (System z servers)**
- **Dataset FlashCopy (System z Servers)**
- **Multiple Relationship FlashCopy**
- **FlashCopy Nocopy**
- **Incremental FlashCopy**
- **Fast Reverse Restore**
- **FlashCopy consistency groups**

Metro Mirror

Synchronous mirror of data volumes to a remote location for DS6000, DS8000

- **What is Metro Mirror used for?**
  – For applications that need zero data loss, Metro Mirror provides the tightest Recovery Point Objective
  – About 1 ms delay per 100km

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

Less Than 300 KM
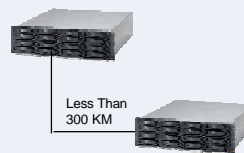
Global Mirror

Asynchronous disk mirror at any distance with continuous data integrity. Its performed at a volume level

- **What is Global Mirror used for?**
  – Do not impact Production I/O performance when sufficient bandwidth and / or other resources are unavailable**.**

  – Creates tertiary (data consistent) copy every 3-5 seconds using incremental FlashCopy

Greater Than 300 KM

FlashCopy
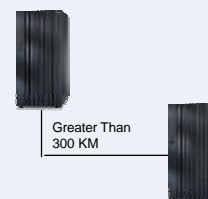
SAN Volume Controller Delivers Value

- **Reduces the cost & complexity of managing storage**
  - Creates tiers of storage
  - Enables multi-vendor strategies

- **Improves Business Continuity**
  - Change storage without interrupting applications
  - Allocate more storage to applications automatically

- **Improves Storage Utilization**
  - Combines storage capacity into a single resource
    – from multiple vendors

- **Improves Personnel Productivity**
  - Manage a single storage resource from a central point

- **Improved application availability**
  - Eliminate many of the causes of storage-related downtime
  - Designed to migrate data without disruption

- **Designed to ease deployment of tiered storage**
  - Helps use lower-tier storage for greater range of applications

Host

VDisk

MDG MDG MDG

SAN Switch   SVC Node Pair   SAN Switch

11
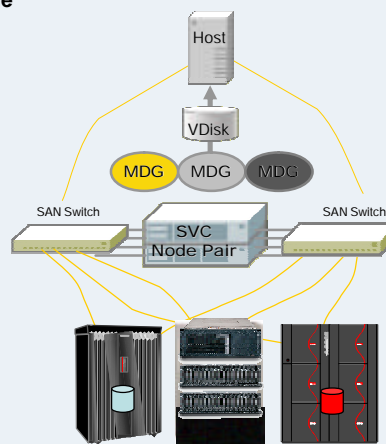
SVC High Availability Architecture Design

SVC (2145 8G4 engine) based on IBM System 3550
– *Two dual-core Intel Xeon 5160 processors at 2.33GHz*
– *8GB of cache per Node*
– *Four 4Gbps FC ports per Node*
– *Pre-loaded Virtualization Software*

Uninterruptible Power Supply (UPS)
– *Two 1U per cluster connected to a separate branch circuit*

Service Engine
– *Master Console for Management*
– *ICAT SMIS and GUI Webserver*
– *RAS Features – Remote Support, Call Home*
– *Two 4Gbps FC ports*

Setting a New Standard in Cost Effectiveness
– *Performance* – SVC 4.2 delivers 272,500 SPC-1 IOPS and 7080 SPC-2 MBs
– *Scalability* – SVC cluster is scalable up to 8 nodes
– *Flexibility* – FC and FATA disks fit cost/access
– *Storage Management* – Full set of interfaces/tools
– *Resiliency* – Industry leading advanced functions

SVC Node Pair | Management Node | DS6800 | UPS Pair | 19" Rack Mounted

---
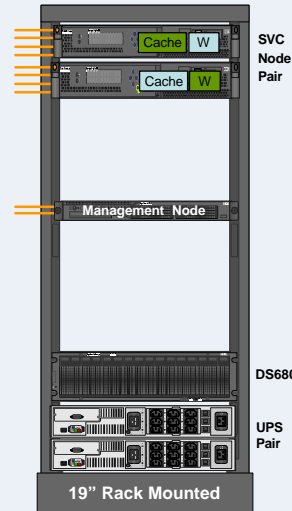
Example Migration to a SAN Volume Controller

To migrate data in a production environment using the SVC Image Mode based method the following steps are performed:

1. *Host:* Uninstall the old drivers and shut down. The maintenance window begins for the host and applications residing on it.
2. *Switch:* Change zoning so the host only sees SVC and SVC sees the storage controller.
3. *Storage controller:* Un-map disks from the host. Map disks to the SVC.
4. SVC: Define the host and the disks.
5. Host: Install SDD; disks should be visible as before. The maintenance window ends.
6. *Rediscover:* LUN to Host

▪ **It is easy**: No advanced features are needed from storage controllers.
▪ **It is simple**: Once the LUNs on the storage controller are presented to the host system through SVC, the rest of the procedure is always the same no matter which storage controller it is.

Host | VDisk | MDG MDG MDG | SAN Switch | SVC I/O Group Cache WC | SAN Switch | Storage | Storage

12
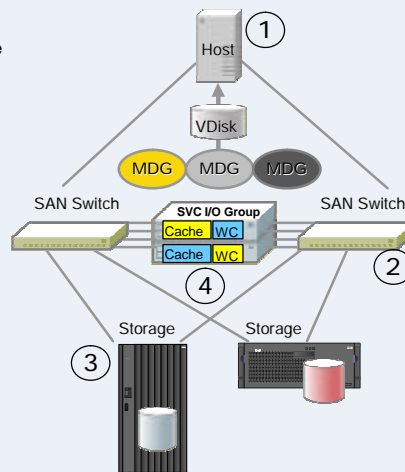
SAN Volume Controller: Did You Know?

IBM has shipped **over 9000 SVC engines** running in more than **3100 SVC systems**

There are more than **130 customer references** for SAN Volume Controller

- IBM has **40 years experience** in virtualization technologies

- Mature, stable solution – 4+ years

- Enterprise class reliability

- Simple to install, easy to use

- SAN Volume Controller delivers the **FASTEST** storage performance benchmark ever recorded for ALL controllers

- **15+ PB** of client data managed today and growing!

- SAN Volume Controller can virtualize IBM and **non-IBM** storage (over 75 systems from EMC, HP, HDS, Sun, Dell, STK, NetApp)
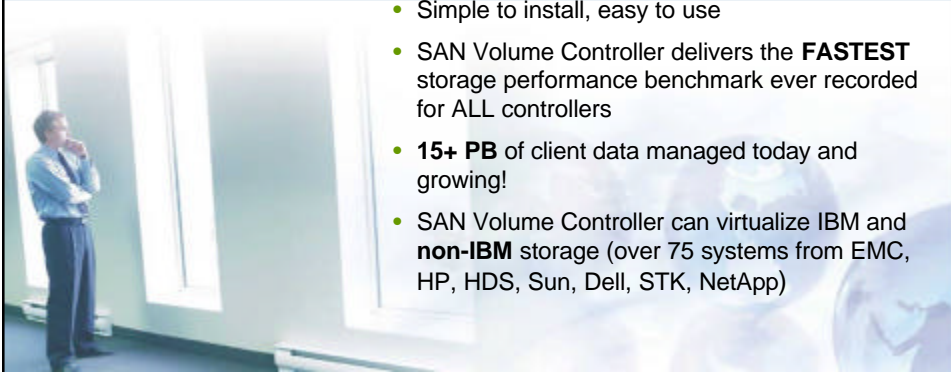
SVC Business Continuity Features

IBM SVC FlashCopy

? Target is time-zero copy of the source

? Up to 3855 Source virtual disk to Target virtual disk available (almost) immediately

? Up to 256 Consistency groups supported

? Target may be updated independently of the source

IBM SVC Metro Mirror

? Synchronous remote copy between virtual disks geographically dispersed on two SVC clusters

? **Intra-cluster** and **Inter-cluster** remote copy is supported

? SVC supports Consistency Group at the level of the **cluster**

IBM SVC Global Mirror

? Long distance asynchronous remote mirroring function
  - Practically unlimited distances for business continuity

? Does not wait for secondary I/O before completing host I/O
  - Minimizes performance impact to applications
  - Supports consistency groups which span IO groups within a cluster

? Designed to maintain consistent secondary copy at all times

SAN

SAN

SAN Volume Controller

SAN Volume Controller

RAID controller 1   **RAID controller 2**

RAID controller 1   **RAID controller 2**
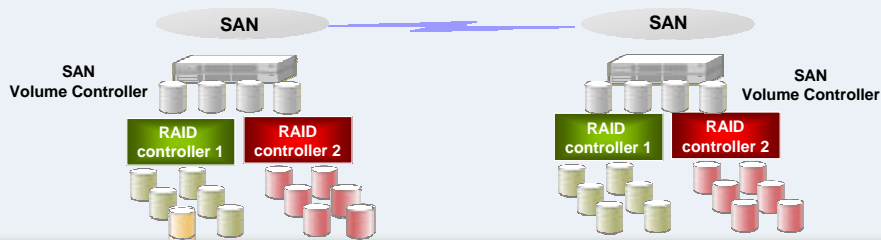
System Storage SAN Volume Controller v4.2
Supported Environments

NEW!

Host Platforms   **1024 Hosts (256 per I/O Group)**

Intel/AMD x86,x64,ia64 | POWER ™ | N Series / NetApp ® | SPARC ® | Alpha | PA-Risc ® | MIPS ® | System z ™

IBM AIX
HACMP/XD,GPFS,VIO
Linux
Redhat/SLES with LVM
HP/UX
ServiceGuard Cluster
Tru64 UNIX
OpenVMS
Sun Solaris
VCS Cluster, Sun Cluster
Microsoft Windows
MSCS,MPIO,VSS,GDS
x64, ia64
Novell Netware
Cluster
SGI IRIX
VMWare
Via RPQ
zVM, zVSE

McData (CNT)   **SAN**   Brocade Cisco

SAN

4096 vdisks (1024 per I/O Group)

Continuous Copy
Metro Mirror
Global Mirror
Point-in-Time Copy
Full volume
Multiple targets

IBM SAN Volume Controller
4 I/O Groups (8 nodes)

NEW!

Cache Disabled Virtual Disk
Enables Array-based Copy Services

Intended as an overview only.
For the most complete information,
visit **ibm.com**/storage/software

Disk Systems

4096 mdisks from up to 64 storage controllers

IBM	**EMC**	**HDS**	**HP**	**Sun/STK**	**NEC**	**NetApp**	**Bull**	**Fujitsu**
ESS	Symmetrix	Lightning	StorageWorks	StorEdge	iStorage	Filers	StorWay	Eternus
DS	Clariion	Thunder	(XP,EVA,	FlexLine		Gateways		
FAStT		Tagmastore	MA/EMA,MSA)	D-Series				
N Series		AMS, WMS						

NEW! (NEC)   NEW! (Bull)   NEW! (Fujitsu)

© IBM Corporation 2007   2007 System z Technical Conference

SVC Delivers ROI Benefits for System z

- **Support for Linux operating environments implemented on IBM System z9 and zSeries platforms with storage managed by SAN Volume Controller. This capability is intended to:**

  – Enable System z9 and zSeries servers to utilize open systems storage pools created by SAN Volume Controller

  – Create a tiered storage environment for Linux operating environments in System z9 and zSeries environments

  – Simplify volume management by combining mainframe and open systems storage environments into a single pool, managed from a central point

  – Expand the host environments managed by SAN Volume Controller

  – Improve application availability by enabling changes to the storage infrastructure without disrupting the applications running on the hosts

© IBM Corporation 2007   2007 System z Technical Conference
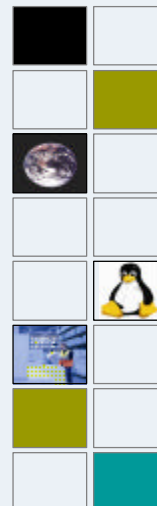
Data Migration – Different Shapes and Sizes

Technology Refresh

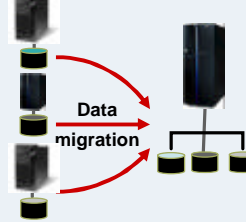Server/Storage Consolidation

- **Low Impact**
  - **Architecture independent**
  - **File system independent**
  - **Storage manufacturer independent**
  - **Local or remote**
  - **Migrates Volumes, Disk Partitions or Files**
  - **Allows physical formats and LUN sizes to be changed**
  - **Can be used to Migration to New Volume Manager Configurations**
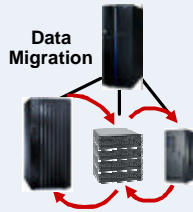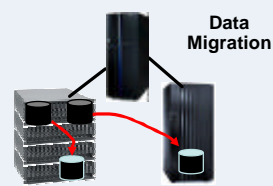 - **Can be used with Existing Networks**

Old Array → **Data Migration** → New Array

Data migration

Facilitate Tiered Storage

Data Migration

Improve Performance

Data Migration

TDMF (IP) Architecture overview

File system

File system

TDMF driver

Buffer, for async staging

Volume manager

TCP/IP

Volume manager

Persistent store (pstore)

Mounted source volumes

Unmounted target volumes
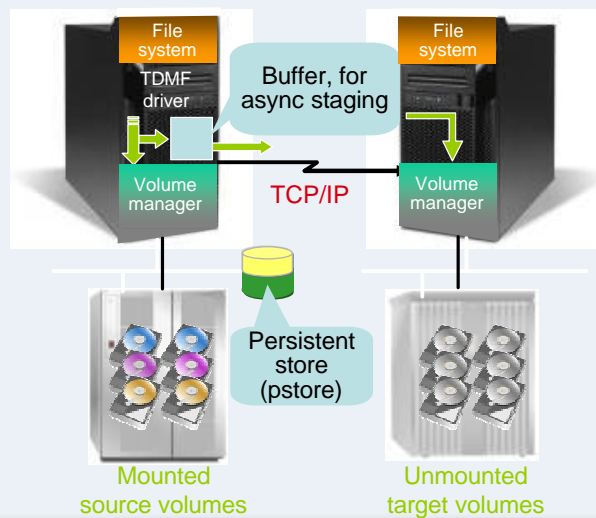
Over 50 Years of Tape Innovation

- **Tape is an integral part of the storage hierarchy**
- **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**

Starting in 1952 - IBM 726

1848
George Boole
invents
binary algebra

1928, IBM invents
the 80 column
Punch Card

1952
IBM 726
1st magnetic tape drive

1959
IBM 729
1st read/write drive

1964
IBM 2401
1st read/back drive

1984
IBM 3480
1st cartridge drive

1995
IBM 3590

1999
IBM 3590E

E N T E R P R I S E

2000
LTO Gen1

2002
LTO Gen2

2004
LTO Gen3

O P E N

2003
3592 Gen1

2005
3592 Gen2
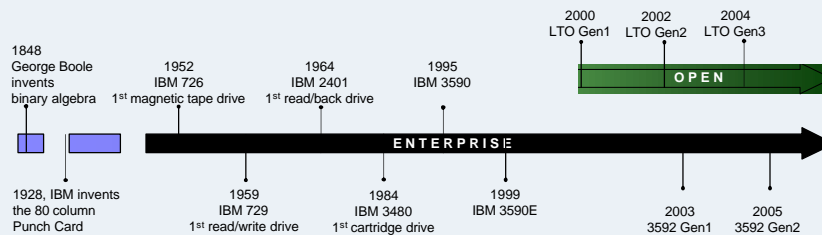
IBM System Storage Tape

Tape Storage Continuum

TS1030
(LTO4)

TS1120
(3592)

TS3200

TS3100

TS3310

3494

TS3500

TS7510
(distributed)

TS7740
(mainframe)

Tape Drives	Tape Libraries	Virtualization

LTO4 tape drive
– Encryption capable
– Supports up to 800 GB cartridge
– Up to 120 MB/sec throughput
 	
TS1120 tape drive/controller
– Second generation tape drive
– Controller supports ESCON & FICON
– Tape drive data encryption
– 100, 500 and 700GB cartridge capacity
– Up to 104 MB/sec throughput
– Auto Virtual Backhitch
	TS3100 tape library (up to 17.6TB[1])
	TS3200 tape library (up to 35.2TB[1])
	TS3310 tape library (up to 316.8TB[1])
	– Stackable modular design
	TS3400 tape library (up to 38TB[1])
	TS3500 tape library (up to 16.5PB[1] with LTO4 or up to 13.1PB[1] with TS1120)
	– Linear, scalable, balanced design
	– High Availability
	– Fastest robotics in industry
	– LTO and TS1120 tape drive
		VTO (Virtual Tape Open)
		– Up to 4.8 GB/sec throughput
		– Up to 884 TB cache capacity
		
		Hydra (Virtual Tape Server)
		– Up to 900 MB/s throughput
		– 6 TB native cache
		  • 12 TB with GRID (PtP)
		– Standalone or GRID (PtP)
		  • Synchronous data replication
		  • Third site support in plan
		  • GDPS support
		– Advanced cache management

IBM Tape Drives
What determines Best of Breed?

Technology Competencies

IBM LTO Drives

MR Heads
Media
Mechanics
Electronics / Servo
Recording Technology
Error Correction
Dynamic Compression Look Ahead
Speed Matching
Attach Interface Architecture
Dynamic Load Balancing
Native Fibre Channel
SARS (**S**tatistical **A**nalysis **R**eporting **S**ystem)

50+ Years of Tape Technology

IBM 3592 Technology

IBM Ultrium 4

- **800 GB Native Physical Capacity** (1.6 TB compressed) on LTO Ultrium 4 media

- Up to **120 MB/s** native data transfer rate

- 4Gbps Fibre Channel, Ultra160 LVD SCSI* and/or **New 3Gbps SAS**** attach

- **Encryption capable** for LTO4 SAS and Fibre Channel

- Digital Speed Matching (30, 48, 66, 84 103, 120 MB/s)

- 256 MB Internal Buffer (128 MB for IBM Ultrium 3)

- Several continued features/functions from IBM Ultrium 3
  - WORM technology
  - Dual stage 16-channel head actuator
  - Independent tape loader and threader motors
  - Graceful dynamic braking
  - SARS (Statistical Analysis and Reporting System) and ECC (Error Correction Code)
  - Same 5 ¼" form factor
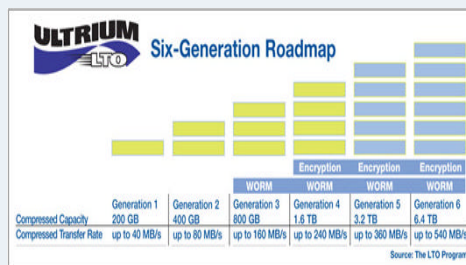
*Available only for TS2340, TS3100, and TS3200
**Not available with TS1040 (TS3500)

ULTRIUM LTO Six-Generation Roadmap

	Generation 1	Generation 2	Generation 3	Generation 4	Generation 5	Generation 6
				Encryption	Encryption	Encryption
			WORM	WORM	WORM	WORM
Compressed Capacity	200 GB	400 GB	800 GB	1.6 TB	3.2 TB	6.4 TB
Compressed Transfer Rate	up to 40 MB/s	up to 80 MB/s	up to 160 MB/s	up to 240 MB/s	up to 360 MB/s	up to 540 MB/s

Source: The LTO Program

LTO Cartridge Interchangeability

		Gen 1	Gen 2	Gen 3	Gen 4
	Capacity	**100GB**	**200GB**	**400GB**	**800GB**
Gen 1 Drive	**Read**	**15 MB/sec**			
	Write	**15 MB/sec**			
Gen 2 Drive	**Read**	**15 MB/sec+**	**35 MB/sec**		
	Write				
Gen 3 Drive (*Gen 3 HH*)	**Read**	**15 MB/sec+**	**35 MB/sec+**	**80 MB/sec** (*60 MB/s*)	
	Write				
Gen 4 Drive	**Read**		**35 MB/sec+**	**80 MB/sec+**	**120 MB/sec**
	Write				

New IBM TS2340 Tape Drive Express Model
(Machine Type 3580, Model L43, S43)

- **Single IBM Ultrium 4 Full High External Drive**

- **120 MB/sec native data transfer rate**

- **Single Ultrium cartridge**
  – 800 GB native physical capacity (1.6 TB with 2:1 compression)

- **Server Top Packaging or Rack Mount Shelf**

- **LED Display and Indicator**

- **Two models:**
  – L43 - Ultra160 LVD SCSI attachment
  – S43 – New 3Gbps dual port SAS (Serial Attached SCSI) attachment

- **3-year warranty in most countries (customer replacement exchange)**

- **Available in HVEC (PNs 3580L4X and 3580S4X)**

- **Customer Setup**

The IBM 3592.... the 1TB Tape Technology

2006 SearchStorage Product of the Year

Drive Capacity

□ 3590 □ 3592
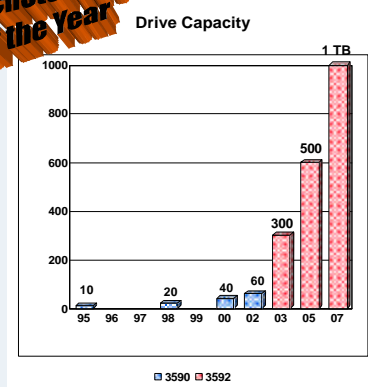
IBM TS1120 Highlights

- Enterprise class drive
- **ENCRYPTION!!**
- Open systems / mainframe
  – Dual port 4Gb fibre
  – FICON/Escon attachment
- 300/500/700GB native
- 100GB Scaled
  – WORM capability too
- 104 MB/s native
- 512MB Buffer
- Speed Matching
- Virtual Backhitch
- High speed data search
- High-resolution directory
- Media and drive health statistics
- Media reuse (read and write)

IBM ADVANTAGE!!!
One drive for both high-capacity and fast-access!!!
One drive for Encryption and non-encryption requirements!
One drive for mainframe and open.

IBM TS1120 Enterprise Tape Drive Roadmap

	3590 Generations			3592 Generations		
3590/3592 Tape Drive Models	Gen 1 3590 B	Gen 2 3590 E	Gen 3 3590 H	Gen 4 3592 JIA	Gen 5 3592 E05	Gen 6 3592 xxx
Native Capacity	10/20 GB	20/40 GB	30/60 GB	60 or 300 GB	500 or 700 GB	900 - 1100 GB* *with M media
Native Data Transfer Rate	9 MB/sec	14 MB/sec	14 MB/sec	40 MB/sec	100 MB/sec	100-160 MB/sec
Cartridge Type	J & K	J &K	J & K	L	L & M	L & M
Server Attachment	Ultra-SCSI ESCON FICON	Ultra-SCSI Fibre Channel ESCON FICON	Ultra-SCSI Fibre Channel ESCON FICON	2 GB Fibre ESCON FICON	4 GB Fibre ESCON FICON	10 GB Fibre ESCON FICON

Product Road Map: These statements represent IBM's current intent, are subject to change or withdrawal, and represent only goals and objectives.

Value of Virtual Backhitch

- **Higher Performance**
- **Higher Capacity**
- **Higher Reliability**
- **No Charge**
- **AUTOMATIC!!**

MBs/Sec

7
6 — 6
5
4
3
2
1 — 0.4
0

256 MBs

■ TS1120 ▨ 3590 Fibre

3592 Drive Portability

3494

Silo

3592

3584

TS3400

Rack
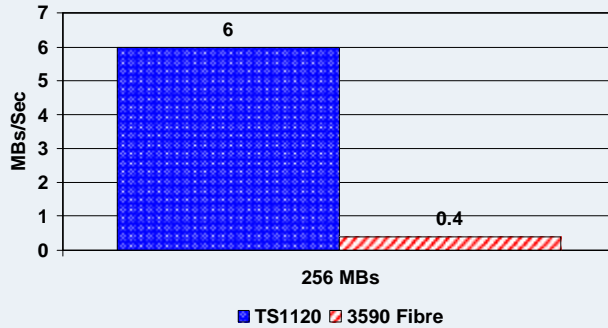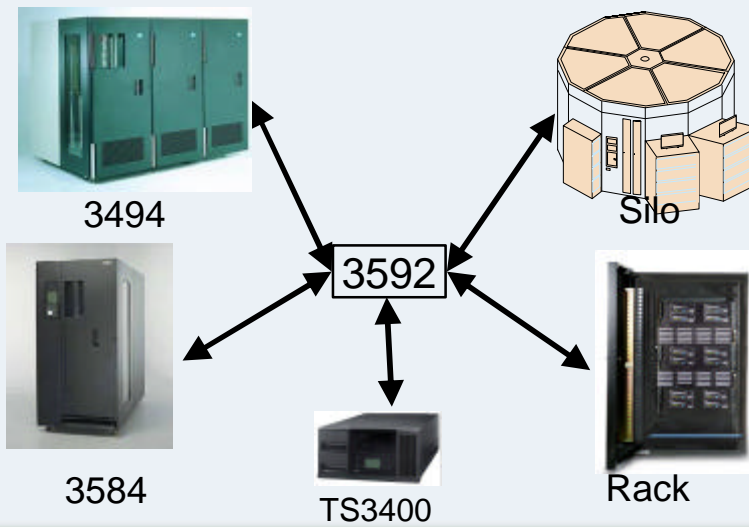
Function	Attribute	IBM LTO Generation 4	IBM 3592-E05 / TS1120
Performance	Native tape throughput	**120 MB/s**	104 MB/s
	Maximum sustained with compression	280 MB/s	280 MB/ s
	Drive buffer capacity	256 MB	**512 MB**
	Virtual backhitch support	No	**Yes**
Capacity	Native Cartridge Capacity	**800 GB**	700 GB
	Compressed cartridge capacity at 2:1	**1.6 TB**	1.4 TB
	Number of data tracks	896	896
Function	Speed Matching	6 speeds; **30 – 120 MB/s**	6 speeds; 35 – 104 MB/s
	Attachments	Single 4 Gbps fibre LVD Ultra 160 **3 Gbps SAS**	Dual 4 Gbps fibre **FICON[†]**
	System z support[†]	No	**Yes, FICON or ESCON**
	Encryption support	Symmetric	**Asymmetric**
	High resolution tape directory	No	**Yes**
	Scaled media support for fast access	No	**Yes**
	Longitudinal position sensing	No	**Yes**
	Reuse of old media at new capacity	No	**Yes; 300GB → 500 GB**
Mechanical Performance	Load to Beginning of Tape (BOT)	15 seconds	**13 seconds**
	Average File access time	54 seconds	**27 / 46 seconds**
	High speed search	8 meters / second	**10 meters / second**

Reasons Customers Virtualize Tape

1. To improve backup process
   - Address slow LAN backup
   - Address slow SAN server backup
2. To improve recovery time objective
   - Eliminate physical tape movement
   - Reduce queuing for resources
3. To improve recovery point objective
   - By creating incremental backups more frequently
   - By writing incremental backups to cache
4. To augment the existing business continuance infrastructure
   - Reduce bottlenecks and utilize tape assets more efficiently
   - Redeploy disk2disk pools and improve operational efficiency

Virtual Tape Server

IBM #1 since 1997

User's View

32-256 Virtual 3490E Automated Tape Drives

500,000 Virtual Volume Automated Storage

359x Highlights
4-18 Drives
Performance
Data Availability
Capacity
Attachment

Outboard Storage Subsystem View

Self-contained, self-managed, independent, sharable data server

Value
- ✓ *Target SystemZ marketplace*
- ✓ *90% Reduction in Cartridges*
- ✓ *30-50% Batch run improvement*
- ✓ *No JCL Changes*

Escon and/or Ficon

Automatic resource, performance, and content management

Systems Managed Tape

ACS Routines

Tape Data Sets

Logical View:

Data Class	Storage Class	Mgmt Class

Physical View:

Storage Group

System-Managed

Library Manager

3490 STD.	3490 ENH.
3590 3592	3590 3592

10-700GB

Native

VTS

400-4000MB

Virtual Vols

Automated Systems
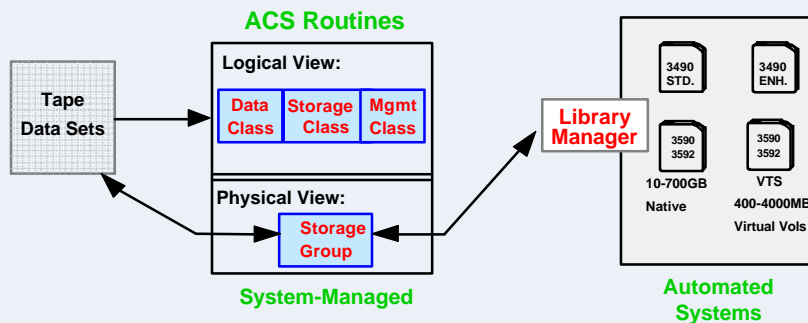
- ■ **Simplifies Installation and Management**
- ■ **NO ADDED COST!!**
  - → **MVS support through DFSMS/MVS**

Peer-to-Peer (PtP) VTS: Overview

Physical View

Logical View

Virtual Tape Controller Frames
(4-8 controllers)

zSeries Host

VTS

Peer-to Peer VTS

zSeries Host

IBM Tape Libraries
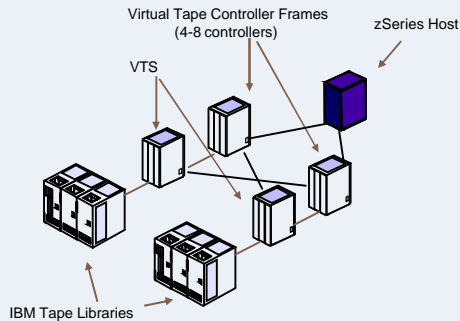
IBM offering outships competition by 10:1 margin* (recent Gartner report)
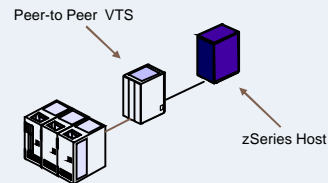
✓ **PtP Value**
✓ High Performance
✓ Increased Security
✓ High Availability
  → Maintenance
  → Upgrades
  → Disaster Recovery
  → Outages

© IBM Corporation 2007

2007 System z Technical Conference

TS7700 Grid Configuration - 1st Release

- **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 1 Gbps Ethernet links**
  - **RJ45 Copper (Cat 6)**
  - **Standard TCP/IP**
- **Policy-based replication management**
- **Can be configured for disaster recovery or higher availability environments**

TS7700

Hosts

Optional connections

I/P Links

TS1120/TS3500

Optional Hosts

TS1120/TS3500

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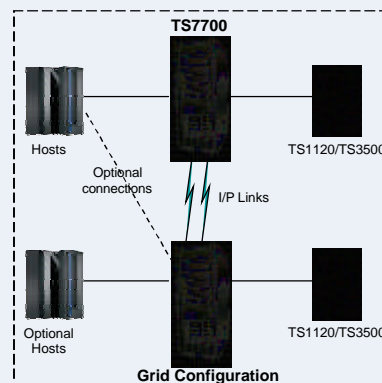
I/P replication may greatly simplify the infrastructure and management needed
for a disaster recovery solution as compared to IBM's existing PTP VTS

© IBM Corporation 2007

2007 System z Technical Conference

TS7700 Virtualization Engine Specifications

Specification	TS7740	Model B10		Model B20		Model B18		
Number of Virtual Devices	**256**	**64**		**128**	**256**	**64**		**128**
Usable Cache Capacity	**6 TB**	**216 – 432 GB**		**864 GB to 1.7 TB**		**72 GB to 1.7 TB**		
Compressed Cache Capacity (3:1)	**18 TB**	**648 GB to 1.2 TB**		**2.4 TB to 5.2 TB**		**216 GB to 5.2 TB**		
FICON	**4**	**2**	**4**	**4**	**8**			
ESCON Channels		**2**	**4**	**8**	**8**	**16**	**2**	**4**	**8**
TS1120/3592 Tape Drive Attachment	**4 - 16**	**4 - 12**		**4 - 12**				
3590 Tape Drive Attachment		**4 - 6**		**4 - 12**		**3 - 6**		
Number of Virtual Volumes	**500,000**	**250,000**		**500,000**		**250,000**		
Supports upgrade path	**planned**			**planned**				

Statements of IBM future plans and directions are provided
for information purposes only. Plans and direction are subject to change without notice.

TS7520 Virtualization Engine

New Advanced Function
- Hosted Backup
- Network Data Mgmt. Protocol (NDMP)
- Enhanced Caching
- HW assisted Compression
- Local Replication
- Encryption
- iSCSI
- Control Path & Data Path Failover
- eSA/RETAIN
- Secured Tape

Specifications
- Up to **884TB** cache w/500GB SATA drives
- **Up to 1.3PB cache** w/750GB SATA drives (Preview)
- Up to **4.8 GB/sec** Performance
- 24 Fibre channel Support
- Up to **512 Virtual Libraries**
- Up to **4096 Virtual Drives**
- Up to **256,000 Virtual Cartridges**
- Active Failover

Releases
- Limited Edition (LE)
- Enterprise Edition (EE)

Platform Support
- **System i**
- **System p**
- **System z (Linux)**
- **System x**
- **Sun***
- **HP***

***Selected environments**

Designed to scale (up to 12 frames)!!

TS7520 Virtualization Engine Components (continued)

- **Optional TS7520 software features**
  - **Network Replication**
    - Supports policy based virtual volume replication
  - **Network Encryption**
    - Helps protect sensitive customer information
    - Encrypt data in flight (data transmission)
  - **Optional Network Compression**
    - Designed to improve replication performance
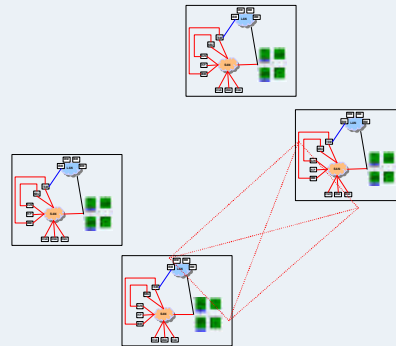
IP Link

Reasons Open System Customers Virtualized Tape

Customer Pain Point	TS7520 Feature(s)	Benefit
Backup/Restore Window	- Local Replication - Tape Caching - Compression	- Utilizes incremental backup technology to reduce backup and recovery points in time - Minimizes net time needed for backup by serving as high performance cache - Reduces backup window by increasing data transfer speeds
Data Growth	- Scalable Architecture - 128,000 virtual volumes	- Large scalable cache integration allows configuration flexibility for growth - Provides extensive access to virtual capacity
Reliability	- Call Home - ESA/Retain	- Allows the system to report problems before they cause downtime - Provides full function capability for RAS
Regulatory and Legal Requirements	- Encryption Support - IP Encryption	- Data at rest encryption to meet strict regulatory and security requirement - Provides security of data over the network
Offsite Copy	- Remote Copy	- Allows for the copied tape to reside on a remote TS7520 or vault
Recovery Point Objectives	- Local Replication	- Utilizes incremental backup technology to reduce backup and recovery points in time
Business Continuity	- IP Replication - HA Options (CPF& DPF)	- Allows for remote vaulting and the protection of business information assets - Provide various degrees of availability depending on your requirements
Hodgepodge of Solutions	- Advanced Function	- Provides various levels of configuration support and availability

25

IBM TotalStorage Productivity Center

An open storage infrastructure management solution designed to:

- *Centralize, single point of management* and control of storage infrastructure (disk, data, fabric) providing asset, capacity, performance and availability management

- *Reduce the effort of managing* complex multi-vendor heterogeneous environments

- *Improve administrator efficiency* & storage utilization

- *Provide analytic reporting* on performance impacts and configuration changes

System P   BladeCenter   System X

Brocade Director

Cisco Director

EMC   HITACHI DATA SYSTEMS   IBM   IBM Tape Library
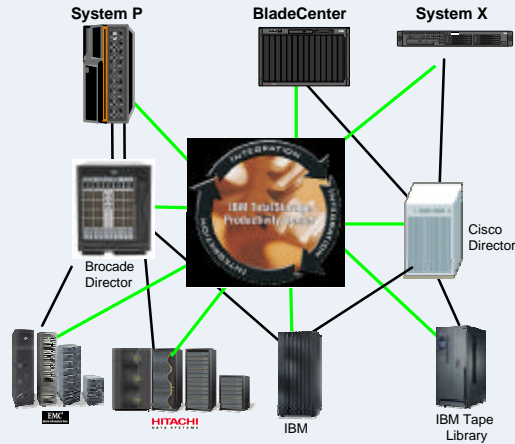
IBM TotalStorage Productivity Center

TotalStorage Productivity Center for Data
- Data collection and analysis, file systems and databases
  - What is the current utilization ?
  - Why is storage growing ?
  - How much worthless data ?
- Reporting, chargeback and quotas
- Automated actions
- Support for Heterogeneous disk (IBM, EMC, HDS, HP, Engenio)
- IBM 3584 Tape Asset Reporting

TotalStorage Productivity Center for Fabric
- SAN Topology Display and Management
- Event reporting
- Performance reporting
- Zone control
- Heterogeneous Fabric Support (Brocade, Cisco, etc.)

TotalStorage Productivity Center for Disk
- Disk Subsystem Management
- Performance Management – IBM and Heterogeneous Storage
- Volume performance advisor
- Storage Provisioning – IBM and Heterogeneous Storage
- Advanced performance management – IBM DS4000/6000/8000/SVC
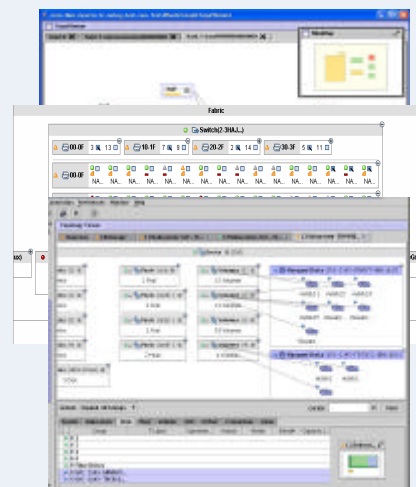
Comparing TPC Standard Edition To TPC Limited Edition

Function	Limited Edition	Standard Edition
Storage Infrastructure Configuration/Status Reporting		
Device Discovery/Configuration	✳	✳
Topology Viewer and Storage Health Management	✳	✳
Launch of Device Element Managers	✳	✳
Storage Reporting		
Basic Asset & Capacity Reporting	✳	✳
Capacity Analysis/Predictive Growth		✳
Customized and Detailed Capacity Reporting – including Chargeback and Database Reporting		✳
Performance Management		
Performance Reporting on IBM and SMI-S storage arrays		✳
Threshold Reporting on IBM and SMI-S storage arrays		✳
Fabric performance reporting and monitor		✳
SAN Volume Controller Performance Management		✳
Full Provisioning (including Fabric zoning and Disk LUN assignment		✳
Shipped with IBM hardware	✳	

SMI-S and IBM TotalStorage Productivity Center

- One SMI-S Management Application can now rule multiple devices, including those from different vendors

- SMI-S has a scheme for finding what devices are out there – service location protocol (SLP)

- IBM brings together storage management and infrastructure management

- Ensure IBM hardware provides Standards-Based Open Interfaces to administer and configure IBM hardware consistently

- Prepare for and provide Infrastructure Orchestration, Provisioning and Storage Process Management software that can be used to automate infrastructure management workflows.
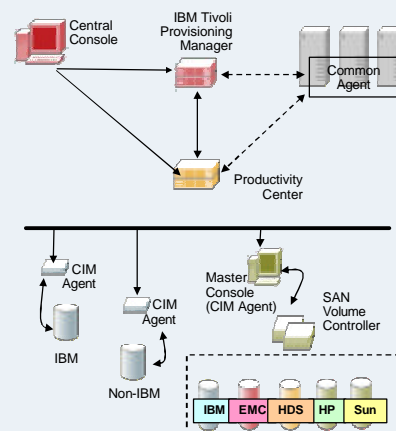
SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION

IBM TotalStorage® Productivity Center

Leverage Information
- Centralized global view of storage
- Enterprise wide reporting on databases + file level data
- Over 400 reports

Optimize IT
- End to End Storage management of Disk and Fabric
- Single interface
- Integrates with TSM + ITSM

Mitigate Risk
- Threshold monitoring
- Monitor backups + archives
- Automated status + problem alerts
- Topology viewing with health overlays
- Automate + monitor replication services

Enable Business Flexibility
- Open Standards Based
- Heterogeneous Device Support
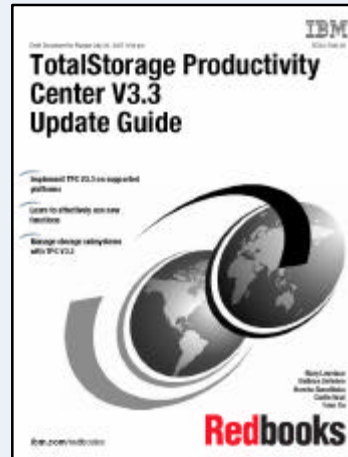
The IBM Tivoli Storage Manager

Complete Storage Management

- Backup/Restore
- Disaster Recovery
- Bare Machine Recovery
- Database and Application Protection

- Archive / Retrieve
- Space Management
- Data Retention

- TSM database catalogs stored data objects and location, facilitating intelligent, policy-based storage management
- Automated, object-level policy allows high granularity in managing data
- Progressive incremental backup reduces bandwidth/storage requirement and simplifies restore
- Choices offer flexible backup/recovery: sub-file, image, journal-based, backup sets
- Integrated storage hierarchy provides efficient device utilization
- Policy-based collocation, migration, and reclamation for efficient storage utilization and fast restore

Servers, Clients, Application systems   WEB   LAN   SAN   TSM Server   Log   Database   Storage Repository
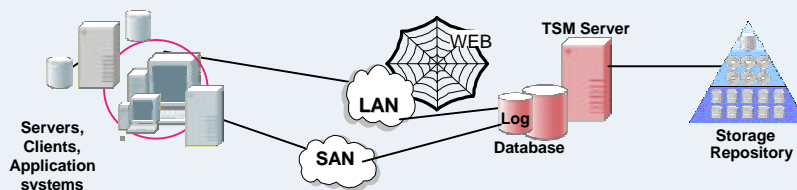
TSM Server on zSeries Linux

- **TSM on zLinux is more like TSM on Open Systems**
  - **TSM on zLinux leverages the zSeries hardware**
  - **Must use SCSI-attached tape drives**
    - Tape drives on traditional zOS systems are ESCON/FICON, not SCSI
  - **Tape libraries which must perform work on both zOS and zLinux must be partitioned**
    - Adds complexity over dedicated tape libraries
  - **Tape drives within library must be dedicated to either zOS (ESCON) or zLinux (SCSI)**
    - Tape drive sharing between zLinux and zOS not possible
  - **Migration of TSM from zOS to zLinux is no easier than migration to other platforms**

LPAR or IFL

z/VM

TCP/IP	TCP/IP
Linux on zSeries	Linux on zSeries
TSM Server	TSM Client

TSM Database

TSM Storage Pools

FCP/SCSI Support Required for Tape Library
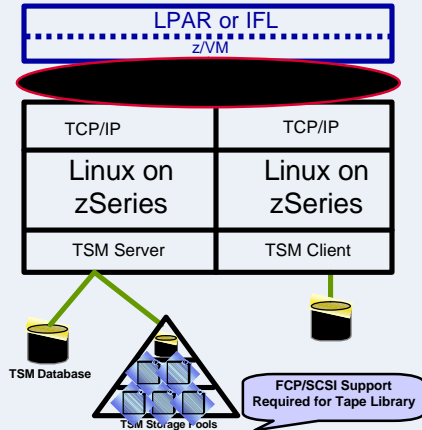
A Unified IBM System Storage and System z Approach

Take back control of your infrastructure

- **Lower costs with tiered storage offerings for effective information lifecycle management**
  - Leverage IBM's enterprise continuum of disk storage, SAN and tape technology to best fit the accessibility needs and value of your data
- **Support data security, business continuity, regulatory compliance**
  - IBM Systems support data security and end-to-end world class business continuity solutions encompassing SAN, disk, tape, server, software and services
- **Reduce costs and complexity through infrastructure simplification and consolidation**
  - High performance, highly scalable systems that include features like 4 Gb FICON/FCP connectivity can enable you to unify your infrastructure to help lower costs, simplify management

Resilient, cost effective, world-class solutions. IBM's integrated infrastructure for On Demand Business.

Thank You

Trademarks

- **The following terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.**

• AIX	• eServer	• ON (button device)	• ServerProven
• AIX 5L	• FICON	• On demand business	• System z9
• BladeCenter	• FlashCopy	• OnForever	• System p5
• Chipkill	• GDPS	• OpenPower	• System Storage
• DB2	• Geographically Dispersed	• OS/390	• Tivoli
• DB2 Universal Database	Parallel Sysplex	• OS/400	• TotalStorage
• DFSMSdss	• HiperSockets	• Parallel Sysplex	• TotalStorage Proven
• DFSMShsm	• i5/OS	• POWER	• TPF
• DFSMSrmm	• IBM	• POWER5	• Virtualization Engine
• Domino	• IBM eServer	• Predictive Failure Analysis	• X-Architecture
• e-business logo	• IBM logo	• pSeries	• xSeries
• Enterprise Storage Server	• iSeries	• S/390	• z/OS
• ESCON	• Lotus	• Seascape	• z/VM
			• zSeries

- **Linear Tape-Open, LTO, LTO Logo, Ultrium logo, Ultrium 2 Logo and Ultrium 3 logo are trademarks in the United States and other countries of Certance, Hewlett-Packard, and IBM.**
- **Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.**
- **Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.**
- **Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States and/or other countries.**
- **UNIX is a registered trademark of The Open Group in the United States and other countries.**
- **Linux is a trademark of Linus Torvalds in the United States and other countries.**
- **Other company, product, or service names may be trademarks or service marks of others.**