




IBM System Storage for System z

*Announcement Update
DS8700 & Virtual Tape*

Volker Widder
Storage Principal FSS
widder@de.ibm.com



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


DS8700

- IBM DS8700 HW
- System z Synergy Items
- Solid State Disk
- Preview Rel. 5.1
(Automated Data Relocation)

Tape News

- System z Tape Offerings
- TS7700 Release 1.5 & 1.6
- TS7680 ProtecTIER Gateway
(Data Dedup for Mainframes)

IBM Storage Reloaded

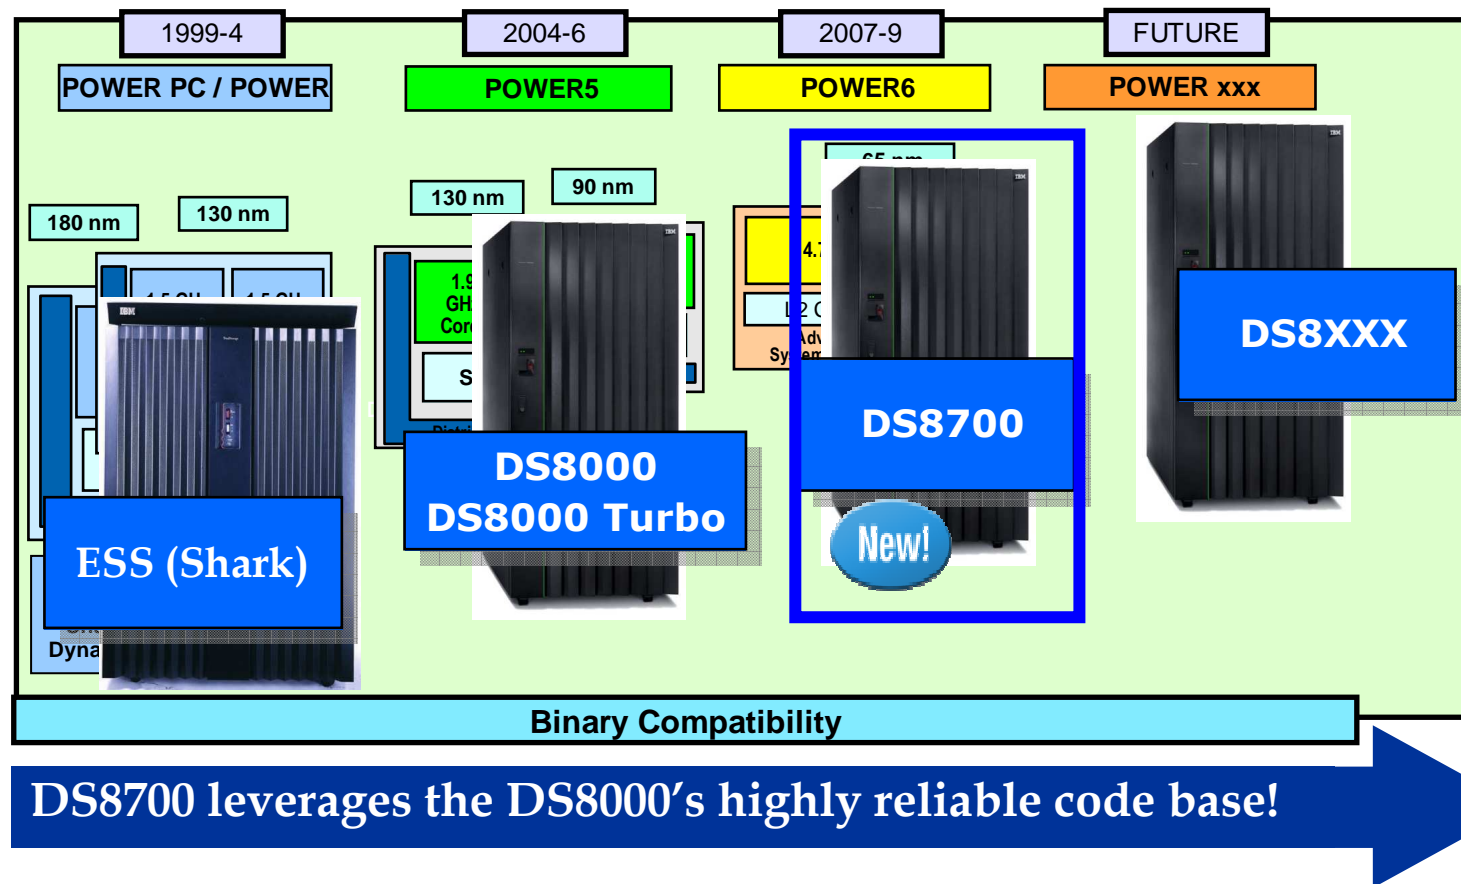


IBM POWER and DS8700 Reliability and Resiliency



The IBM POWER processor has been behind the success of IBM enterprise storage beginning with the Enterprise Storage Server (Shark) in 1999

Benefit: a steady, 10-year lineage of RAS improvements!



Introducing the New DS8700

The Next Chapter in IBM's Flagship Disk Platform

What's New

- **Performance**
 - Up to **over 150% performance boost** with **POWER6**-based controllers
 - New, **faster PCI Express (PCI-E) internal fabric**
 - Almost **70% faster ASICs** on the device adapters
 - **Increased FlashCopy performance** across the board
 - **Increased SSD performance** for sequential reads
- **Availability**
 - **Single model, scalable via concurrent upgrade of all components**
 - Faster concurrent microcode updates (2h-3h vs. 8h)
 - **Better than 99.999% availability**
- **Investment Protection and Scalability**
 - Single model, scalable via concurrent upgrade of all components
 - Full interoperability with prior model's hard drives, drive enclosures, tools, scripts, and copy services
 - **SSD optimization for more effective tiered storage**
- **Simplified Management**
- **Security**
 - **Full Disk Encryption** enhancements help address PCI-DSS compliance



DS8k R5



Changing from DS8000 to Follow on



Change in I/O structure



- New IO bays are directly connected to the servers via point-to-point PCI-e cables.
- IO Towers no longer share common 'loops', they connect directly to each server via separate cables and link cards.
- The server to server communication path continues to use RIO, but is now isolated from IO traffic.

- 4.7 Ghz P6 570 CECs
- Cache 32GB to 384GB
- 4Gb/s FC / FICON

Proven Server-Based Architecture

New IBM POWER6 processor complex offers...

– Exceptional performance

- Up to over 150% performance improvement for sequential workloads
- Up to 100% performance improvement for random workloads

– Improved Reliability, Availability, Serviceability

– Higher Scalability

- Enables higher scalability through POWER6 design, increased cache size (384 GB), and higher RAS

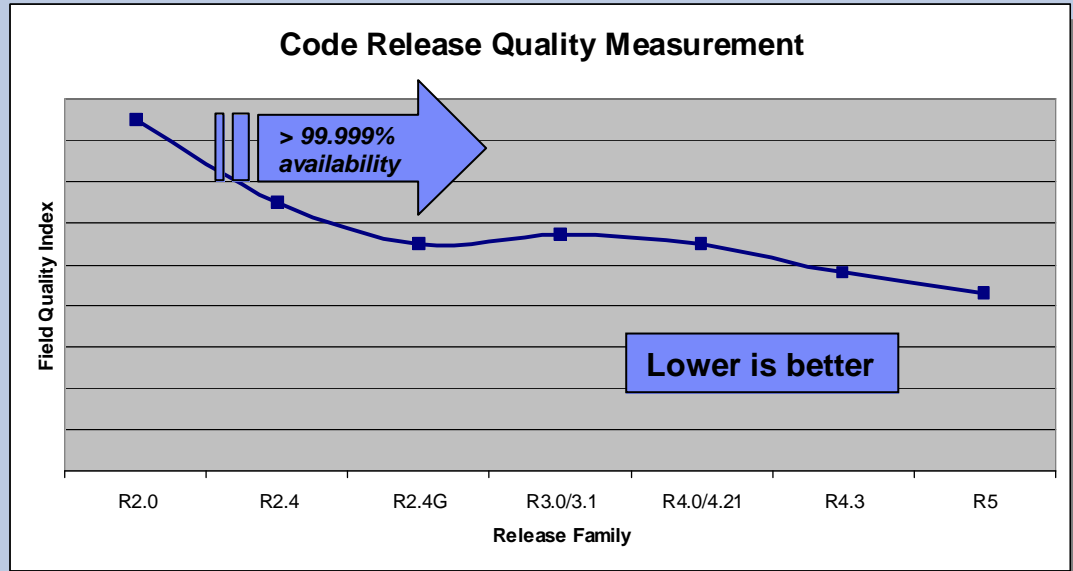
– Greater Energy efficiency

- Over 50% more IOPS/Watt

	Units		DS8700 vs DS8300
Rd Seq	GB/s		2.20x
Wr Seq	GB/s		2.50x
DB z Ficon Cache Std.	K IO/s		1.45X
DB open (50/50/50)	K IO/s		1.54x

**All these benefits with minimal changes to existing firmware and codebase ...
... which translates to higher system reliability and information availability**

22% improvement in the projected quality over prior release, which already boasted extremely high reliability



Continuous quality improvements

- Market-proven IBM POWER architecture
- Design simplification and reliability enhancements from POWER6 processor and PCI-express interface
- Higher* than five-nines availability (99.999+%)

Improved and Simplified Management

GUI enhancements

The screenshot displays the IBM System Storage DS8000 Storage Manager interface. The main content area is divided into several sections:


- System Summary:** Shows system configuration for Storage Unit <75H1350>. The state is "Powered on". Capacity summary shows 99% usage. Performance metrics include Raw Capacity (147,392 GiB), Allocated Capacity (146,690 GiB), and Available Capacity (702 GiB).
- Management Consoles:** A section for monitoring system health and performance.
- Performance:** Four line graphs showing Host MBps, Host KIOps, Rank MBps, and Rank KIOps over time. The next update is scheduled for 23. The legend indicates SI 1 Read (green) and SI 1 Write (blue).
- Racks:** A visual representation of the storage unit's physical components, showing three racks.

New System Summary Panel
 – Overview tab displays system configuration and real-time performance data for a 5 minute sampling interval

Updated every 60sec.

Number of racks shown matches racks in the storage unit

Why Host + Disk System Synergy Matters



Synergy means cooperation. *IT value is maximized when parts work well together.*

Host↔Storage synergy can *significantly enhance IT operations.*

More efficient, more cost-effective IT operations enhance application delivery, *benefiting the overall business.*

System z and Storage Synergy



- System z[®] is about maximizing availability, performance, consolidation, security, scalability, and more, to support mission-critical applications.



- System z, especially z/OS[®], likely provides more specialized storage functions to support these objectives than any other OS on the planet.
- These specialized functions work *only if* the disk system is designed to support them.

System z



- Joint IBM storage + server effort: design, development, test, customer support
 - Competitors don't have access to this collaboration and experience. Even if a given function is eventually delivered by a competitor, that is not an indication of comparable quality, performance, or a design with future IBM-planned enhancements in mind.
- IBM is best positioned for earliest delivery
 - Examples: support for new server models, new OS versions/releases, new synergy functions
- IBM licenses selected functional specifications to competitors. Specifications ≠ Code.

System z Synergy = Value

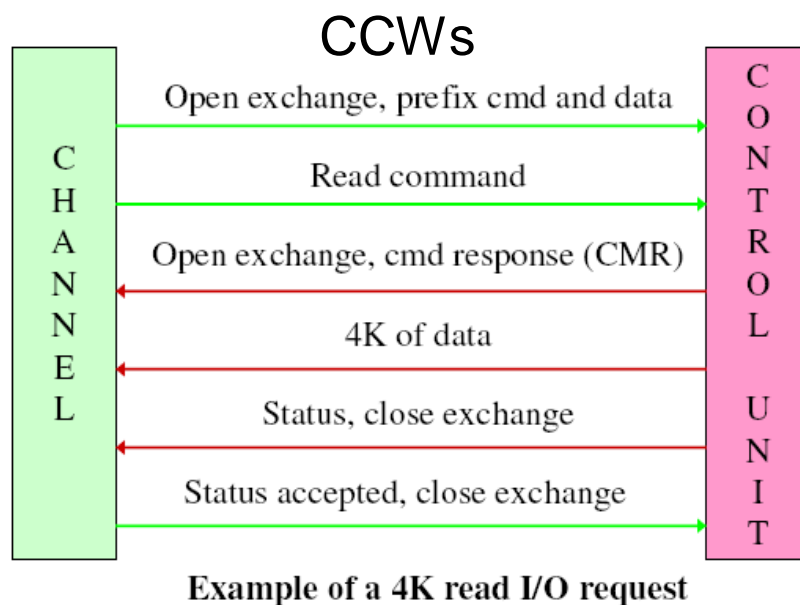


	Availability	Performance	Ease of Mgmt	Growth	Security	TCO
High Performance FICON®		✓	✓	✓		✓
High Performance FICON - multitrack		✓	✓	✓		✓
Larger 3390 volumes (EAV)			✓	✓		
Dynamic volume expansion – 3390s	✓		✓	✓		
Space-efficient volume copies				✓		✓
DFSMS™ recognition of SSDs			✓			
Basic HyperSwap™	✓					✓
Remote Pair FlashCopy®	✓	✓	✓			✓
MIDAW		✓	✓	✓		✓
z/OS Metro/Global Mirror (MzGM) incremental resync after HyperSwap	✓	✓				✓
Self-encrypting disks			✓		✓	✓

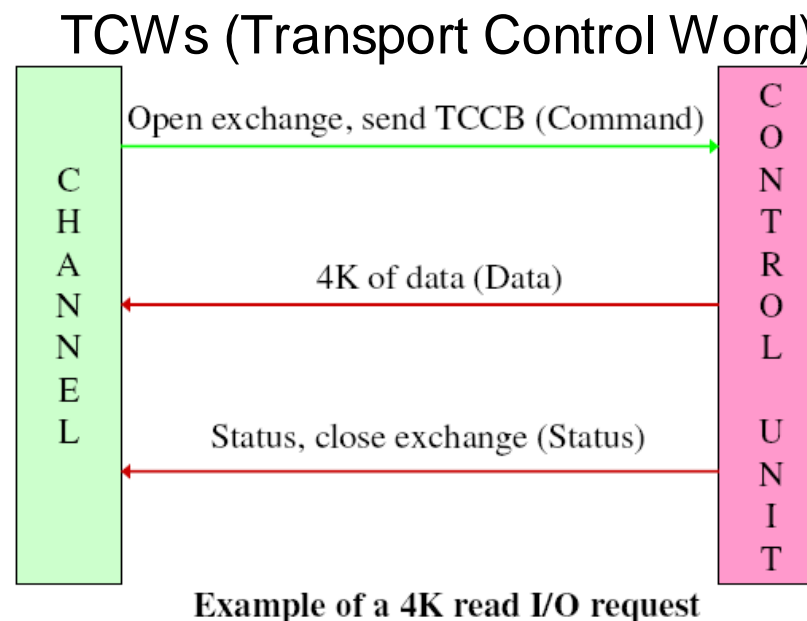
Additional Information: IBM System z & DS8000 Technology Synergy, Robert F.Kern, July 21,2009

FICON vs z High Performance FICON (zHPF)

FC Link Protocol for FICON



FC Link Protocol for FCP & zHPF



Improves FICON efficiency, scale, value, and RAS

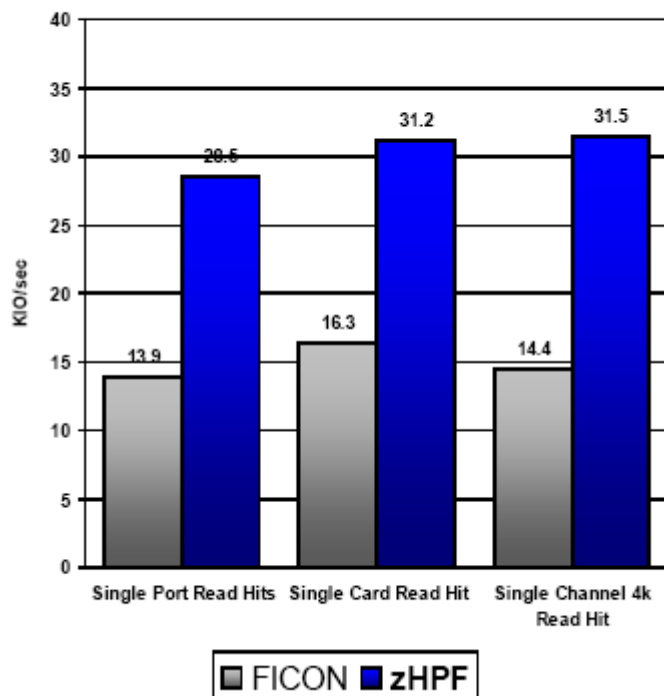
- Significant performance enhancement for small block IOs
- TCO savings (do more with less)
 - reduce the number of channels, switch ports, CU ports, optical cables

High Performance FICON for System z (zHPF)

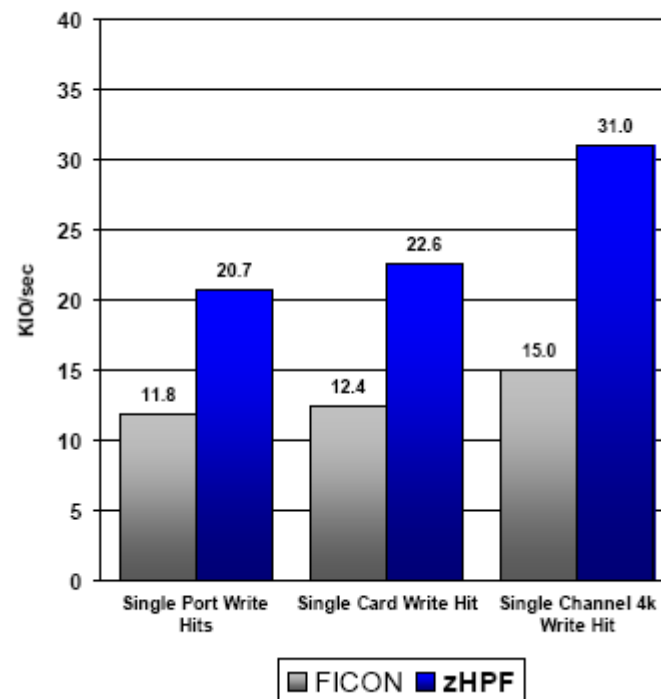
VALUE

- Up to 2x higher I/O throughput per channel
 - Realistic workloads with a mix of data set transfer sizes can see 30–70% of FICON I/Os utilizing zHPF, resulting in **up to a 10-30% channel utilization savings.**
- Can reduce manual tuning
- TCO: May require fewer channels on the z server and on the disk system

4k Read Hit Performance



4k Write Hit Performance



High Performance FICON zHPF – Multi-Track Support

▪ In R4.3

- Higher performance for System z workloads by allowing more input/output operations that can be converted to zHPF protocol
 - Allows to **read/write more than one track of data per channel command** and utilize zHPF
 - Improves channel efficiency
- Higher system resilience to workload spikes
- Helps reduce overhead and improve end-to-end system reliability, availability, and serviceability
- Multi-track is exploitable by **DB2 and VSAM** when using **Extended Format Datasets** for software striping. It is also exploited by the **zFS file system** for USS, Java and WebSphere applications.

▪ zHPF Host Support


- High Performance FICON licensed feature # 0709
- DS8000 series Function Authorization feature # 7092
- z10 Processors
- zOS V1.9+ higher
- APARs OA26084 and OA29071 are required to enable multi-track support

Dynamic Volume Expansion for System z

Innovation that *Matters!*

Grow with Ease Using Online Expansion

- **Simplifies management by enabling easier, online, volume expansion to support application data growth.**
- Yields more highly available, simplified volume expansion
 - No longer need to backup data, bring volume offline, delete and recreate volume to expand it
 - Can be done on-line with a single command or GUI screen
 - Copy services relationships must be removed before volume expansion
 - Can be used to non-disruptively migrate from smaller to larger z/OS volumes




**Supported with
z/OS 1.7 with PTFs**

Extended Address Volumes


Value proposition:

EAV simplifies deployment of large System z installs

- Eliminates volume capacity constraints for large System z installs
- Manage fewer, large volumes as opposed to many small volumes
- First supported with VSAM or applications that use VSAM data sets (e.g., DB2 and CICS)
 - intent to support larger volume sizes and additional data set types and access methods in future*
- Included in the R4 microcode at no additional charge
- Requires z/OS V1.10
- **DS8000 HyperPAV function complements EAV by allowing the ability to scale the I/O rates against a single, larger volume**
- DS8000 automatically converts 3390-9s to the new EAV models when growing a 3390-9 by using DS8000 dynamic volume expansion



DFSMS Recognizes Solid State Disks (SSD)



As part of the DFSMS support, a new device performance capability table exists for volumes backed by Solid State Drives. A user who wants to favour SSD selection over non-SSD can assign a storage class with a direct MSR value of 1 and a direct bias of R.

Value

- **Helps users allocate new data sets to the desired type of disk**
- **Supports tools and commands that report on types of disk installed**

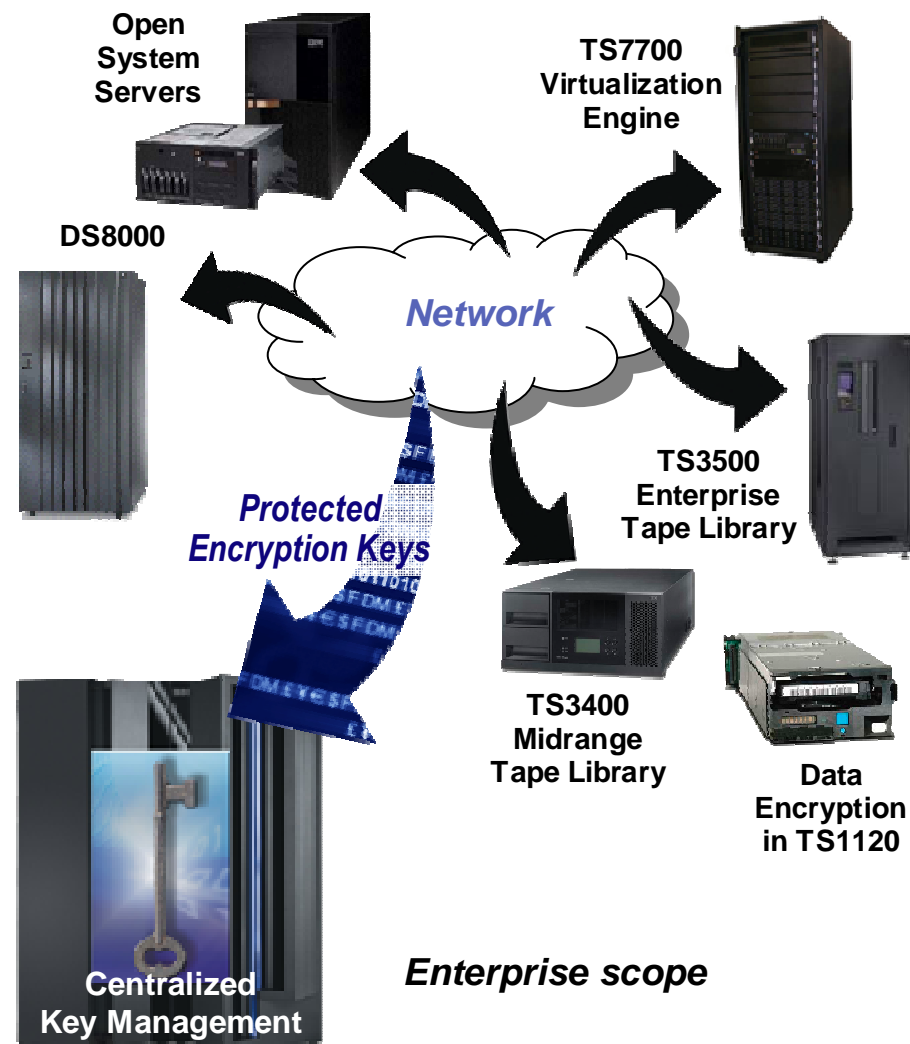
IBM DS8000 Disk Encryption


- Customer data at rest is encrypted
 - Data at rest = data on any disk or in any persistent memory
- Customer data in flight is not encrypted
 - Data in flight = on I/O interfaces or in dynamic memories (Cache, NVS)
- Uses Encrypting Disk
 - Encryption hardware in disk (AES 128)
 - Runs at full data rate, no performance impact
 - 146/300/450 GBs 15K RPM
- Integrated with Tivoli Key Lifecycle Manager (TKLM)
 - DS8000 network attachment to TKLM
 - DS8000 automatically communicates with TKLM when configuring encryption group or at power on to obtain necessary encryption keys to access customer data
- Supports cryptographic erasure data
- Key attack vectors prevented:
 - Disk removed (repair, or stolen)
 - Box removed (retire, or stolen)

The System z and System Storage encryption solution

Delivers integrated security

- z/OS encryption controlled via Data Policy (SMS) and user Policy (JCL)
- Storage Encryption of Tape and DS8000 managed by IBM Tivoli Key Lifecycle Manager (TKLM) for z/OS V1.0
- Protects data on disks removed due to failure, or in a redeployed or retired system
- No performance impact on disk
- Unique key per drive





DS8000

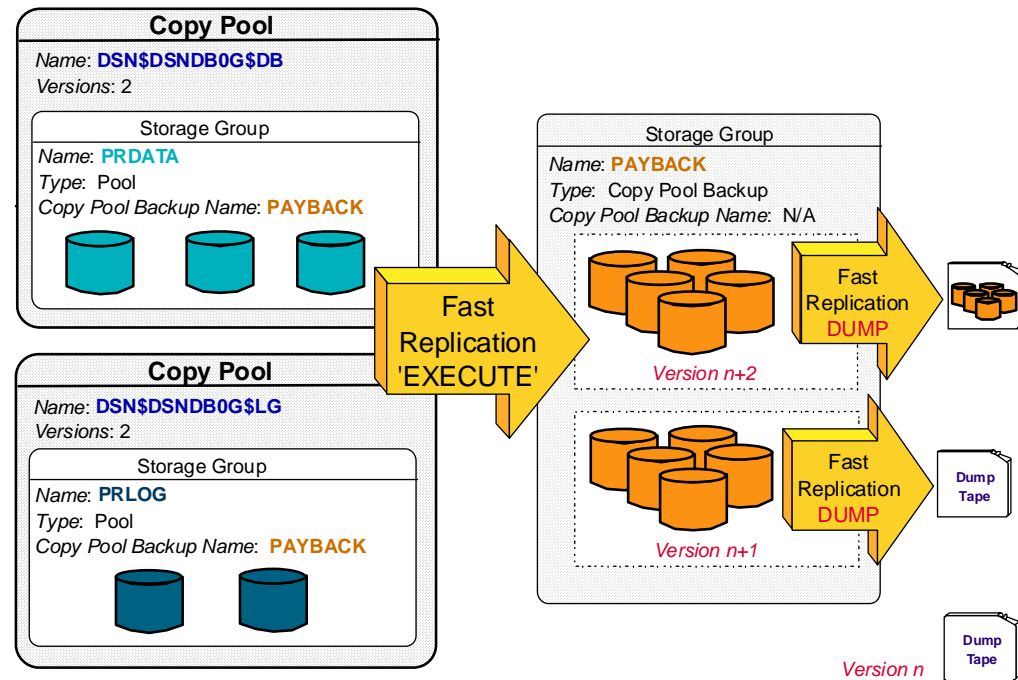
- FlashCopy® – Initialization
 - A target is allocated the same size as the source
 - Application I/O is paused only while a new version of the block map is created
- FlashCopy® – Full Copy
 - All data blocks are physically copied
- FlashCopy® – Incremental
 - Only the new data blocks are physically copied since the Full Copy
- FlashCopy® – Copy on Write (NOCOPY)
 - Data blocks are physically copied only as they are updated
- FlashCopy® Space Efficient (aka SnapShot)
 - Only the changed data blocks need extra space

- **DFSMSdss**
 - COPY
 - DUMP
 - DEFrag
 - CONSOLIDATE

- **DFSMShsm**
 - FRBACKUP
 - FRRESTORE

- **DFSMSrmm**
 - CDS Fast Replication

- **DB2**
 - BACKUP SYSTEM
 - RESTORE SYSTEM
 - DB2 Utilities



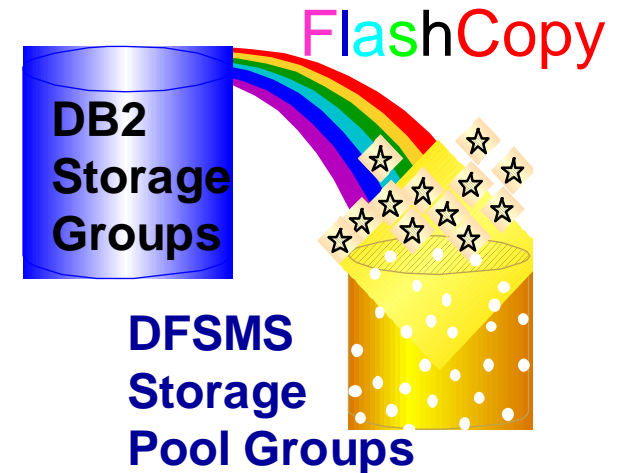
DFSMSHsm fast replication

- In previous releases of z/OS, you could use DFSMSHsm fast replication to create backups for a set of storage groups in a copy pool on target volumes.
 - You could not, however, use DFSMSHsm to dump the fast replication backups to tape.
 - Now you can use DFSMSHsm to dump the fast replication backup copies to tape.
 - You can recover fast replication backups at the volume level from dumps or DASD copies, or at the copy pool level from DASD copies.
- A further enhancement in z/OS V1R8 allows you to recover individual data sets from fast replication backups
- To recover individual data sets from copy pool backups, you use the FRRECOV command, which is expanded with new parameters that allow you to specify one or more fully or partially qualified data set names.
- You can recover data sets from FlashCopy target volumes or dump tapes (now supported for recovery)

Virtual Concurrent Copy

- Concurrent Copy (CC)
 - Uses Control Unit Cache
 - Larger volumes cause constraint problems
- Virtual Concurrent Copy (SnapShot)
 - Data set is “snapped” to Working Space Data Set (WSDS)
 - After snap, backup is *logically* complete
 - After snap copy is moved to target device, backup is *physically* complete
- Virtual Concurrent Copy (FlashCopy) *New*
 - Supported on volumes that are CC and FlashCopy V2 capable
 - Data is Flashed using NOCOPY to WSDS
 - After Flash, backup is *logically* complete
 - After standard I/O is used to move the data to target device, backup is *physically* complete
 - Supported for Data Set Backup and Full-Volume Dump

- New utilities to enhance DB2 backup and recovery
 - Backup System utility
 - Restore System utility
- DB2 backups are managed by DFSMSHsm
 - DB2 Storage Groups are mapped to DFSMS Copy Pool Groups
 - Up to 50 backup versions on disk (BSDS limit)
 - DFSMS uses FlashCopy to non-disruptively make copies of DB2 Storage Groups
 - DB2 logs are adjusted so data integrity is maintained **without DB2 having to take quiesce points**
- Recovery of DB2 Storage Groups is very quick
- Prerequisites:
 - z/OS V1.5 DFSMSHsm
 - DB2 UDB V8 for zOS
 - DS8000 FlashCopy



DB2 Utilities utilize Dataset FlashCopy

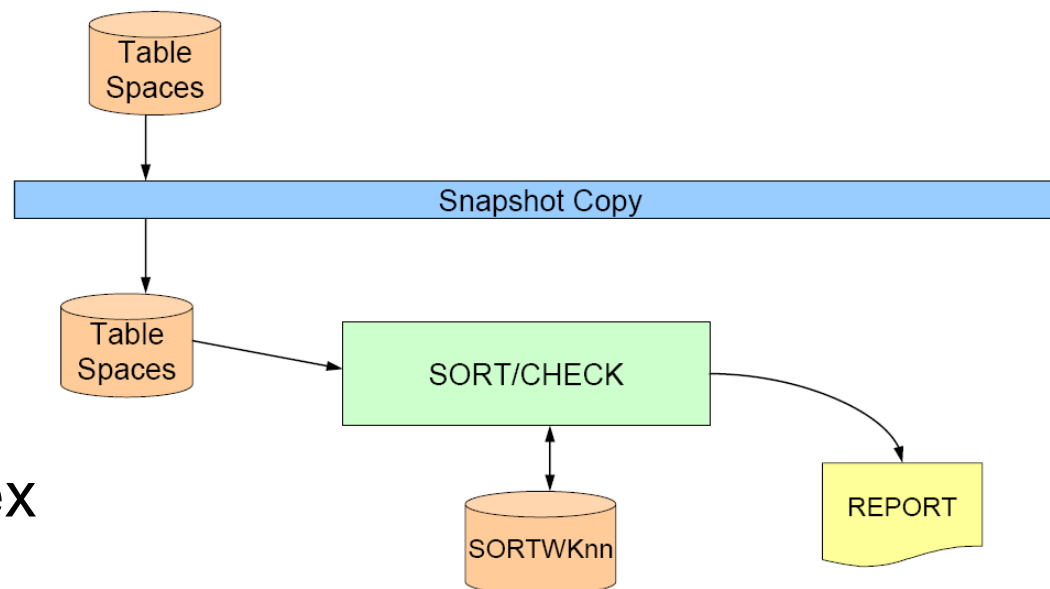
■ Utilities

–DB2 V8

- Online Check Index

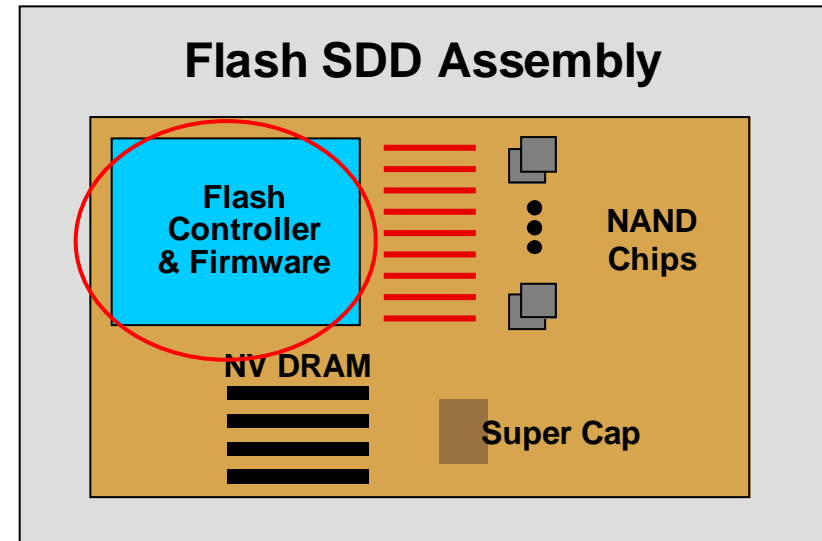
–DB2 V9

- Online Check Data
- Online Check LOB



What are Solid State Disks?

- Solid State Drives:
 - Use semiconductor based flash technology (not magnetic, not optical)
 - Are plug compatible with common HDD form factors
 - 3.5”, 2.5”, 1.8”
 - Use no mechanical moving parts
 - Use less power than HDDs
- Flash **chip technology** has been available for years in consumer devices like USB flash drives
- Flash **controller technology** is enabling flash chip technology to achieve capabilities for use in enterprise storage and server deployments
 - Improved write endurance
 - Improved write performance
 - Improved reliability



Currently available in 73GB and 146GB drive sets for DS8700


DS8000 SSD Performance

DB2 Sync I/Os Running on z/OS

DB2
for
z/OS

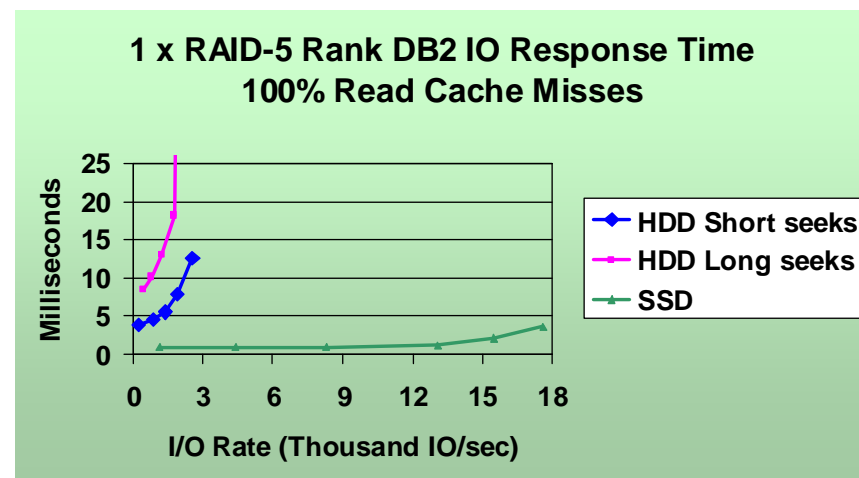



DB2 Sync I/O Wait Time (reads) (Microseconds)



DB2 SSD Performance Improvement

- Synchronous database reads
- List pre-fetch
- Random updates and inserts
- Incremental copy and incremental system backup
- Faster recovery due to faster log apply
 - Fast Log Apply uses list pre-fetch and batched writes
- Utility updates of non-partitioned indexes



SAP and IBM - Core Banking Day Posting Workload

Improved performance with zHPF and SSD's (Solid State Drives)

SSD Measured results	Baseline HDD only	HDD + SSD	% Improvement	HDD + SSD + zHPF	% Improvement
Throughput	14.3M postings/hour	18.0M postings/hour	26%	18.7M postings/hour	31%
DB request time	1.13 sec	0.682 sec	40%	0.605 sec	41%
DASD response time	5.18 ms	3.35 ms	35%	2.85 ms	45%

Feature	Benefits
SSD (Solid-State Drive)	<ul style="list-style-type: none"> • Increased data throughput • Improved Database response time • Improved data center environmental results, (e.g., reduced electrical energy needs, facility space, emissions, etc.)
zHPF (High Performance FICON for System z)	<ul style="list-style-type: none"> • Exploits a new channel protocol especially designed for more efficient I/O operations • Designed to satisfy the performance requirements of bandwidth-hungry applications

IBM “End to End” Solid State Storage Strategy

- **Workload monitoring**
 - Monitor and analyze workloads to determine which workloads can most benefit from solid state technology
- **Smart data placement**
 - Identify and place the right data in solid state storage that will get the most benefit from the technology
- **Application optimization for flash**
 - Adjust structure settings in applications to take maximum advantage of solid state storage
- **Flash optimized architecture**
 - Architect the infrastructure to maximize solid state technology benefit
- **Advanced solid state technology**
 - Develop new solid state technologies further enhance performance and lower cost


IBM plans to introduce dynamic data migration capabilities in our DS8700 disk array

This capability can be done at the full LUN or Volume level, or for smaller pieces of LUNs



Smart data analysis tools


- **New DS8000 performance analytics** for identifying **sub-volume** data to exploit tiered storage for both open and zOS
- **System Monitoring Facility (SMF)** enhancements for analyzing z/OS trends and peak usage. SMF record updates:
 - SMF 42 Subtype 6 (per data set)
 - S42DSRDD (average disconnect time for reads)
 - S42DSRDT (total number of read operations)
 - SMF 74 Subtype 5 (DS8000 cache to backend disk vol operations)
- **IBM FLASHDA** tool based on SAS software for identifying opportunities for SSD exploitation
- **Softek DMCzOS 4.1** uses SMF data to identify candidate datasets for SSD
- **Tivoli Storage Productivity Center for Disk** for real-time performance monitoring and historical performance analysis
- **Omegamon XE** for analyzing dataset and volume perf.






Smart data *placement* tools

- **Upcoming D8000 automated data relocation function*** to dynamically move hot and cold sub LUN/sub volume data across drive tiers
- **DB2 Online Reorg** function for migrating DB2 table spaces to appropriate drive tier
- **z/OS DFSMS** to migrate data to the right tier via HSM policies
- **Softek Data Mobility Feature for z/OS (zDMF)** for real time data analysis and data movement without disrupting applications
- **Softek Data Migration Services** can migrate mainframe and open data non-disruptively



Relocating only 10% of existing data to SSDs, can result in increased system throughput by 300% **



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

** Internal IBM Performance benchmark run with TPC-E-like workload


DS8700 Rel 5.1 Preview

Improve Enterprise Storage System
using
IBM Automated Data Reallocation
for Smart Storage Tiering

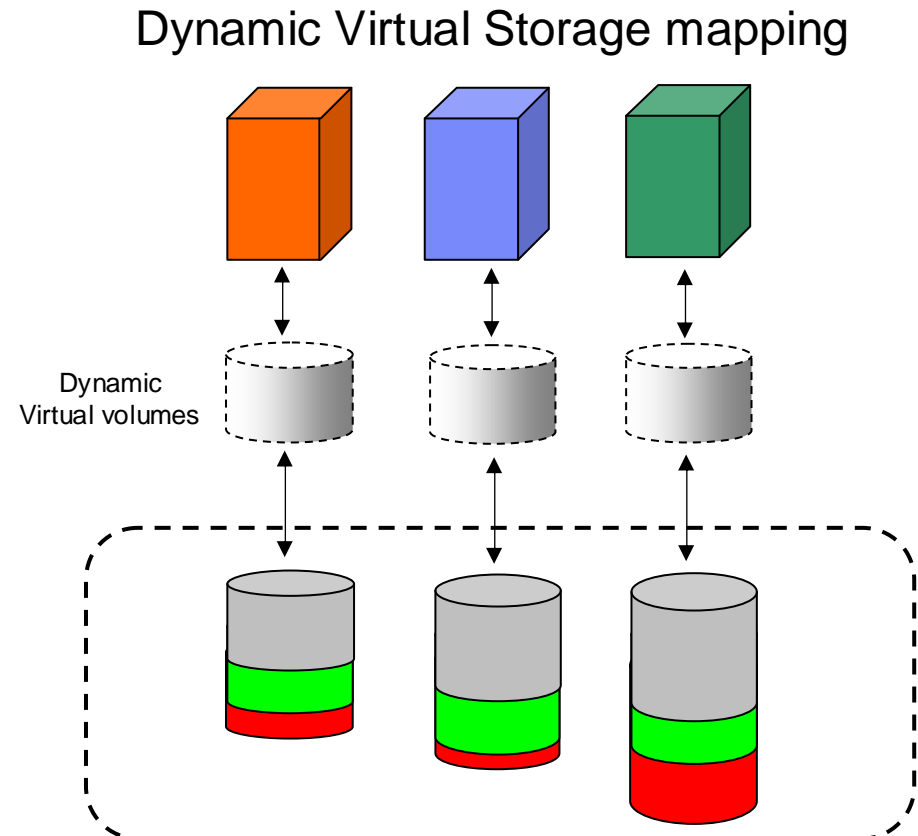


The information on new product features is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product features is for informational purposes only and may not be incorporated into any contract. The information on the new product features is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

Traditional Disk Mapping vs. *Dynamic Virtual Storage Mapping*



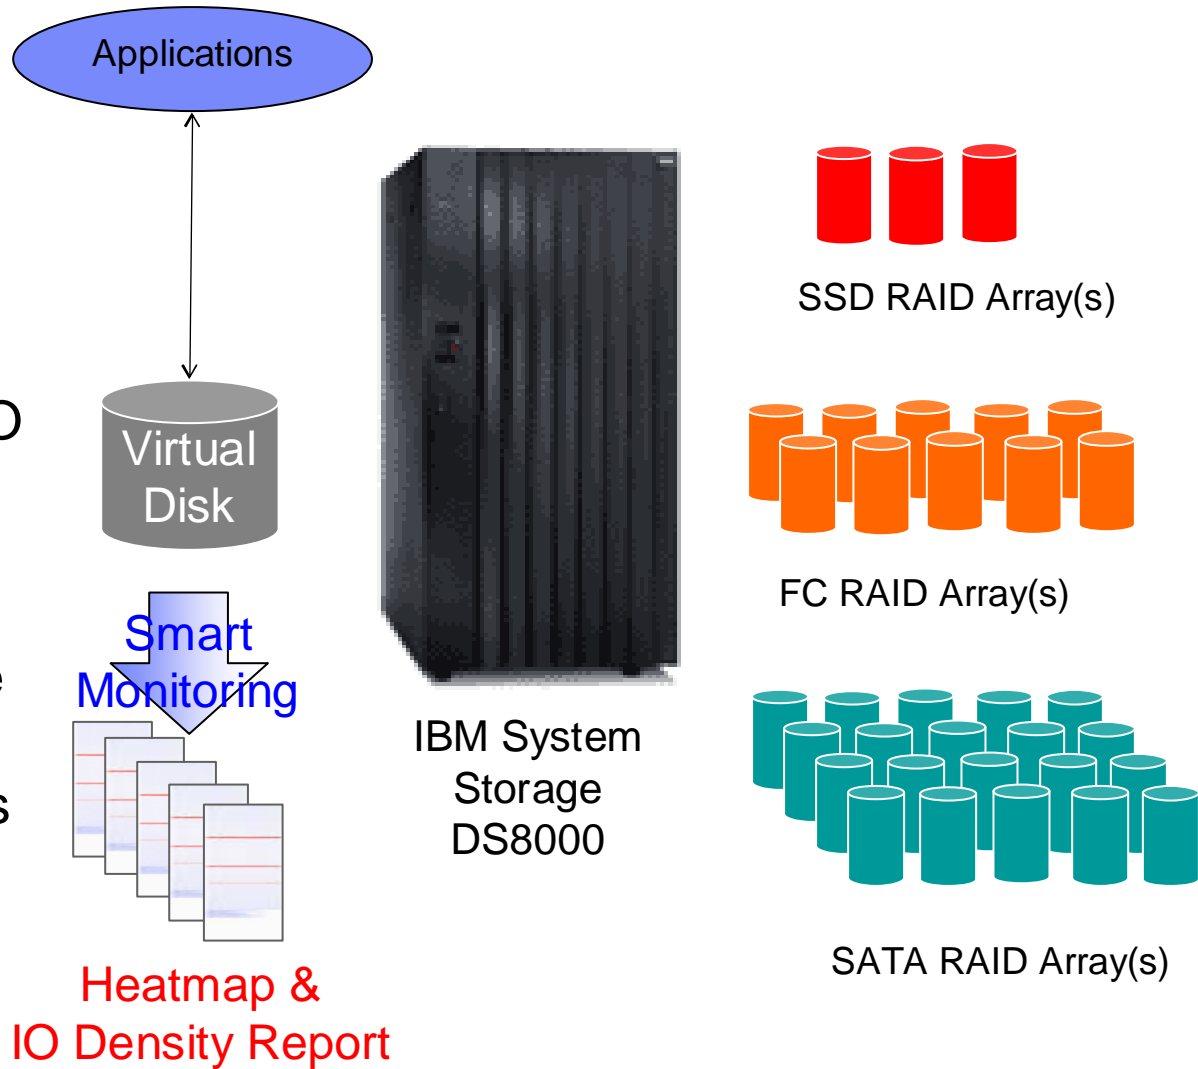
Volumes have different characteristics. Applications need to place them on correct tiers of storage based on usage.



All volumes appear to be “logically” homogenous to apps. But data is placed at the right tier of storage based on its usage through smart data placement and migration.

Fine Grain performance monitoring and reporting available to track the IO demand from application and IO service time from storage device.

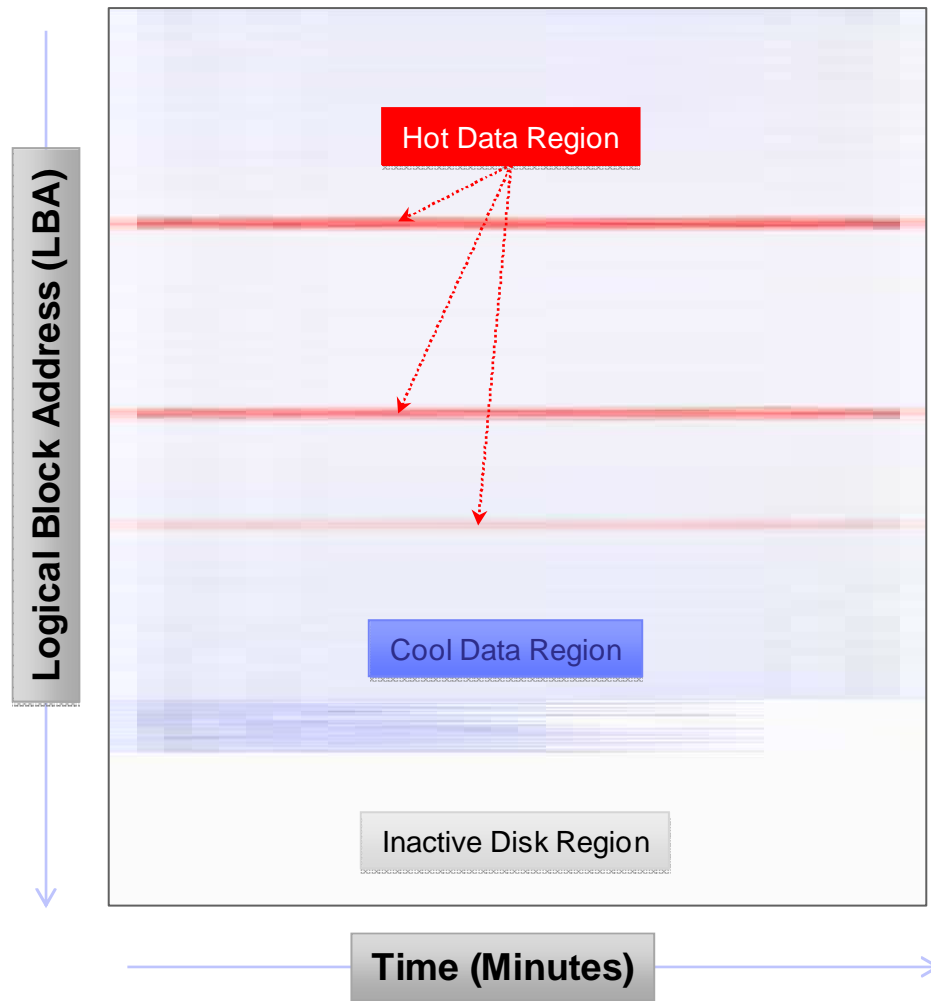
Performance data are collected for multiple durations, hours, days and weeks.



Workload Learning through Smart Monitoring




DS8000 Smart Monitoring and Analysis



- Each workload has its unique I/O access patterns and characteristics over time.
- Smart Monitoring and analysis tool allows customers to develop new insight to application optimization on storage infrastructure.
- This workload shows performance skews in three LBA ranges.
- In 1H2010 we will have the ability to automatically relocate data to the appropriate tier... at the sub-volume level

Data Placement Optimization (Sub-Volume Level)

Logical volume with multiple storage tiers



	"hot" extent
	Medium use
	Low use

Physical Storage Layer



SSD ■

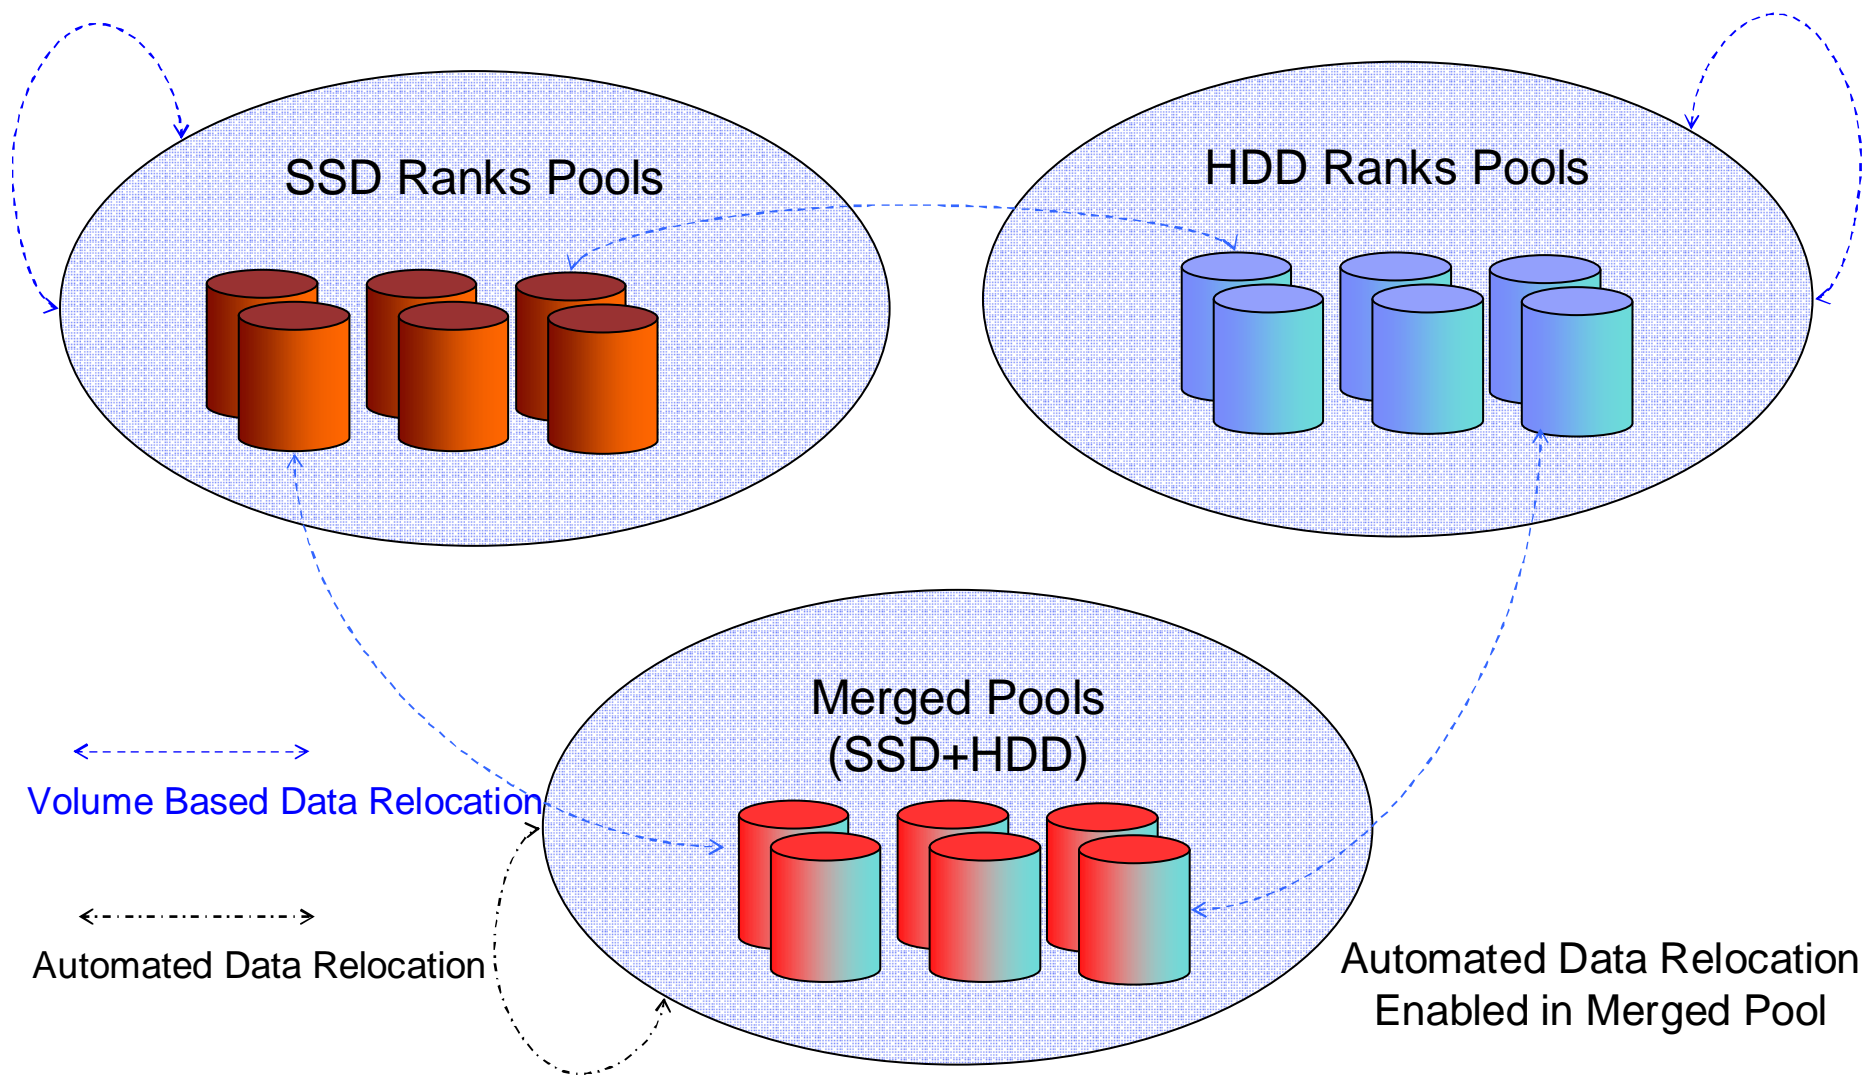


Fibre Channel/SAS ■



SATA ■


Data Relocation – Volume and Sub-volume



The Solution: IBM Easy Tier

Smart Movement and Management of Information for SSD Optimization

- 5% data moved for >50% response time improvement
- SPC-1-like IBM Internal test, Brokerage workload simulation, 2009



The information on new product features is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product features is for informational purposes only and may not be incorporated into any contract. The information on the new product features is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.


Additional Information on IBM Solid-State Drive Offerings

1. **Press Release: Solid-state drive sets speed record, Sept. 2008**
By combining Flash solid-state technology and IBM's storage virtualization technology, the researchers were able to transfer data at more than 1 million Input/Output (I/O) per second.
<http://www.ibm.com/news/us/en/2008/09/03/h532433j77630i97.html>
2. **White Paper: IBM System z® and System Storage DS8000: Accelerating the SAP® Deposits Management Workload With Solid State Drives, 2009**
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101442>
3. **Performance White Paper: Driving Business Value on Power Systems with Solid State Drives, April 2009**
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03025USEN&attachment=POW03025USEN.PDF&apname=STGE_PO_PO_USEN_WH
4. **Redpaper: Ready to Access DB2 for z/OS Data on Solid-State Drives, 2009**
<http://www.redbooks.ibm.com/redpapers/pdfs/redp4537.pdf>
5. **Redpaper: DS8000: Introducing Solid State Drives, 2009**
<http://www.redbooks.ibm.com/abstracts/REDP4522.html?Open>
6. **Techdoc: IBM System Storage DS8000 with SSDs An In-Depth Look at SSD Performance in the DS8000**
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101466>
7. **Press Release: IBM Helps Companies Gain Control of Their Information with Solid State Flash Technology, May 2009**
<http://www-304.ibm.com/jct03001c/press/us/en/pressrelease/27566.wss>
8. **Solid-state Storage Direct Attach option for IBM System x**
The EXP3000 supports 50 GB SATA Solid State Disk drives via the ServerAID controller adapter installed in the System x server
<http://www.ibm.com/systems/storage/disk/exp3000/index.html>
9. **Whitepaper on Tivoli Storage Productivity Center Storage Optimizer - An Introduction and Exploration**
<http://www-01.ibm.com/support/docview.wss?uid=swg21389271>
10. **IBM SSD Landing page**
<http://www-03.ibm.com/systems/storage/solutions/ssd/>

IBM News on Mainframe Tape Virtualization

TS7700 R1.5/R1.6

TS7680 ProtecTIER Deduplication Gateway for System z



IBM's Two Market Segments in Mainframe Virtual Tape

▪ **TS7700**


- Advanced grid technology
- Robust native replication
- Physical tape for long term retention
- Low cost tape/disk hybrid solution
- SMS tape support

▪ **TS7680**

- Data deduplication reduces disk requirements
- Clustered solution for high availability as single foot print
- Reduce / eliminate physical tape requirements
- Up to 1 PB native disk capacity
- SMS tape support / integration

TS7700 4-Way Grid

- Supports up to 4 clusters within a TS7700 Grid
 - **High availability solution** for production site and disaster recovery (DR) site
 - Existing replication / cluster selection policies used to replicate data between all 4 clusters
 - New *Cluster Families* function for improved decisions on => replication and => tape volume cache selection
 - Joining clusters into an existing 2- or 3-way Grid to make a 4-way is supported at GA !
 - Merging of two existing 2-way configurations to create a 4-way configuration **will not be supported at GA**
 - Maximum of 1 million log. volumes supported



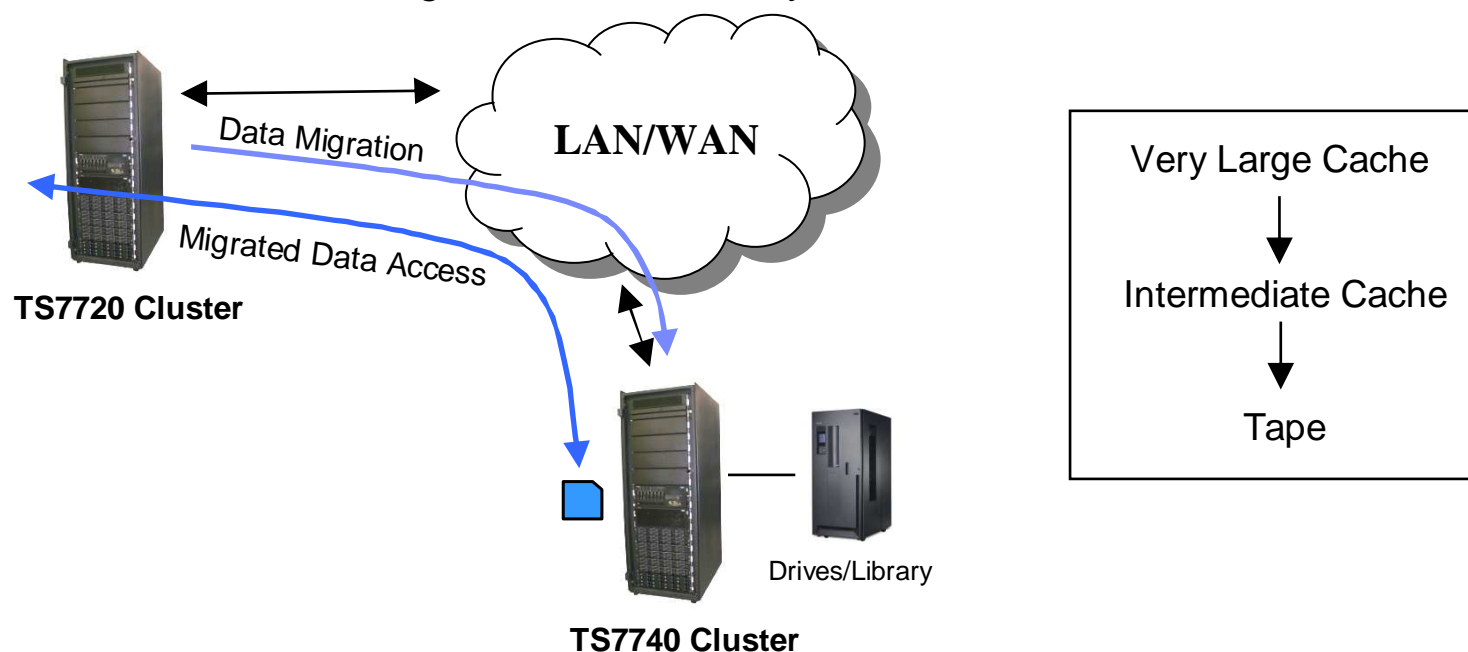
TS7720/TS7740 Hybrid Grid

- TS7720s and TS7740s can be configured within **same** Grid domain
 - Data contained within any TS7700 location can be accessed by any peer TS7700 location
 - independent of TS7700 type
 - Different combinations allow different **High Availability** and **Performance** characteristics (!):
 - TS7740: high capacity physical tape backend, copy export capability
 - TS7720: high performance “read-hits” through larger tape volume cache
 - Grid provides high availability at local and/or remote sites
 - Tape Volume Cache (TVC) Selection
 - TVC will continue to favor a cluster that has the volume resident in cache
 - New *Volume Removal Policy* manages migration of data from the TS7720 sites to TS7740 sites



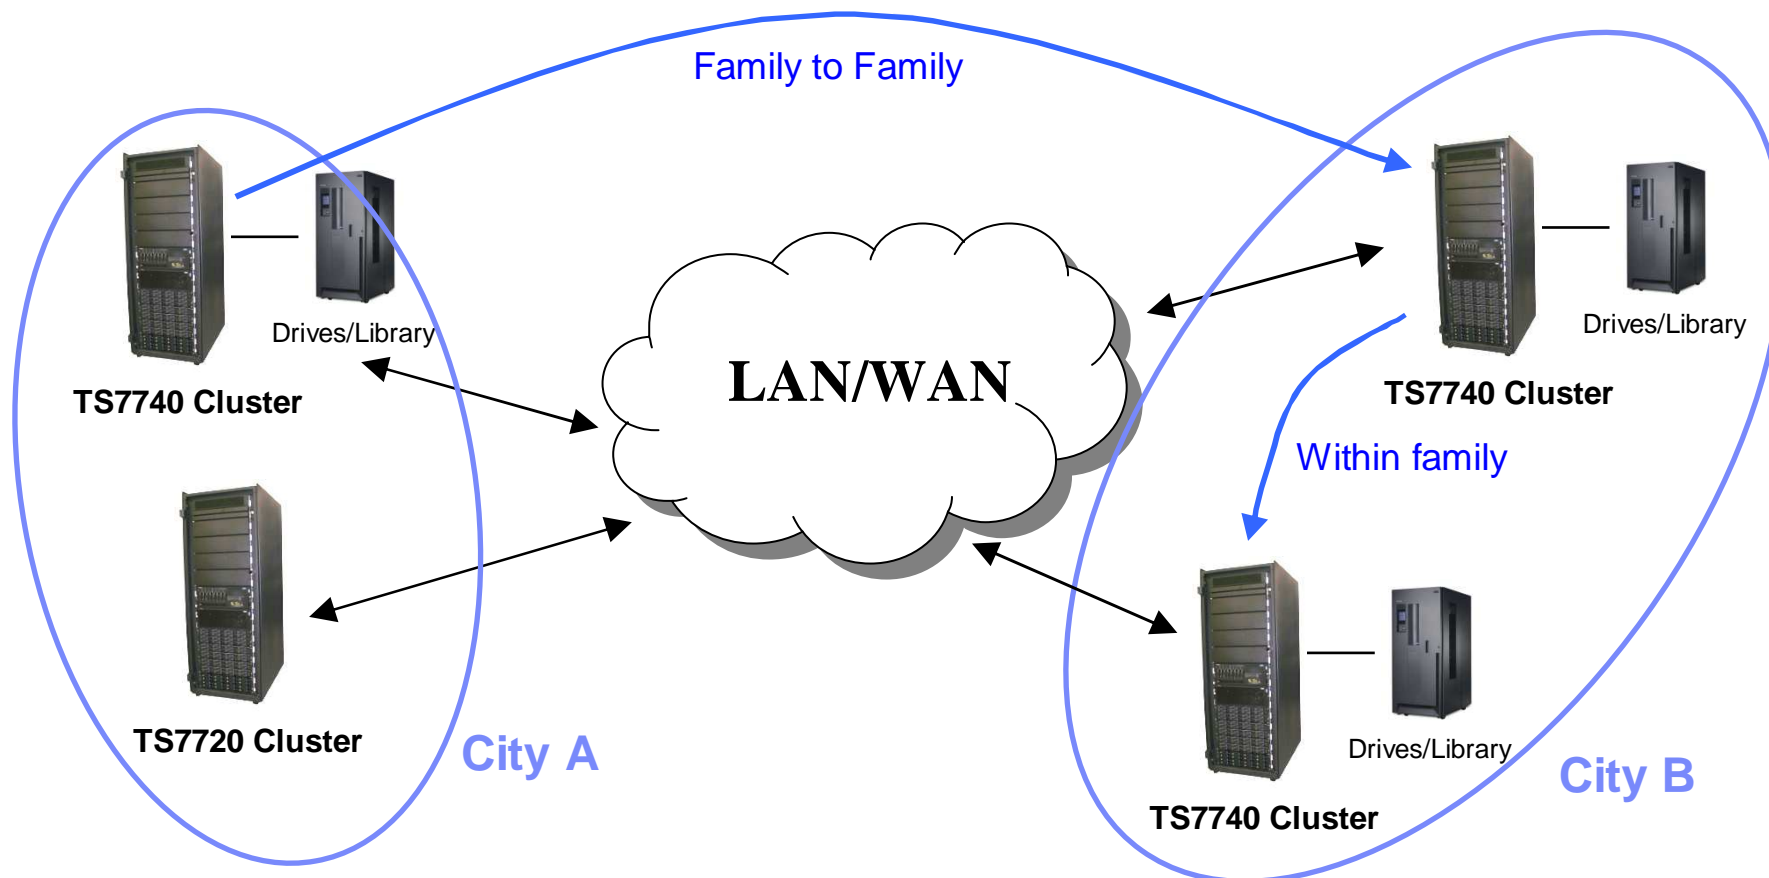
Hybrid Grid Configurations - Summary

- Disk-Centric and Cost Effective Hybrid Grid Configurations
- Automatic Volume Migration and Cache Space Management of the TS7720's cache
 - Volumes are copied from a TS7720 to a TS7740 through normal copy policies
 - When space is needed, the least recently accessed volumes in the TS7720's cache that have been copied to a TS7740 are removed from the TS7720's cache
- Migrated volumes remain accessible through the TS7720
 - TS7720 uses the grid links to remotely access the volume data in the TS7740



TS7700: R1.6 Cooperative Replication

- Copy management uses family information to optimize long-distance copy links
 - For deferred mode copies
 - Prioritizes getting one copy to each family before making copies to members in a family
 - Executes single copy between families, then local copies between family members.



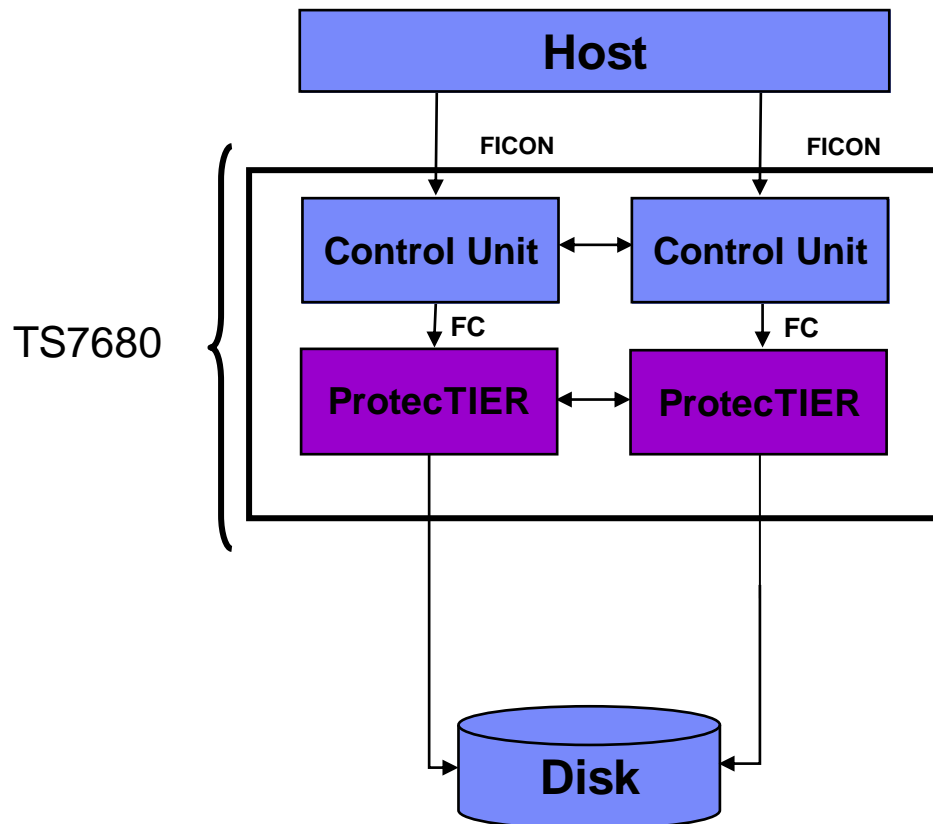
- Support for compliant logical worm (LWORM) tape volumes or a virtual equivalent of “write once read many” (WORM) compliant tape media through TS7700 software tape emulation
 - Host will view the TS7700 as a logical worm compliant library
 - Previously written volumes cannot be upgraded to logical worm volumes !
 - Host software changes will exploit and assist in the management of the logical worm volumes
 - The volume is bound as a LWORM by associating it with a **Data Class construct** indicating WORM and then performing a WRITE from beginning of tape (BOT)
 - Usable in the TS7720 & TS7740

IBM TS7680 ProtecTIER Deduplication Gateway for System z



TS7680 Overview


- IBM System Storage TS7680 – ProtecTIER deduplication Gateway for System z
 - Utilizes enterprise control unit with ProtecTIER to provide a FICON tape device interface
 - Uses the ProtecTIER Gateway cluster for fast, efficient de-duplicated backend storage
 - Uses the availability of the ProtecTIER Gateway to provide high available access to the virtual volumes on disk



- Appears as 3592-C06 device type
- Retains DFSMS functionality
- 128 virtual device support per control unit
- Integrated library mgmt functionality
- Scratch processing delete function

TS7680 Data Deduplication for System z

New Data Stream



HyperFactor



System z


FICON



TS7680

FC

Repository



Disk Arrays

Only a 4GB index is needed to map 1PB of physical disk
"Filtered" data

Dedupe example in a “normal” backup environment



Characterized by: predictable workload volumes, predictable change rates, and known retention times

Multiple generations of a backup policy



Generation 1: 10TB
100% new data; 3:1
Compression

→ 3.3 TB physically written



Generation 2: 10TB
15% unique data; 3:1 compr.

→ .5 TB physically written (1.5 TB / 3)



Generation 3: 10 TB
15% unique data; 3:1 compr.

→ .5 TB physically written

Repository after 3 cycles:

“Nominal” data: **30 TB**
 Physical capacity used: **4.3 TB**
 Dedupe ratio after 3rd gen: **~7:1**

TS7680 Solution

- Scalable to large capacities – large disk arrays and deduplication
 - Based on TS7650G
 - Scalable to 1 PB
 - Disk array of customer choice (DS8000, DS5000, XIV, non-IBM etc.)
- Highly available in a single footprint - redundant controllers sharing single disk repository
 - Concurrent code installs
- Good performance capacity
 - Inline data deduplication
 - Up to 900 MB/s
- Disaster recovery capabilities (Future)
 - Disk replication (unlikely used, available today via disk-subsystem methods)
 - “Native replication” (minimized IP bandwidth requirements) most likely

TS7680 Solution

▪ **Configuration**

- Up to 1,000,000 logical volumes
- Each virtual volume has 100 GB capacity (Host capacity)
- FICON attach– up to 8 ports with 256 logical paths/port
- Up to 256 (2 x 128) virtual drives

▪ **Library manager function integrated**

- Automated Tape Library (ATL) support (MF PLF command support)
- Automated return to scratch processing
- Library Operational State Change notification

▪ **Data deduplication up to 25x factoring (depends on actual user data)**

- In-line deduplication
- Deduplication ratios are data dependant
- Studies indicate System z data will dedupe and 6,5:1 to 10:1 are reasonable assumptions

Questions



Trademarks and Disclaimers



IBM Corporation 1994-2009. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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