IBM System Storage DS8000

+ Code R4.0 updates

A "*New Standard*" *for* Storage Leadership

Dominique SALOMON
European Product Introduction Center (PIC) Leader
Storage Specialist Certified – IBM Montpellier
✉ **Dominique.Salomon@fr.ibm.com**

Introducing IBM® TotalStorage® DS8000

"Game Changing" Capabilities!

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

- **Setting a "*New Standard*" in Cost Effectiveness**
  - ***Balanced Performance*** – **Up to 7X ESS Model 800**
  - ***Linear Scalability*** – **Up to 0.5 PB (with FATA disks)**
  - ***Integrated Solution Capability*** – **Storage System LPARs**
  - ***Flexibility*** – Dramatic addressing enhancements
  - ***Extendibility*** – Designed to add/adapt new technologies
  - ***Storage Management*** – All New Management Tools
  - ***Availability*** – **Designed for 24X7 environments**
  - ***Resiliency*** – Industry Leading Copy and Mirroring Capability
  - ***Long Term Cost*** – Four Year Warranty
    Model to Model Upgradeability
- **Delivered through**
  - Server/Storage Integration – **POWER5+™** Technology
  - Exploitation of **IBM Virtualization Engine™ Technology**
  - Innovation by leveraging IBM technology leadership
  - Extension of a proven microcode base to offer stability but allow exploitation of new technologies
  - Timely integration of new technologies

DS8000 Turbo Hardware Overview

- **2-Way (Model 8100)**
  - Two dual processor servers (POWER5+)
    - Up to **128GB Cache** (16, 32, 64 or 128 GB)
  - 8 to 64 4Gb FC/FICON – 4 to 32 ESCON Ports
  - 16 to **384 Disks**
    - Intermixable 73/146/300 GB 15Krpm, 146/300GB 10Krpm
  - Physical capacity from 1.1TB **up to 192TB**
    - (384 x 500 GB FATA Disks)

- **4-Way (Model 8300)**
  - Two four processor servers (POWER5+)
    - Up to **256GB Cache** (32, 64, 128 or 256 GB)
  - 8 to 128 4Gb FC/FICON – 4 to 64 ESCON Ports
  - 16 to **1024 Disks**
    - Intermixable 73/146/300 GB 15Krpm, 146/300 GB 10Krpm
  - Physical capacity from 1.1TB **up to 0.5PB**
    - (1024 x 500 GB FATA Disks)

DS8000: Hardware Architecture

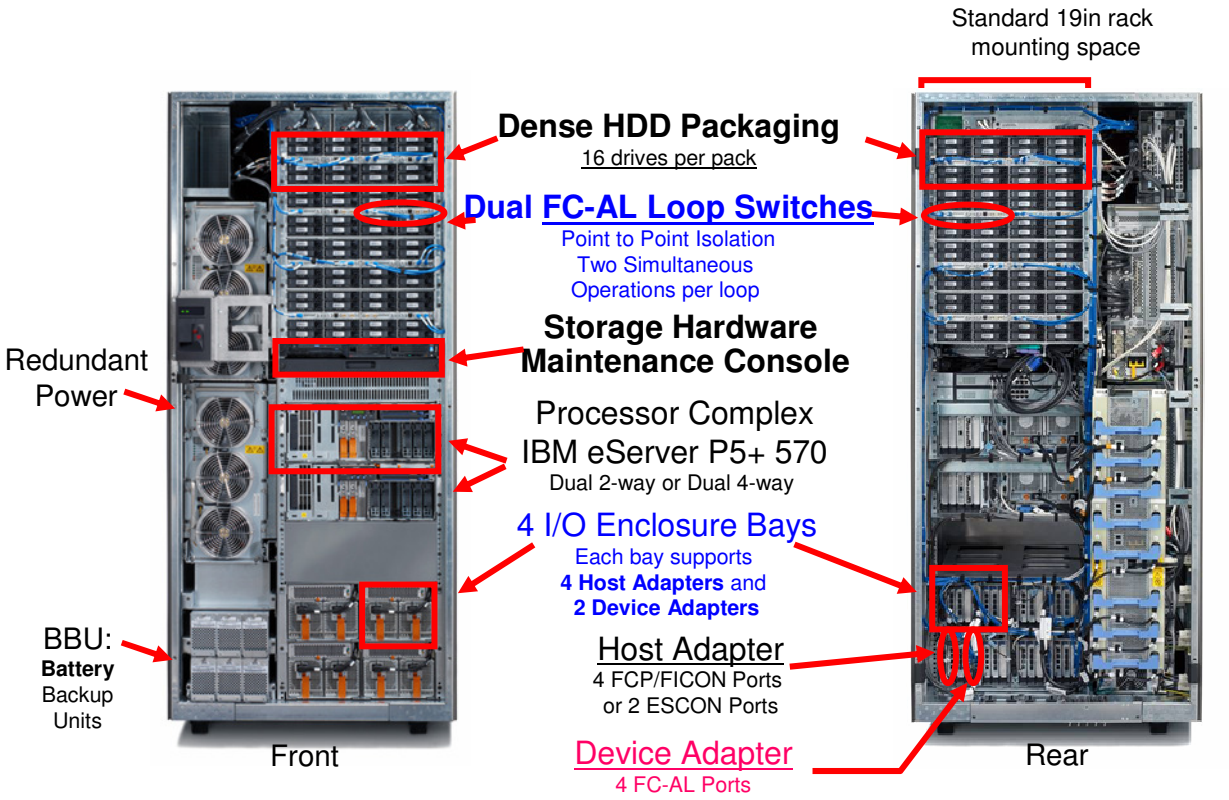

PowerPC™ Server Based Technology

- The server-based architecture brings continuous improvement
  - Performance
  - Scaling
  - Reliability, Availability, Serviceability
  - Features
  - Cost

while preserving the code base
75% of DS8000 code is running in the ESS 800 today
97% of DS6000 code is running in the DS8000

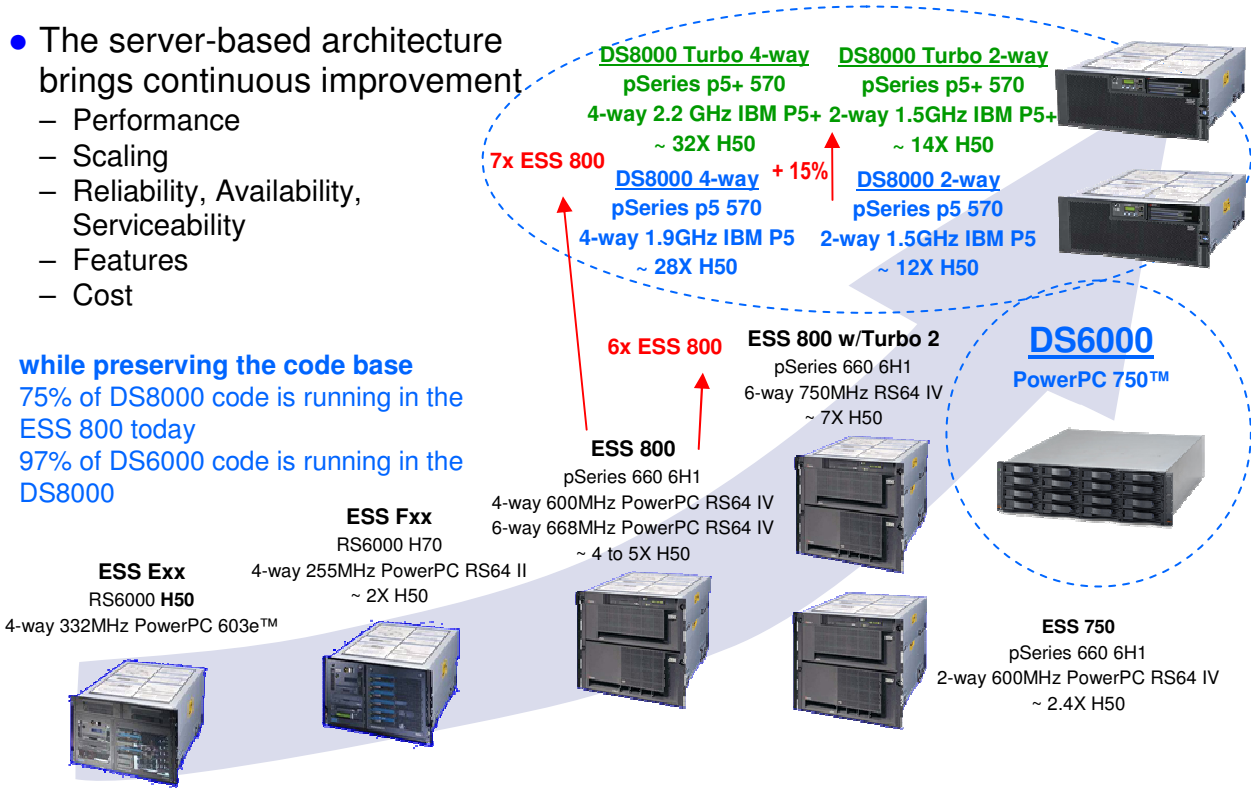

DS8000 - Configuration “Base Frames”

2-way System

128 disks
37 / 64 TB
64 ports
4 I/O drawer
128 GB cache

4-way System

128 disks
37 / 64 TB
64 ports
4 I/O drawer
256 GB cache

DS8000 - Configuration with “expansion frame”

2-way System

384 disks
115 / 192TB
64 ports
4 I/O drawer
128 GB cache

4-way System

384 disks
115 / 192TB
128 ports
8 I/O drawer
256 GB cache

DS8000 - Configuration Maximum

DS8300 - 320 TERABYTES Configuration

4-way System

640 disks
320TB
128 ports
8 I/O drawer
256 GB cache

DS8300 - 500 TERABYTES Configuration

4-way System

3rd and 4th expansion frames

1024 disks
0.5PB
128 ports
8 I/O drawer
256 GB cache

DS8000 supported Operating Systems

- IBM :
  - POWER : OS/400 , i5/OS, Linux & AIX
  - System z : z/OS, z/VM & Linux
- Intel Servers :
  - Windows, Linux, VMware & Netware
- Hewlett-Packard :
  - HP UX
  - AlphaServer : Tru64 UNIX & OpenVMs
- SUN :
  - Solaris
- Apple Macintosch
  - OSX
- SGI Origin Servers :
  - IRIX (April 29 2005)
- Fujitsu Primepower
- + many other OS

Check the **DS8000 series Interoperability Matrix** for complete and updated Information on that subject.

DS8000 Hardware Architecture

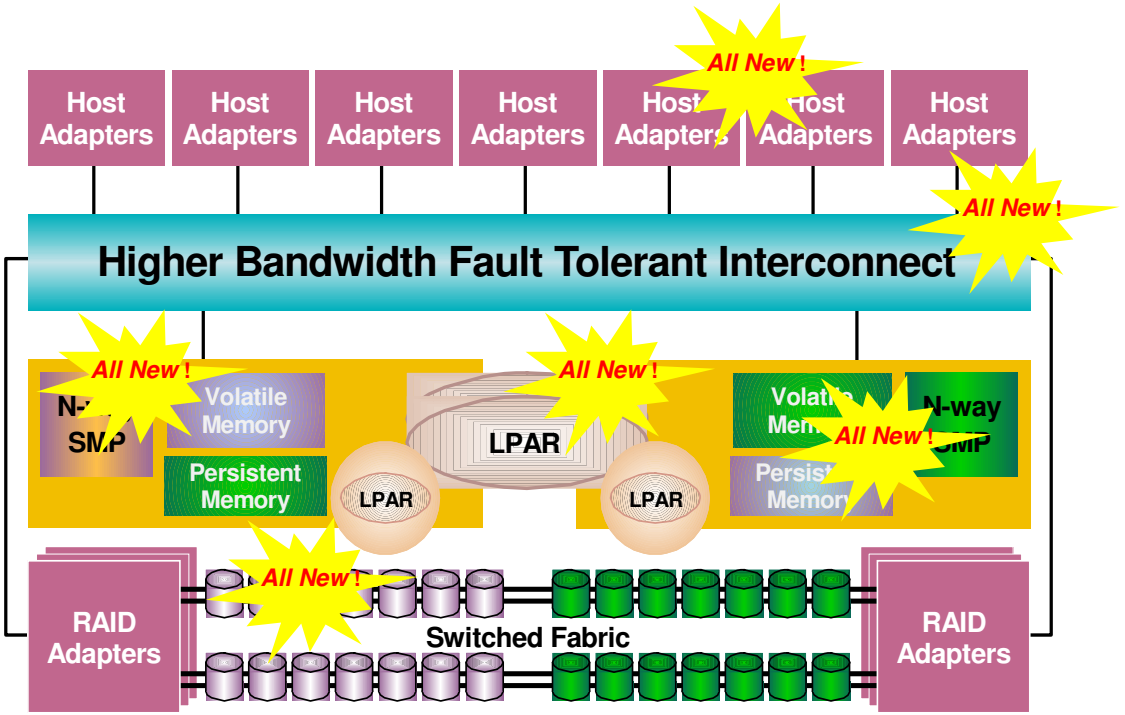

DS8000 eServer p5+ 570 servers

- **DS8000 Turbo** (931, 932 & 9B2) are using **POWER5+** processors
  - 15 % performance improvement / previous models (POWER5)
  - **2.2 GHz for POWER5+ 2 & 4 way**

- Cache memory from 16GB-256GB

- Persistent Memory 1GB to 8GB : dependant on cache size
  - **Battery backed** for backup to internal disk (4GB per server)

- The POWER5+ processor supports Logical Partitioning (LPAR)
  - The p5 hardware and **Hypervisor** manage the real to virtual memory mapping to provide robust isolation between **LPARs**.

- **On DS8000 4-way models, possible to configure each server in 2 LPARs (Storage System LPARs)**
  - Only 50-50% before code R4.0 (2 CPUs / LPAR)
  - now 50-50%, 25-75% or 0-100% with last code R4.0

DS8000 Exploitation of Storage System LPARs

- **Yields Reduced TCO**
- **Provides tangible customer benefits***
  - Exploit investment of a large high-performance storage subsystem
  - Increased storage management productivity and flexibility
  - Simplified infrastructure
  - Reduced floor space
  - Dynamic allocation of resources and workload balancing
  - Greater scale out capability
  - Service Level Attainment
- **IBM leadership**
  - Value to customer
  - Extensibility in future

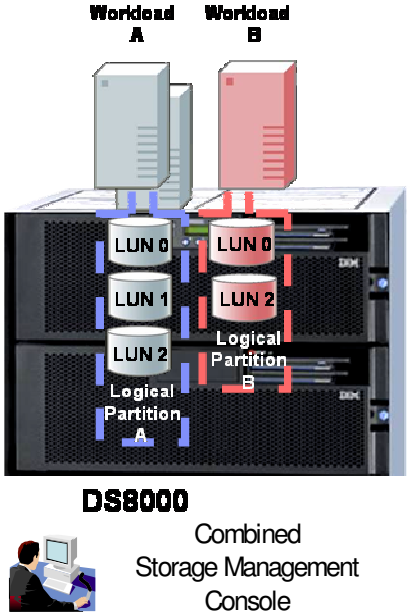

Separate Segregated Storage Sub-Systems

* Third party blind survey, IBM customer councils and individual customer briefings

DS8000 Data flow

- The normal flow of data for a write is the following :
  - 1. Data is written to cache memory in the owning server.
  - 2. Data is written to NVS memory of the alternate server.
  - 3. The write is reported to the attached host as having been completed.
  - 4. The write is destaged from the cache memory to disk.
  - 5. The write is then discarded from the NVS memory of the alternate server.

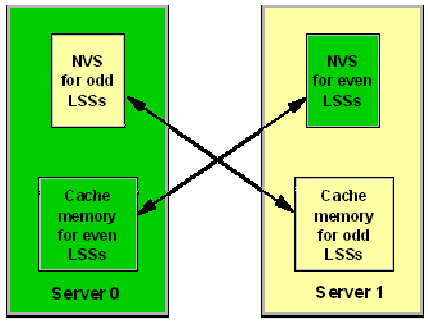

Under normal operation, both DS8000 servers are actively processing I/O requests

DS8000 Server Failover

- Normal flow of data for a write:
  1. Data is written to cache memory in the owning server.
  2. Data is **written to NVS memory of the alternate server**.
  3. The write is reported to the attached host as having been completed.
  4. The write is destaged from the cache memory to disk.
  5. The write is then discarded from the NVS memory of the alternate server.
- After a failover, **remaining server is processing all I/Os but with cache and NVS divided by two** (half for each server load)

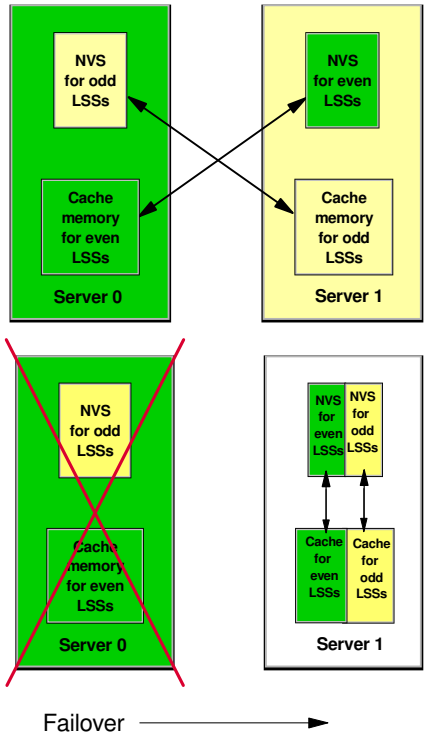

DS8000 New SARC Cache Algorithm

- **SARC** : Sequential Prefetching in Adaptive Replacement Cache
- **Best caching algorithms in industry**
- **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

Benefits of adaptive replacement caching

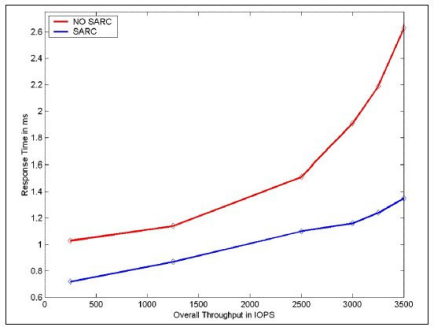

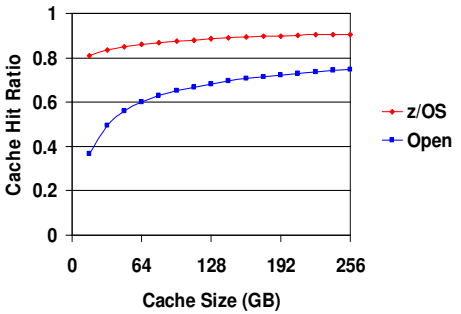

Lower cache-to-backstore ratios with outstanding service times

Management Interfaces to DS8000

- **IBM TotalStorage DS Storage Manager (DS-SM: Web based)**
  - Program Interface to perform logical configurations & copy services Mgt functions
  - Accessed through **Web browser**

- **DS command-line interface (CLI: Script based)**
  - OPEN hosts invoke and manage **FlashCopy**, **Metro** & **Global Mirror** functions
    - **Batch processes and scripts**
    - Check storage unit configuration & perform specific application functions

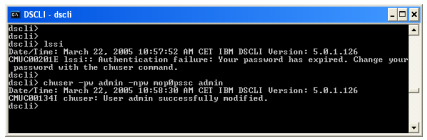

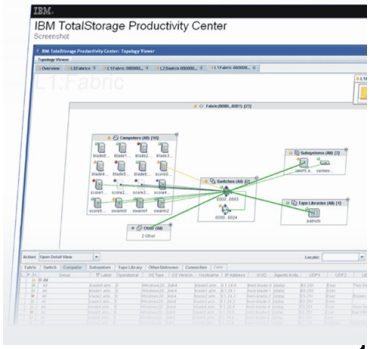

- **DS8000 can be also managed from IBM TotalStorage Productivity Center (TPC)** management software or SSPC Console

IBM – Resilience Offerings on DS8000

Local Point in time Copy :

- **IBM FlashCopy**
  - For local data copy capabilities

Remote Mirroring :

- **IBM TotalStorage Metro Mirror**
  - For **synchronous** remote copy
- **IBM TotalStorage Global Copy**
  - To Move data at a remote location in asynchronous mode
- **IBM TotalStorage Global Mirror and z/OS Global Mirror**
  - For asynchronous remote copy

Disk Storage Mirroring

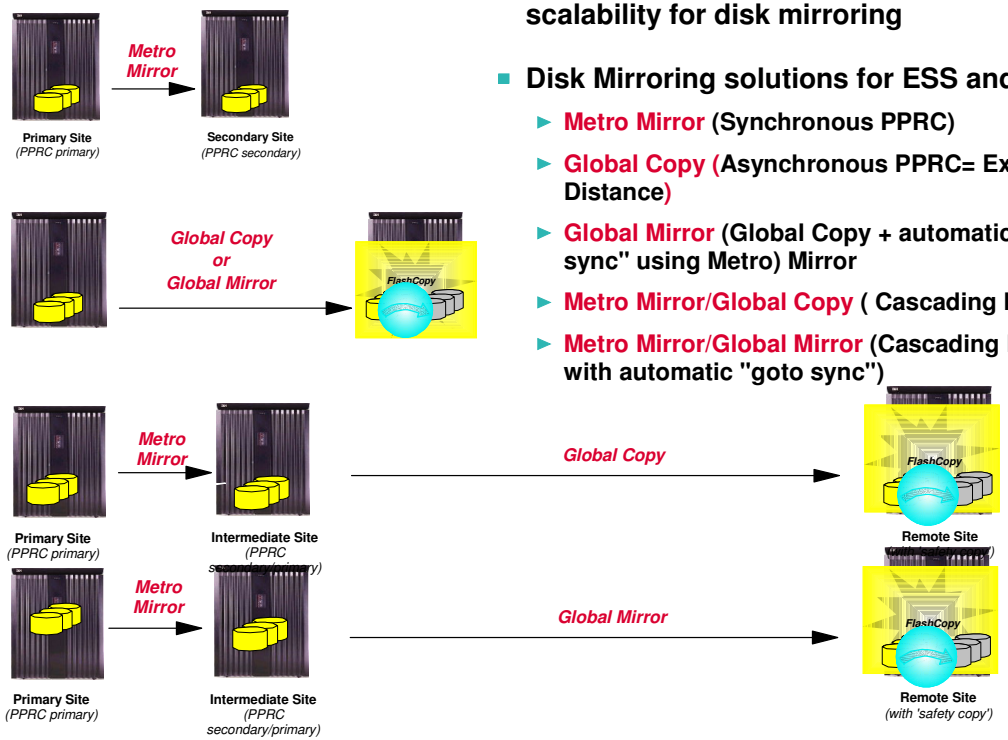

- **Leading solution : performance, and scalability for disk mirroring**
- **Disk Mirroring solutions for ESS and DS**
  - **Metro Mirror (Synchronous PPRC)**
  - **Global Copy (Asynchronous PPRC= Extended Distance)**
  - **Global Mirror (Global Copy + automatic "goto sync" using Metro) Mirror**
  - **Metro Mirror/Global Copy (Cascading PPRC)**
  - **Metro Mirror/Global Mirror (Cascading PPRC with automatic "goto sync")**

IBM® System Storage® DS8000

New code R4.0 (Sep 2008)

DS8000 R4.0 code new features

The IBM System Storage™ DS8000 series **R4.0 new code** provides new functional capabilities, allowing clients to better choose the combination of price and efficiency that is right for their application needs.
New options include:

- **RAID-6:** RAID-6 allows for additional fault tolerance by using a second independent distributed parity scheme (dual parity).
- **450 GB 15,000 rpm Fibre Channel Disk Drive:** Provides additional capacity and another price/performance option for transaction processing workloads.
- **Variable LPAR:** Variable LPAR provides the ability for Dual Shark Images where one image has more processor and cache than the other for increased flexibility.
- **Extended Address Volumes"** Extended Address Volumes (EAV) extends the addressing capability of IBM System z environments to support volumes that can scale up to approximately 223 GB (262,668 cylinders). This capability can help relieve address constraints to support large storage capacity needs.
- **IPv6:** The IBM System Storage DS8000 has been certified by the IPv6 Ready Logo Committee as meeting the requirements of the IPv6 Ready Logo program, indicating its implementation of IPv6 mandatory core protocols and the ability to interoperate with other IPv6 implementations.

RAID 6 Support

Value proposition:

Increasing DS8000's superior resiliency to support today's lower-cost, less-reliable, large-capacity drives

- **RAID 6 - offers protection from double disk failures by using a second independent distributed parity scheme (dual parity)**
  - Can tolerate loss of data in two strips of any stride without data loss
    - 2 Disk Failures
    - 1 Disk Failure + 1 Media Error
    - 2 Media Errors
- **Critical for larger drives (SATA/FATA), which have longer rebuild times and thus longer periods of potential exposure should two drives start failing**
- **Factory support only for R4.0**
- **Field support in R4.2**
- **RAID 6 will be included in the R4 microcode at no additional charge**

Extended Address Volumes

Value proposition:

EAV simplifies deployment of large System z installs

- **Increases support for System z volume sizes by over 400%!**
- **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 Dynamic Volume Expansion can allow non-disruptive migration to larger volume sizes**

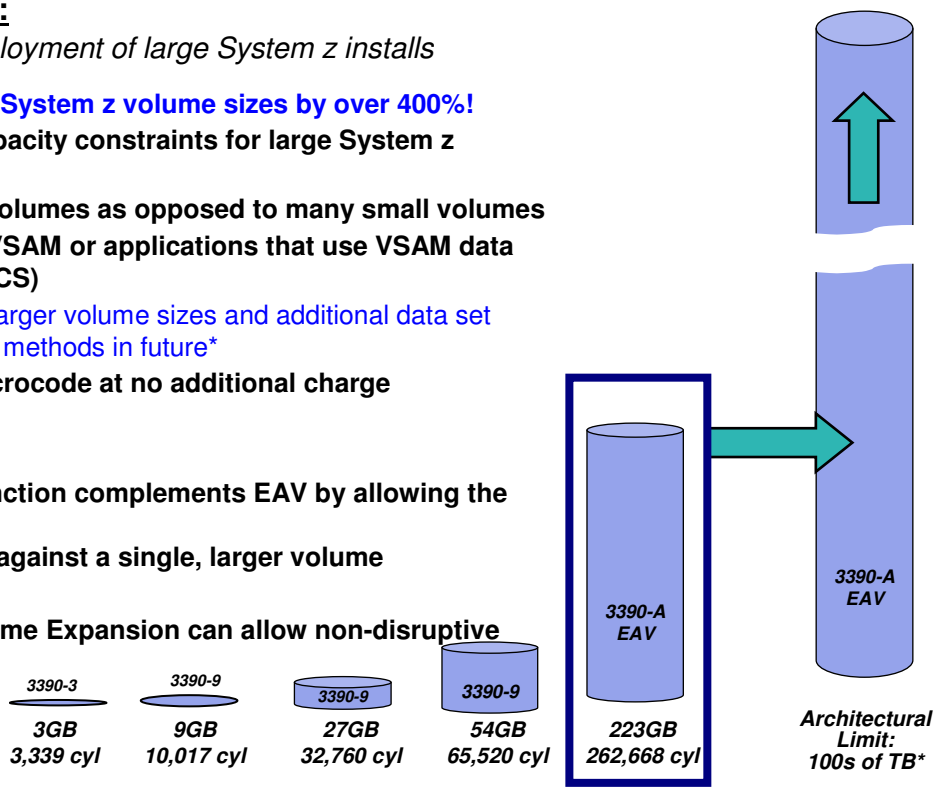

Variable LPAR

Value Proposition:

Advanced flexibility enables more granular virtualization and consolidation

- **LPARs now support 25%, 50%, and 75% of system resources to further enable:**
  - **Simplification through consolidation**
  - **Improved cost of ownership**
  - **Improved management efficiency**
  - **Reduced data center footprint**
  - **Dynamic allocation of resources**
  - **Service Level Management**
  - **Efficient workload balancing**
  - **High availability – storage image independence**

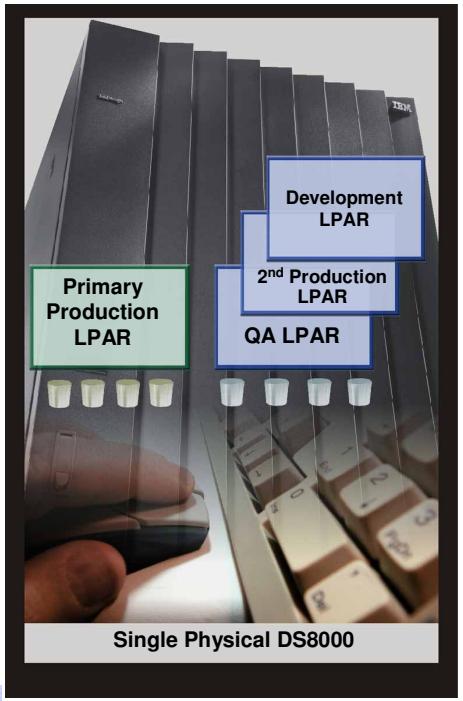

Added flexibility to support today's dynamic enterprise

Certified Secure Data Overwrite Services

Value proposition:

Securely erase data when storage systems are de-commissioned

- **Lab-based service to remove proprietary and sensitive data from storage systems when de-commissioned**
  - Supports ESS750, ESS800, and DS8000 storage systems
- **Data is erased by writing a series of patterns to every sector of the disks**
- **Done with special software that runs inside the DS8000**

Support for Internet Protocol V6 (IPV6)

Value Proposition

Increased addressing space, built-in QoS, and better routing performance and services

- **DS8000 is now certified to meet requirements of the IPv6-Ready Logo program, indicating its implementation of IPv6 mandatory core protocols and the ability to interoperate with other IPv6 implementations.**
  - **Larger address space**
    - **While IPv4 supports 4 billion addresses, IPv6 provides unlimited IP addresses**
  - **Better Security**
    - **IPv6 offers embedded encryption, authentication and data integrity safeguards**
  - **Efficient routing**
    - **The IPv6 Internet has much smaller and more efficient routing tables, which simplifies and brings more reliable packet direction**

Intermix of Drives Speeds in Same Enclosure

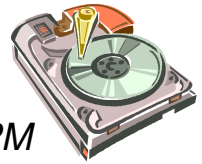

Value Proposition

Flexibility to combine older 10k RPM drives with 15k RPM drives to maximize space in drive enclosures

- **The ability to intermix 10k and 15k drives for customers with partially full DS8000 systems of older 10k RPM disk drives**
- **Provides investment protection by allowing 15k and 10k RPM drives within disk enclosures**

Support for Secure File Transfer Protocol (FTP)

Value Proposition

Higher security for added protection

- **DS8000 now supports the Secure File Transfer Protocol (FTP) for stronger protection when downloading microcode to DS8000 systems**
  - Secure FTP is an industry standard protocol for moving data securely
  - IBM field engineers will now be able to use Secure FTP instead of the standard File Transfer Protocol when installing microcode to the DS8000 Hardware Management Console

DS8000 Future Directions (2008)

Planned in 2008 :

- Release R4.1 :
  - z High Performance FICON

Planned in 2009 :

- First release (R4.2) :
  - **1TB SATA disks**
  - **Thin Provisioning**
  - **Disk Encryption**
  - Cache algorithm improvement
  - RAID6 Field update capability

- Second Release (R5):
  - **POWER6**
  - **Disk Encryption improvements**
  - Solid State Drives (SSD)

IBM DS8000 Disk Encryption Releases

- 1st Release :
  - One Encryption Group per Storage Facility Image (One Key on TKLM)
  - **No intermix of encrypting and non-encrypting DDMs**
  - New factory order only, no field MES
  - **Entire storage facility image is either all encrypted or all not encrypted** - selected when configure first rank

- 2nd Release
  - One Encryption Group per Storage Facility Image (One Key on TKLM)
  - **Each extent pool enabled for encryption or no encryption**
  - Encrypting extent pool must have ranks with encrypting disks
  - **Intermix of encrypting and non-encrypting DDMs**
  - Field MES of encrypting disks to existing DS8000s

BACKUP Slides

DS8000 4Gbps Host Adapters performances

New 4 Gb Host adapters are designed to improve by 50% single port throughput performance.

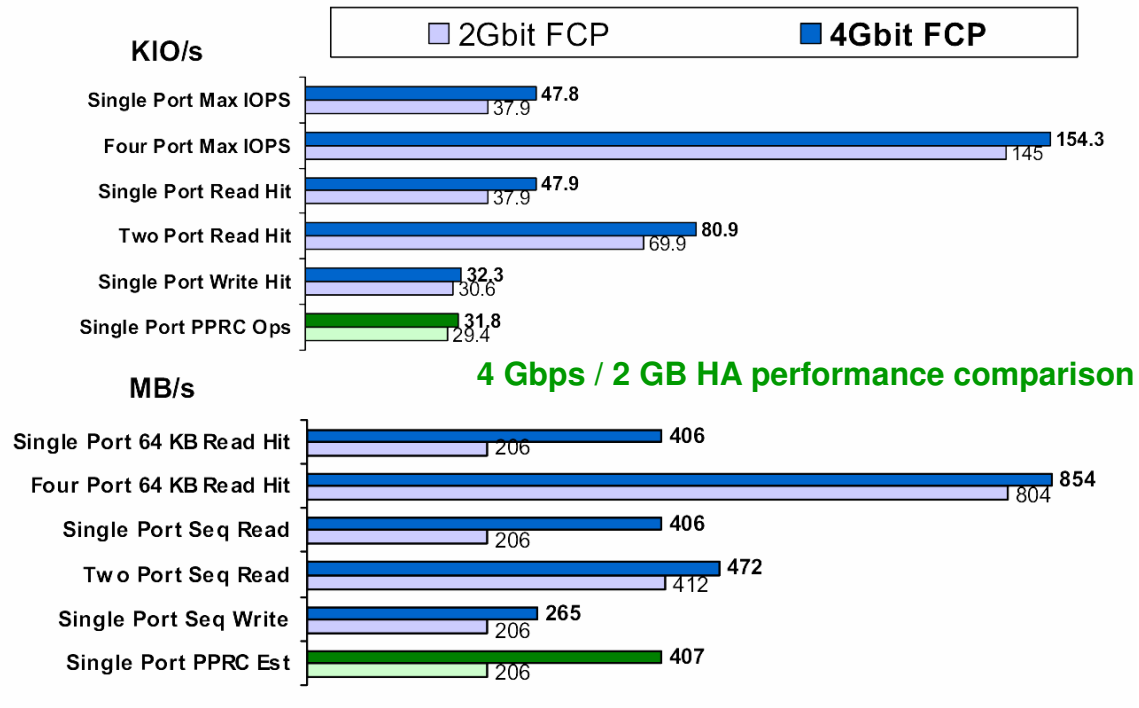

4 Gbps / 2 GB HA performance comparison

DS8000 Device Adapters

- Device Adapters support RAID-5 or RAID-10
- **FC-AL Switched Fabric topology**
- **FC-AL dual ported drives** are connected to FC switch in the Disk Enclosure backplane
- **Two FC-AL loops** connect Disk Enclosures to Device Adapters
- **Array across loops** is standard configuration option in DS8000
  - Two simultaneous I/O ops per FC-AL connection possible.
  - Switched FC-AL or SBOD (switched bunch of disks) used for back-end access
- Device Adapters are attached to a FC Switch with the enclosure
- **4 paths to each drive** : 2 FC-AL loops X dual port access

DS8000 Storage Enclosures

- Enclosure hold **16 DDMs**
  - **Dual ported FC-AL DDMs**
  - 73, 146, or 300 GB DDMs
    - 10 or 15K RPM
  - New FATA Disk drives of 500 GB / 7200 rpm are also supported in the same enclosures.

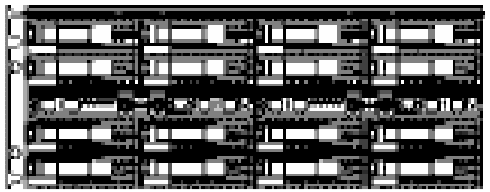

- Drives can be added in groups of 8 drives by DS8000 storage enclosure
- Enclosures act as a **FC switch** connecting drive using point to point connections

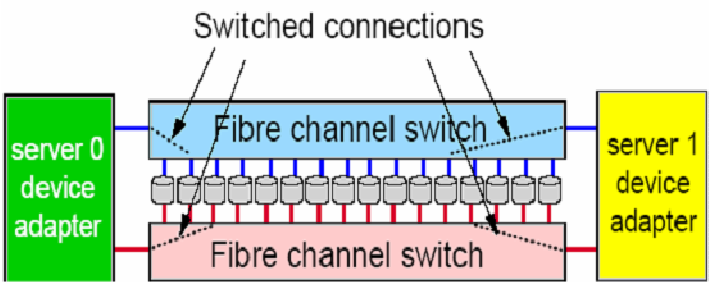

The picture above depicts four simultaneous and independent switched connections, one from each device adapter port.

DS8000 : Storage Enclosure & DA cabling

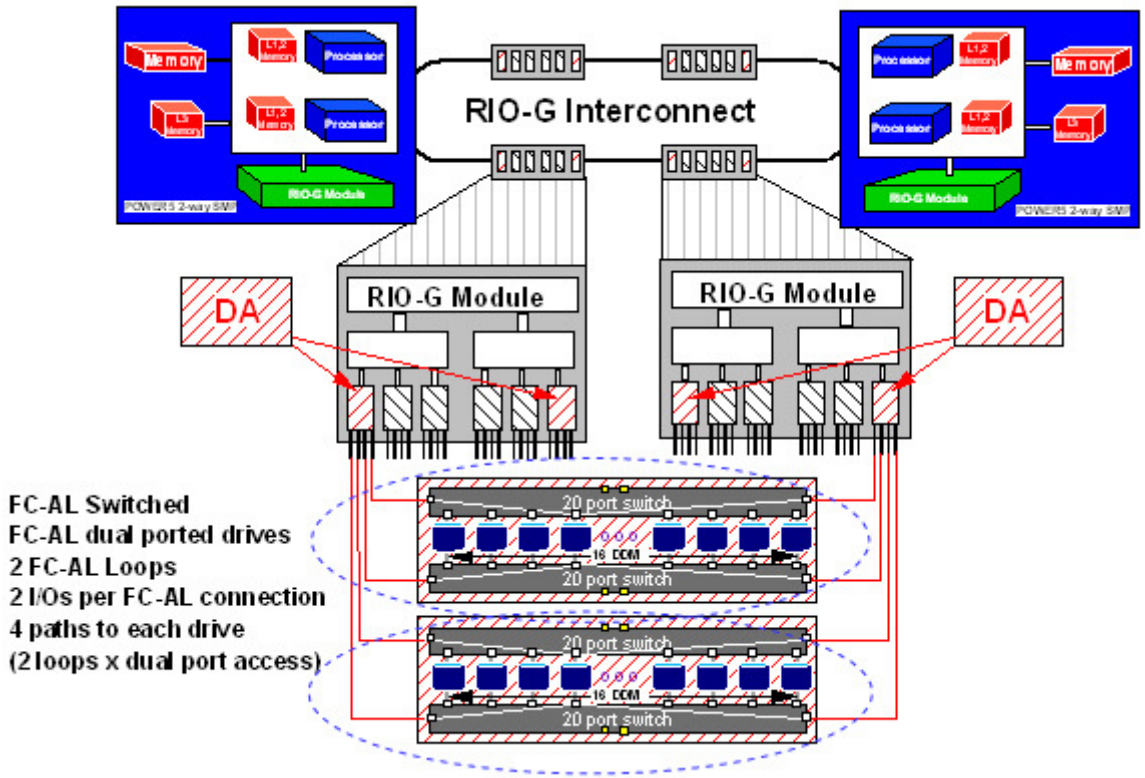

DS8000 OPEN Dual Port Host Attachment

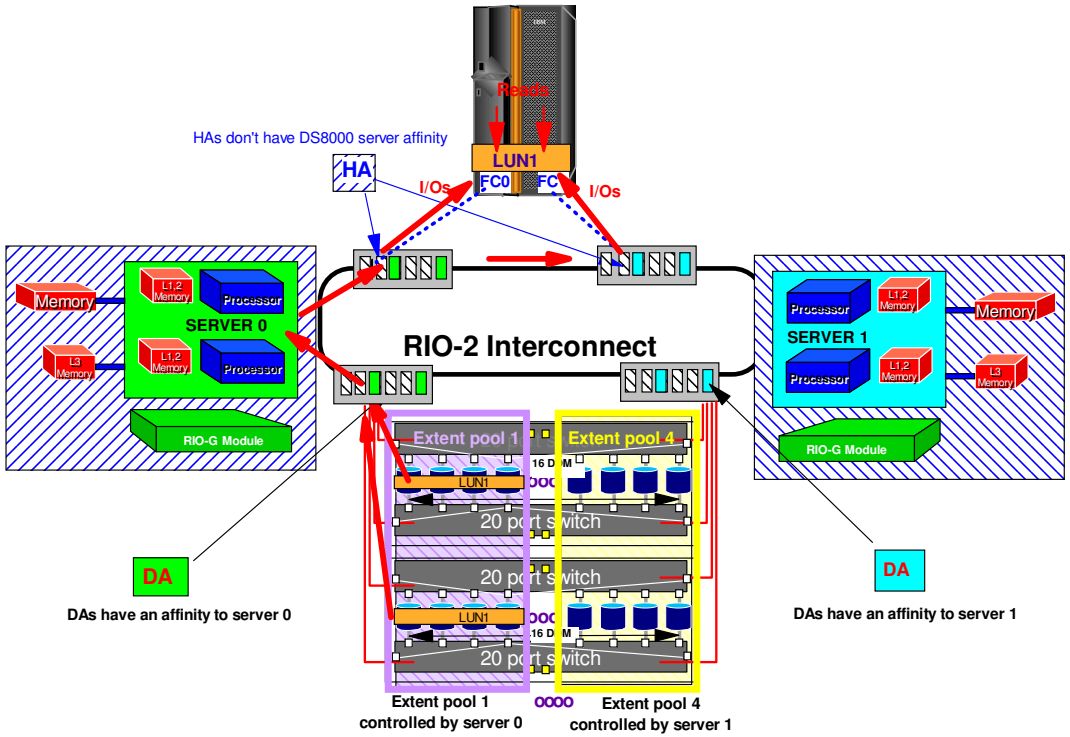

Volume Extents Spread Across RAID arrays

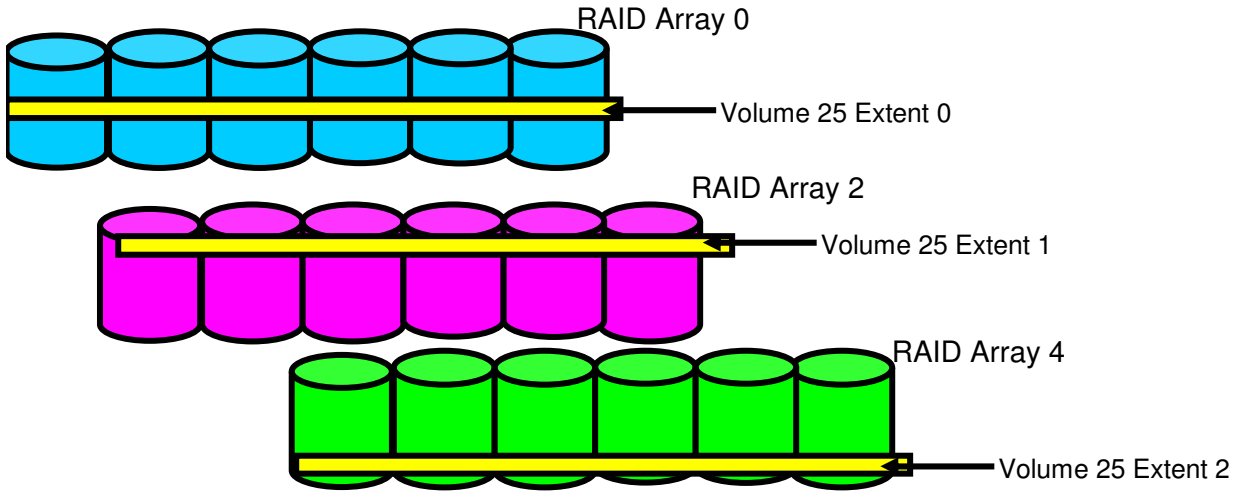

- **Extents striped across multiple RAID arrays**
- **Improves volume throughput for some workloads**

Host Connectivity (Open) : IBM SDD & MPIO

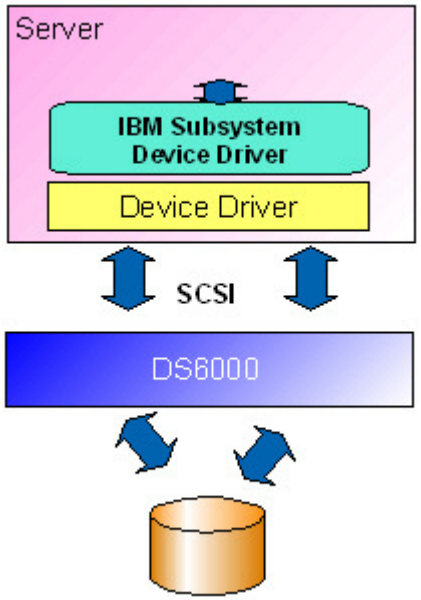

- **SDD provides the following functions:**
  - Enhanced data availability
  - Automatic path failover
  - Dynamic I/O load-balancing across multiple paths
  - Path selection policies for the host system
  - Concurrent download of licensed machine code

- **With DS6000 & DS8000, SDD is supported on the following OS :**
  - Windows,
  - Netware,
  - AIX,
  - HP-UX,
  - SUN Solaris
  - Linux
- Can coexist with **RDAC** (DS4000 multipath driver) on most OS as long they manage separate HBAs.
- Can not be used with most other multi-path drivers (ie Veritas, PV-Links, Powerpath)

MPIO 2.1.0 PCM (Path Control Module) is also supported for
AIX 5.2 ML5 (or later) and AIX 5.3 ML 1(or later)
!!! Default MPIO is not supported : **sddpcm** filesets are required.

New 450GB 15k RPM Fibre Channel Disks

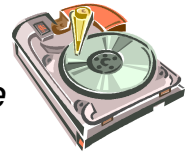

Value Proposition

Provides additional capacity and another price/performance option to enable client consolidation efforts

- **Higher capacity drives increase capacity by 50% over our largest previous fibre channel drives of 300GB**
  - Increased capacity from 307TB to 460TB with a single DS8000 fully loaded with fibre channel drives
- **Larger drive option helps enable consolidation of more storage on highly scalable DS8000 systems**
- **Similar to existing FC drive options, these drives are best for supporting mission-critical applications that serve multiple users, including online transaction processing and intensive database queries.**

1TB SATA 7.5k RPM Drive Support

Value Proposition:

Near-line storage for tier 2 applications that provide lower pric per GB

- **Where to implement SATA drives in tier one storage**
  - **Customers looking for low cost per GB and willing to take some availability risks / performance reduction**
  - **Bandwidth / streaming applications where transfer rate is more important than seek time**
  - **Some applications with low or limited IOPS performance requirements**

- **Where to implement SATA drives in tier two storage**
  - **Disk-based near-line storage between disk and tape**
  - **Disaster recovery target for remote replication**
  - **Fixed content / Managed Retention data**
  - **Temp work space for short-term processing**

Full Disk Encryption and Key Management

Value Proposition

Automatically s*ecure sensitive data at rest when it leaves the data center*

- Full disk encryption (FDE) drives
  - Encrypt data-at-rest with embedded encryption key and password authentication
- Storage system
  - Define secure volume groups, authenticate with the key source, and pass authentication key to the drive
- Key management service
  - Uses same proven key management as TS1120 tape drive to easily and securely manage keys
- Standards for interoperability
  - FDE management support via Trusted Computing Group security protocol
  - Working to create industry standards for the authentication key management protocol

IBM DS8000 Disk Encryption details

- 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
  - 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)

DS8000 Encryption Network Environment