

IBM DS8870 Enterprise Storage News and Review

Nick Clayton

Solution Architect for DS8000

IBM Systems Division



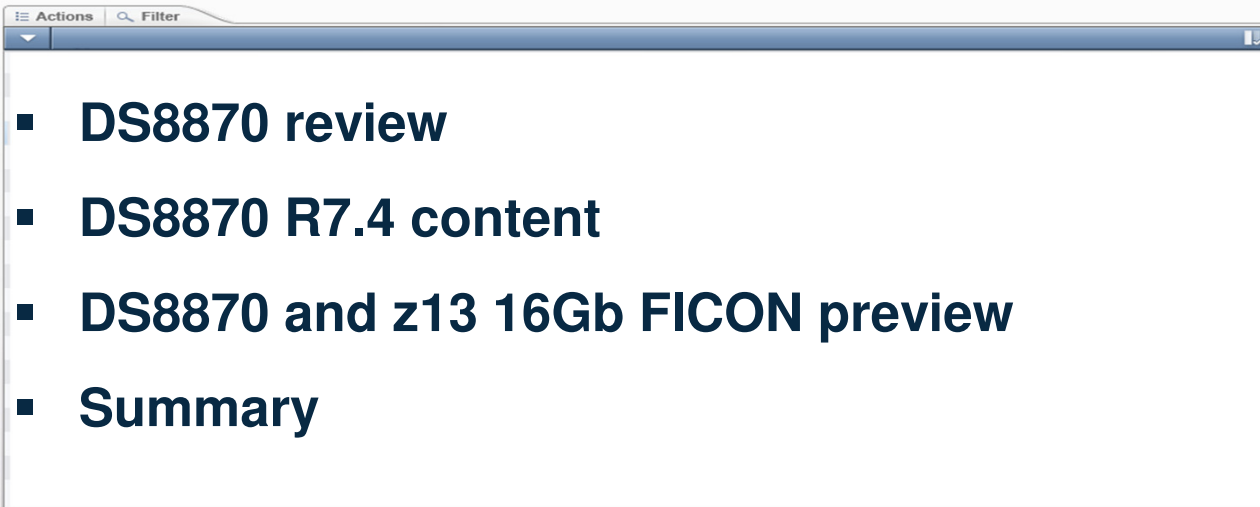
#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides education, professional networking and industry influence.



Agenda

- 
- A screenshot of a presentation slide with a blue header bar containing 'Actions' and 'Filter' icons. The slide content is a bulleted list:
- **DS8870 review**
 - **DS8870 R7.4 content**
 - **DS8870 and z13 16Gb FICON preview**
 - **Summary**



Current DS8870 Family Models

DS8870 – Enterprise Class

- Dual POWER7+ controllers (2-way, 4-way, 8 -way, 16-way)
- 8 Gb/s host adapters
- 8 Gb/s device adapters
- 2.5" Enterprise SAS-2 and SSD drives , 3.5" Nearline drives
- All drives encryption capable
- Up to 1,536 2.5" drives plus 240 1.8" Flash drives
- Up to 1 TB processor memory
- All upgrades concurrent

DS8870 – Business Class

- Dual POWER7+ controllers (2-way, 4-way, 8 -way, 16-way)
- 8 Gb/s host adapters
- 8 Gb/s device adapters
- 2.5" Enterprise SAS-2 and SSD drives , 3.5" Nearline drives
- All drives encryption capable
- Up to 1,056 drives plus 240 1.8" Flash drives
- 192 drives per DA pair populated before next DA pair is needed
- Up to 1 TB processor memory
- All upgrades concurrent but no upgrade path to Enterprise Class

DS8870 – All Flash

- Dual POWER7+ controllers (8-way, 16-way)
- 8 Gb/s host adapters
- 8 Gb/s device adapters
- Up to 240 1.8" Flash drives
- All drives encryption capable
- 512 GB or 1 TB processor memory
- All upgrades concurrent but no change to Enterprise or Business Class




DS8870 Hardware Review

High Performance Flash Enclosures provide significant performance improvements with Flash optimized RAID

Highly resilient DC-UPS power supply with 98% power efficiency and 4 second power loss ride-through (50 seconds with ePLD)

High performance dual core 8Gb Host and Device Adapters provide CPU offload



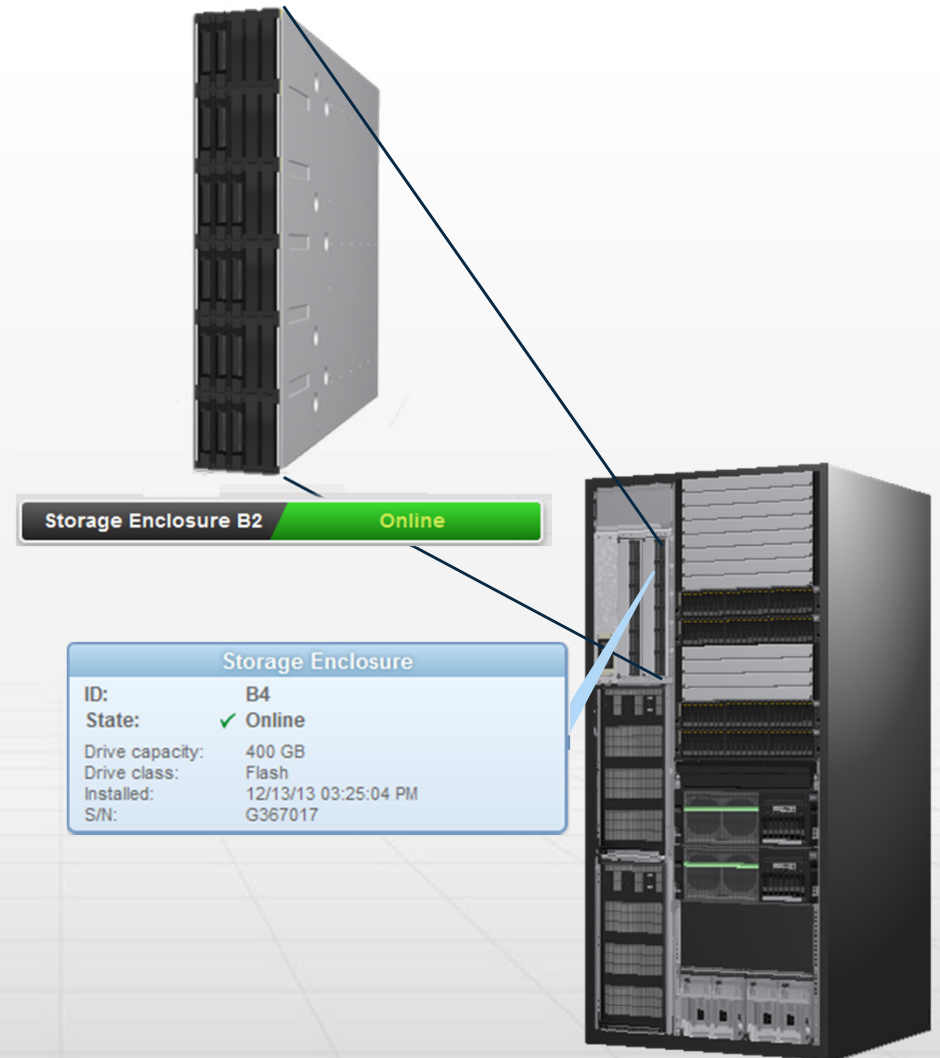
Switched Fibre Channel backend connectivity to 2.5" drives and 3.5" Nearline drives

Dual POWER7+ processors with up to 32 cores and 1TB memory in total

PCI-e fabric provides high speed resilient internal connectivity between components


High Performance Flash Enclosure

- Flash Optimized RAID for ultra-high performance
- 1U Storage Enclosure with up to 30 1.8" 400GB Flash drives
- 9.1 TiB usable RAID5 capacity per enclosure
- Fully redundant with integrated power and cooling
- Supports encryption using FDE Flash drives



HPFE Performance

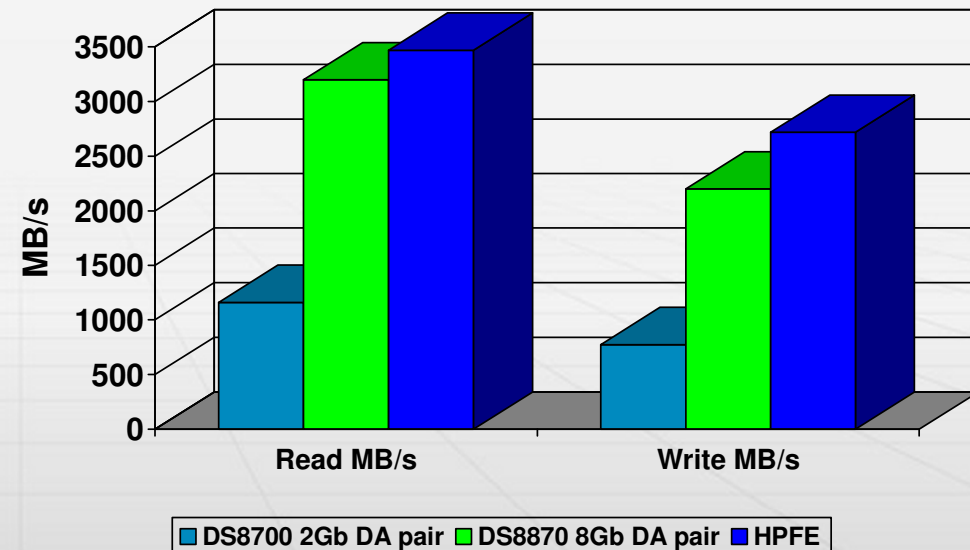
Random IO



- Single High Performance Flash Enclosure has modest increase in sequential throughput capability compared to DS8870 device adapter
 - 16% for sequential read
 - 23% for sequential write
- DS8870 device adapter already provided significant sequential throughput improvements

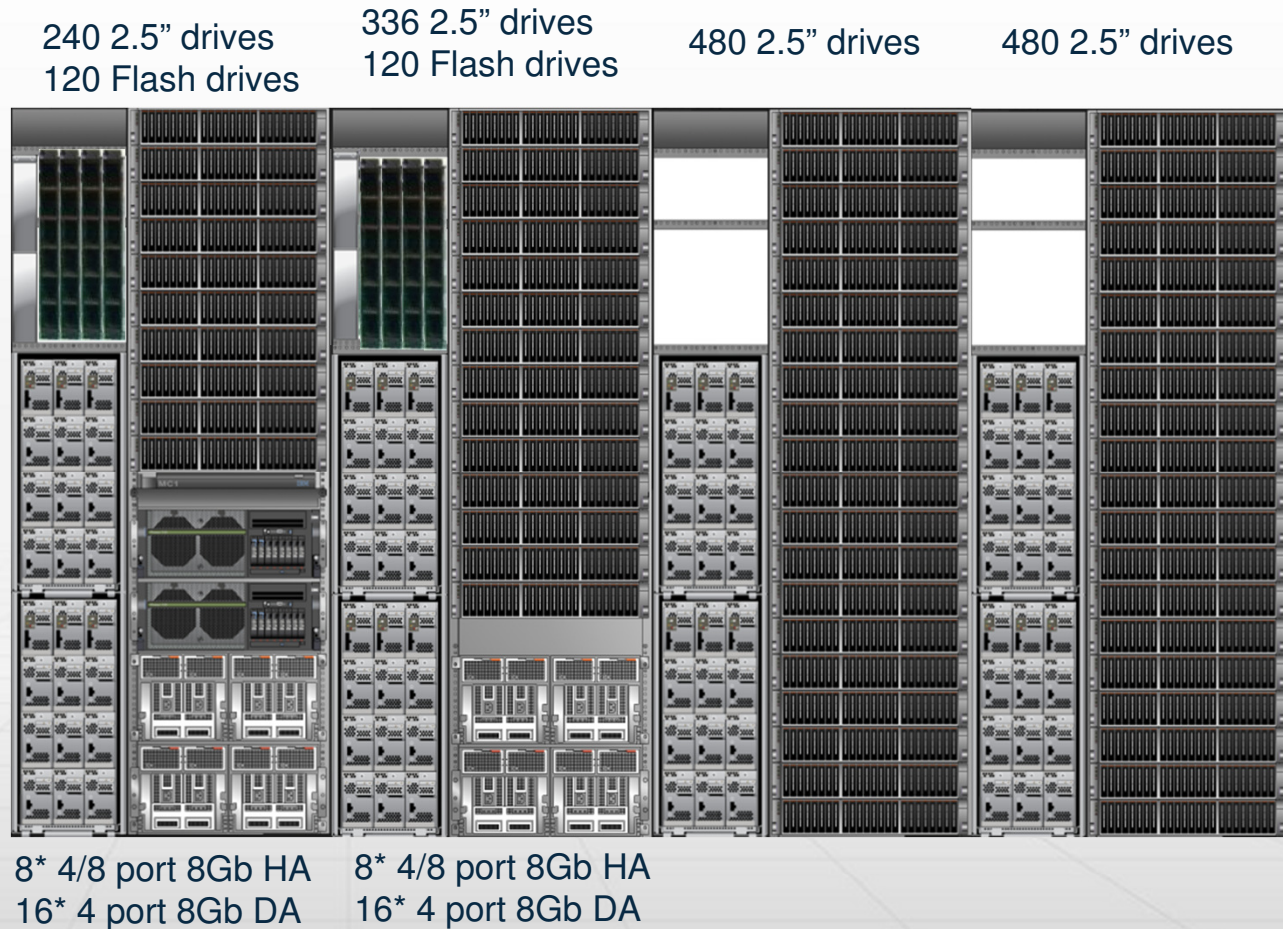
- Massive improvements in throughput for random workloads for High Performance Flash Enclosure compared to DS8870 device adapter
 - Up to 5.5x increase for random write
 - Up to 4x increase for random reads

Sequential IO




DS8870 Scalability

- Up to 4 frames with 1536 2.5" drives and 240 Flash drives
 - Third and fourth frames can be located separately by RPQ
- Up to 128 ports with 4/8 port 8Gb Host Adapters
 - Each port separately configurable as FICON or FC
- Up to 1TB Power7+ memory used for read and write cache
- Up to 65280 volumes



All-Flash DS8870 with HPFE

Up to 8 High Performance Flash Enclosures provide up to 96TB raw capacity with 400GB drives



All 8 IO bays installed in base frame for maximum throughput and host connectivity in single frame

8-core 256GB cache or 16-core 512GB cache or 16-core 1024GB cache

DS8870 Drive Technology

Performance

- Flash – 1.8” in High Performance Flash Enclosure
 - 400 GB drive
- SSD – 2.5” Small Form Factor
 - Latest generation with higher sequential bandwidth
 - 200/400/800GB/1600GB SSD
- 2.5” Enterprise Class 15K RPM
 - Drive selection traditionally used for OLTP
 - 146/300/600GB drives
- 2.5” Enterprise Class 10K RPM
 - Large capacity, much faster than Nearline
 - 600GB and 1.2TB drives
- 3.5” Nearline – 7200RPM Native SAS
 - Extremely high density, direct SAS interface
 - 4TB drives



DS8000 Replication Technologies


FlashCopy
Point in time copy

Within the same Storage System



Metro Mirror
Synchronous mirroring

Primary Site A Metro distance Site B




Global Mirror
z/OS Global Mirror
Asynchronous mirroring

Primary Site A Out of Region Site B



Metro Global Mirror
Metro z/OS Global Mirror
Three site and Four Site synchronous & asynchronous mirroring


Primary Site A Metro Site B Out of Region Site C/D



Complete your session evaluations online at www.SHARE.org/Seattle-Eval

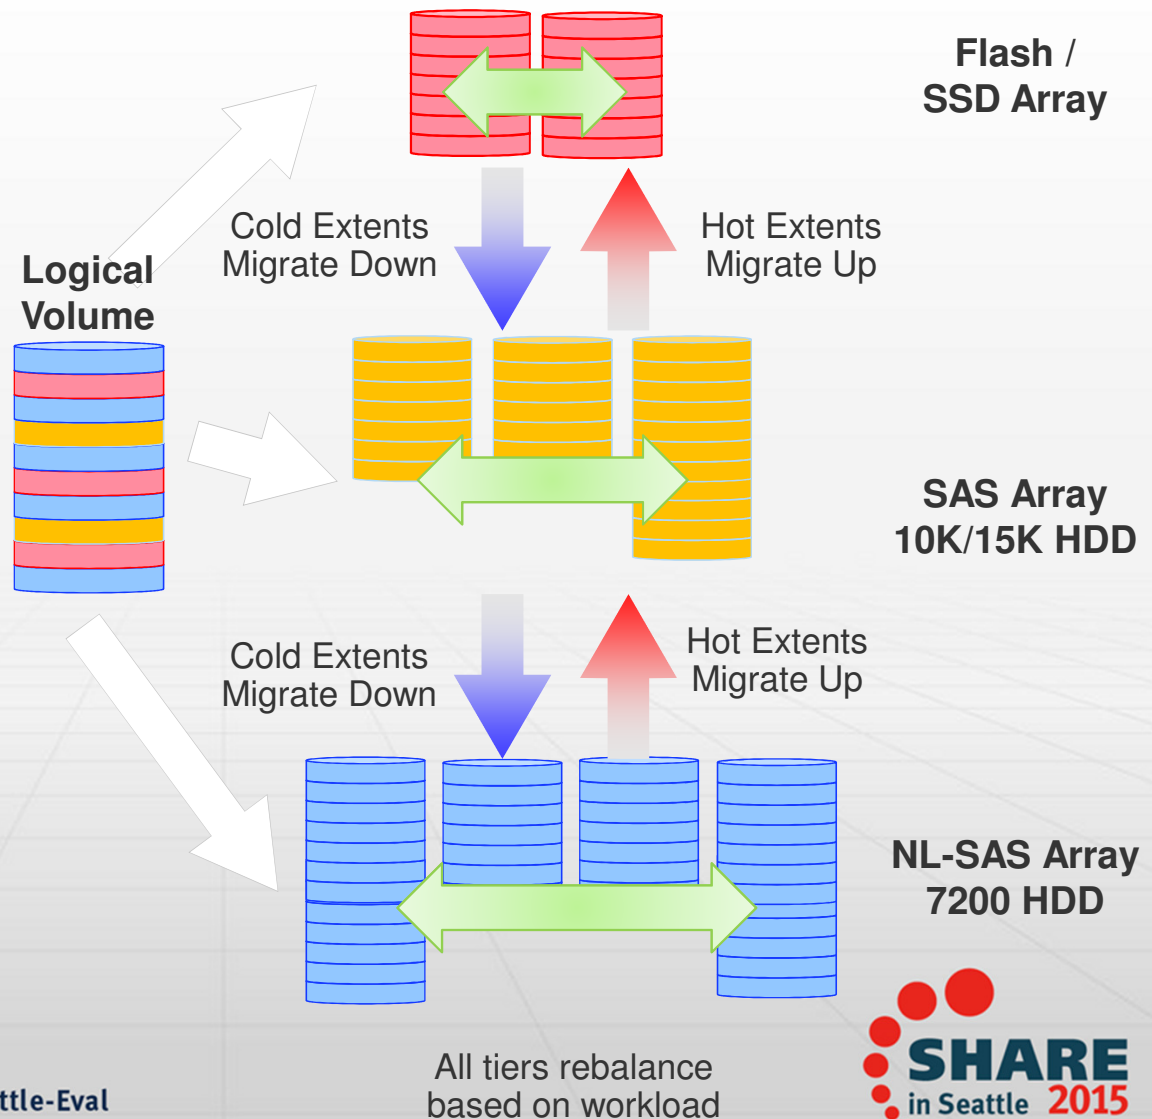
HyperSwap for Continuous Availability

- Substitutes storage secondary to take the place of failed primary device
 - Non-disruptive - applications keep running
 - Key value add to HA/DR deployments
- Unplanned HyperSwap
 - Continuous Availability in case of storage failures
- Planned HyperSwap
 - Storage Maintenance without downtime
 - Storage migration without downtime
- Typical swap times of a few seconds
- Supported on z/OS, zLinux, AIX and IBM I
 - Delivered by GDPS, TPC-R and PowerHA



Easy Tier automated tiering

- Optimisation of backend storage resources based on historical performance data
- SubLUN granularity using native DS8000 extents for any volume type
- Flexible configurations with any combination of drives of any size and speed
- Easy Tier Application provides APIs for policy and proactive actions
- Easy Tier HeatMap transfer enables workload history to be transferred for replication scenarios (DR, migration etc)
- Easy Tier Server integrates server side caching with Flash Adapter 90 and other devices




DS8870 Release 7.4 Content

☰ Actions 🔍 Filter

- **Key Hardware Enhancements to the DS8870**
 - High Performance Flash Enclosures in First Expansion Rack
 - Encrypting 1600GB 2.5" SSDs
 - Encrypting 600GB 2.5" 15K RPM HDDs
- **Key Function Enhancements to the DS8870**
 - MultiTarget Mirroring
 - Global Copy Collision Avoidance
 - Multiple Incremental Flash Copy
 - zHyperWrite : DB2 Log Write Acceleration with Metro Mirror
 - XRC Workload Based Write Pacing
 - Easy Tier Policy controls
 - zEasy Tier Application DFSMS and DB2 integration
 - Enhanced User Interface
 - Multi-thread Performance Accelerator



Multi-target Mirroring



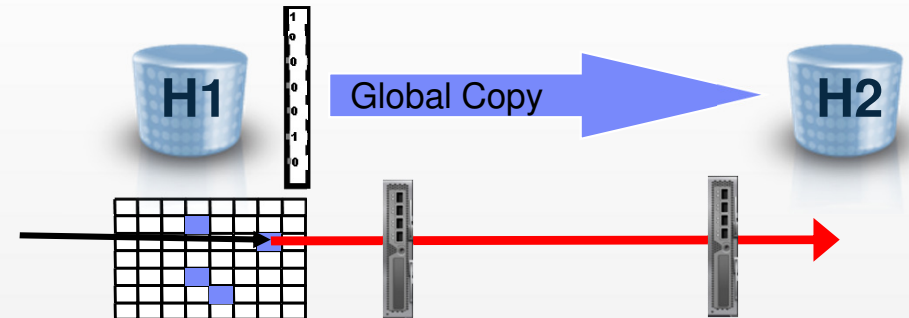
- Allow a single volumes to be the source for more than one PPRC relationship
- Provide incremental resynchronisation functionality between target devices
- Use cases include
 - Synchronous replication within a datacentre combined with another metro distance synchronous relationship
 - Add another synchronous replication for migration without interrupting existing replication
 - Allow multi-target Metro Global Mirror as well as cascading for greater flexibility and simplified operational scenarios
 - Combine with cascading relationships for 4-site topologies and migration scenarios
- TPC-R and GDPS support for Multi-target Metro Mirror

Global Copy Collision Avoidance

- A collision for Global Mirror or Global Copy is where there is a conflict between the requirement to send a track from primary to secondary and a subsequent update to that track
- A sidefile was introduced for Global Mirror collisions in R6.2 which avoided the requirement to delay writes to a track in the current consistency group until the track was sent to the secondary
- This enhancement introduces the same capability where the track is already in the process of being sent to the secondary when the update occurs and is for both Global Copy and Global Mirror
- The functionality is also available as an RPQ from R7.2 and R6.3


Prior to R7.4

Track in the process of being sent is locked to prevent writes from occurring.



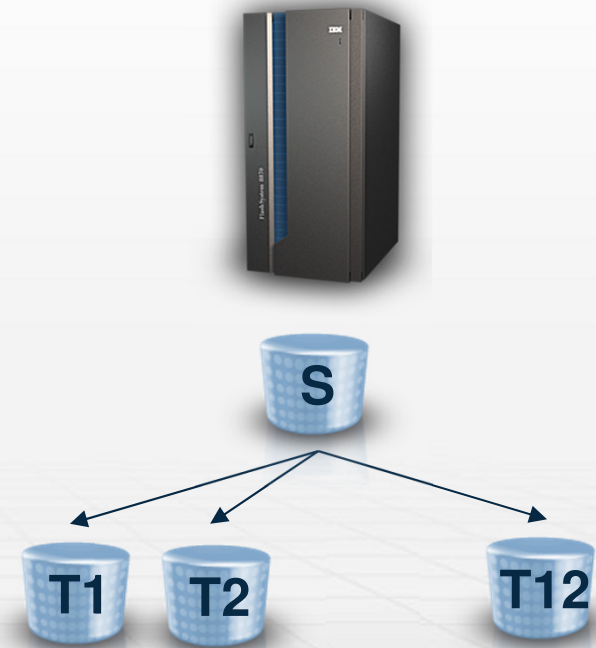
With R7.4

Track is locked until data is sent to host adapter and is then released. Collision sidefile is used when required for Global Mirror transfers



Multiple Incremental FlashCopies

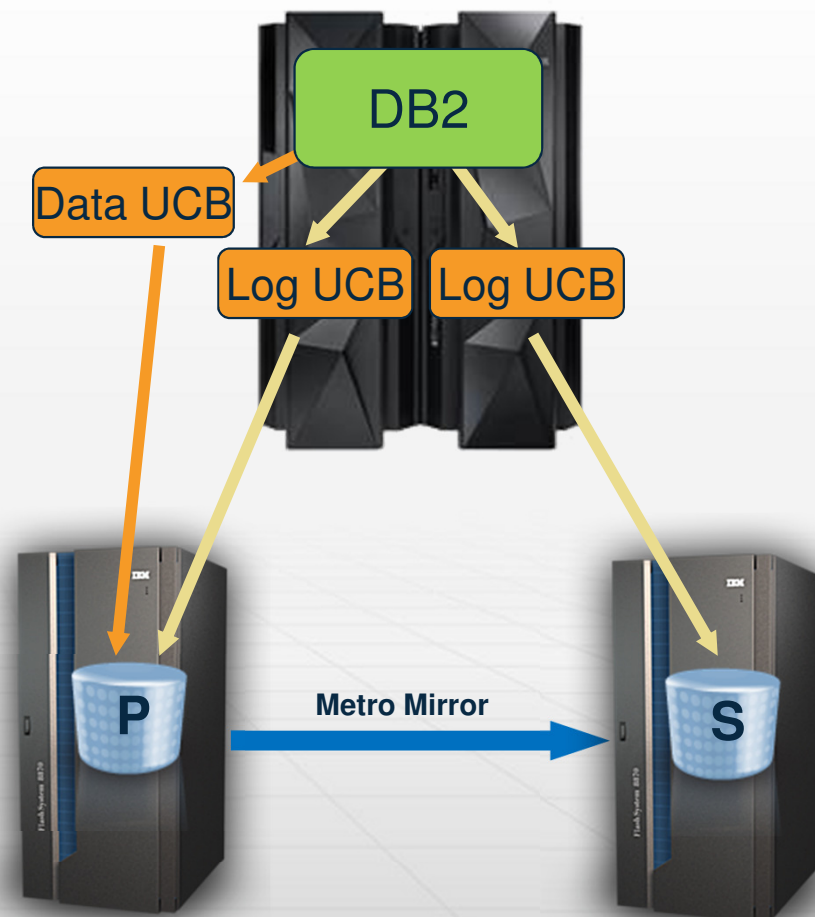
- Previously only a single incremental FlashCopy was allowed for any individual volume
- This provides the capability for up to 12 incremental FlashCopies for any volume
- A significant number of clients take two (or more) FlashCopies per day for database backup both of which can now be incremental
- The Global Mirror journal FlashCopy also counts as an incremental FlashCopy so the testing copy can now also be incremental
- The functionality is also available as an RPQ from R7.1.5



zHyperwrite

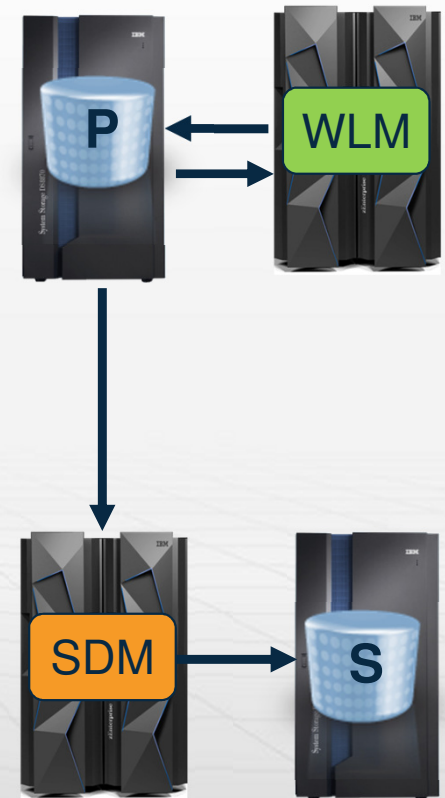
DS8870 Log Write Accelerator for DB2

- Improved DB2 Log Write Performance with DS8870 Metro Mirror
 - Reduces latency overhead compared to normal storage based synchronous mirroring
- Reduced write latency and improved log throughput
 - Up to 60% improvement in testing
- Supports HyperSwap with TPC-R or GDPS
- Enabled through
 - SYS1.PARMLIB(IECIOSxx)
 - SETIOS command
 - DS8000 microcode, IOS, DFSMS PTF's



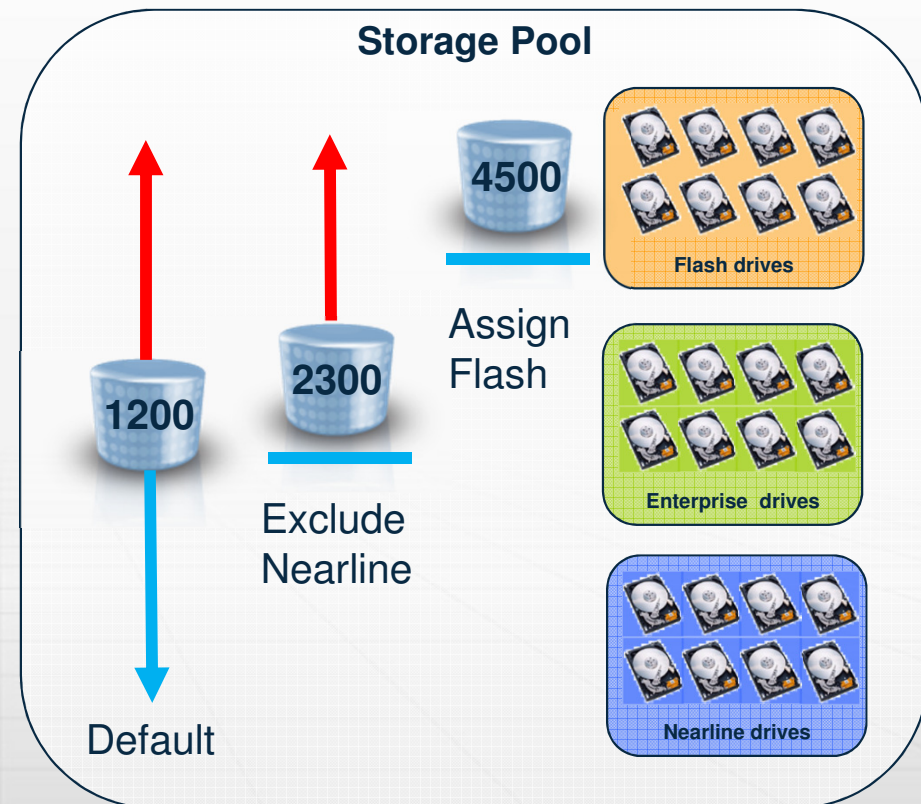
Workload Based Write Pacing for zGM

- Software Defined Storage enhancement to allow System Z Workload Manager to control XRC Write Pacing
- Reduces administrative overhead on hand managing XRC write pacing
- Avoids the need to define XRC write pacing on a per volume level allowing greater flexibility in configurations
- Prevents low priority work from interfering with the Recovery Point Objective of critical applications
- Enables consolidation of workloads onto larger capacity volumes



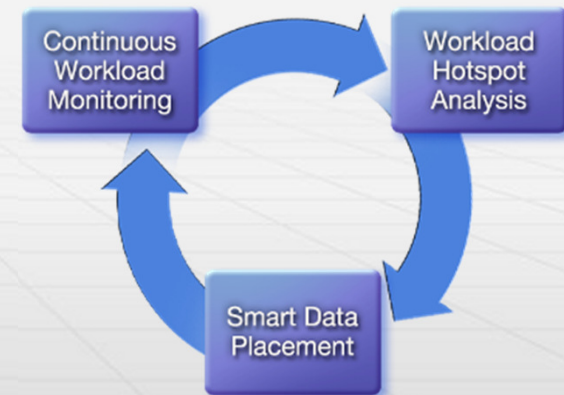
Easy Tier Application Policies

- New Exclude Nearline tier assignment policy
- Prevents the extents of a volume from being demoted to Nearline arrays
- If data is already on Nearline it will be promoted to Enterprise drives
- Three common use cases for Easy Tier Application policies
 - Default – optimise use of all tiers
 - Exclude Nearline – avoid potential low performance
 - Assign Flash – high performance guaranteed
- Also possible to assign to Enterprise or assign to Nearline but less common use cases



Easy Tier Controls

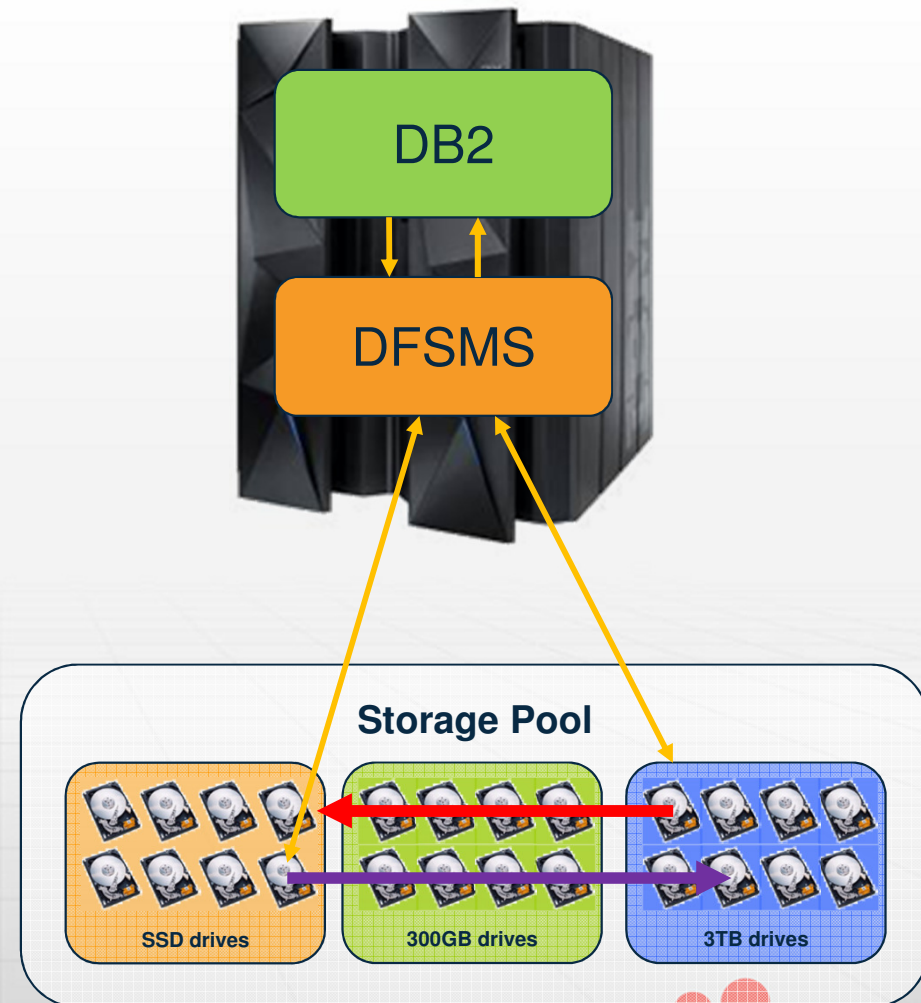
- In the majority of environments Easy Tier is able to successfully use the history of workload performance to predict the future requirements
 - There are however cases where this is not true
- Easy Tier Controls provide mechanisms for proactively and reactively modifying Easy Tier behaviour to handle these situations
- Controls include
 - Pause and Resume Easy Tier learning for volume or pool
 - Reset Easy Tier learning for volume or pool
 - Pause and Resume Easy Tier migration for a pool



Easy Tier Application

Integration with DFSMS and DB2

- Easy Tier currently optimises data placement and tiering based on workload history and this does not always reflect the future performance requirements of the data
- Easy Tier will provide interfaces to enable software such as DFSMS and DB2 to provide hints when data has been created, moved or deleted
- This will avoid performance degradation following maintenance activities such as database reorganisation
- DB2 integration function PTFs not yet available



DS8000 Unified User Interface

Next generation user interface providing unified interface and workflow for IBM storage products

7.4 release of functionality including

Logical configuration

System Health

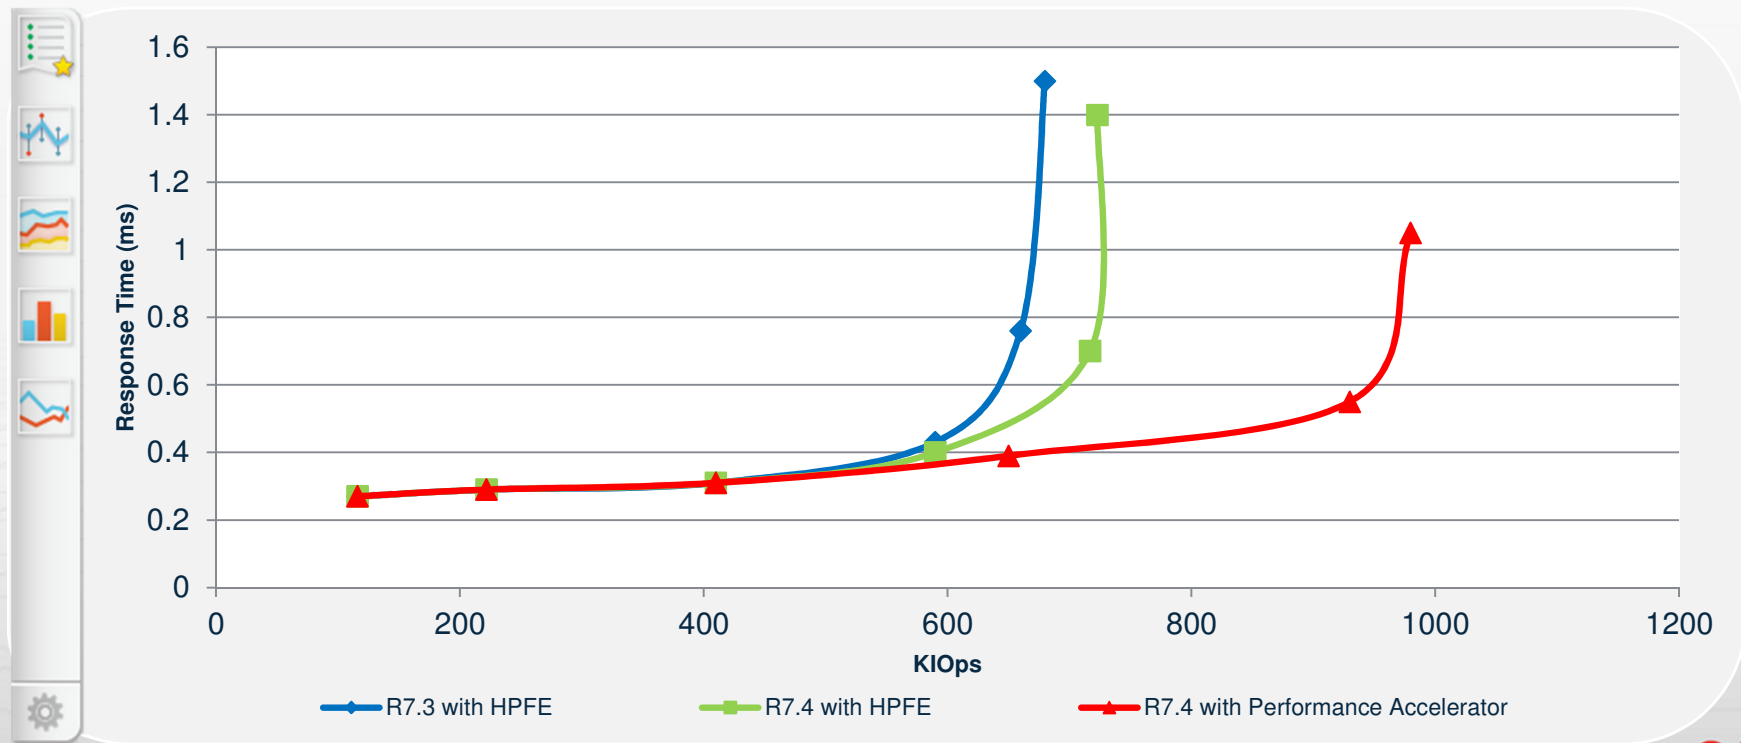
Events reporting

Designed in commonality with San Volume Controller, Storwize V7000, XIV, Tape, and SONAS products.




Multi-thread Performance Accelerator

- The IBM Multi-thread Performance Accelerator provides up to 45% performance improvement when used with the High Performance Flash Enclosures
- Up to 10% improvement will be seen without the Performance Accelerator

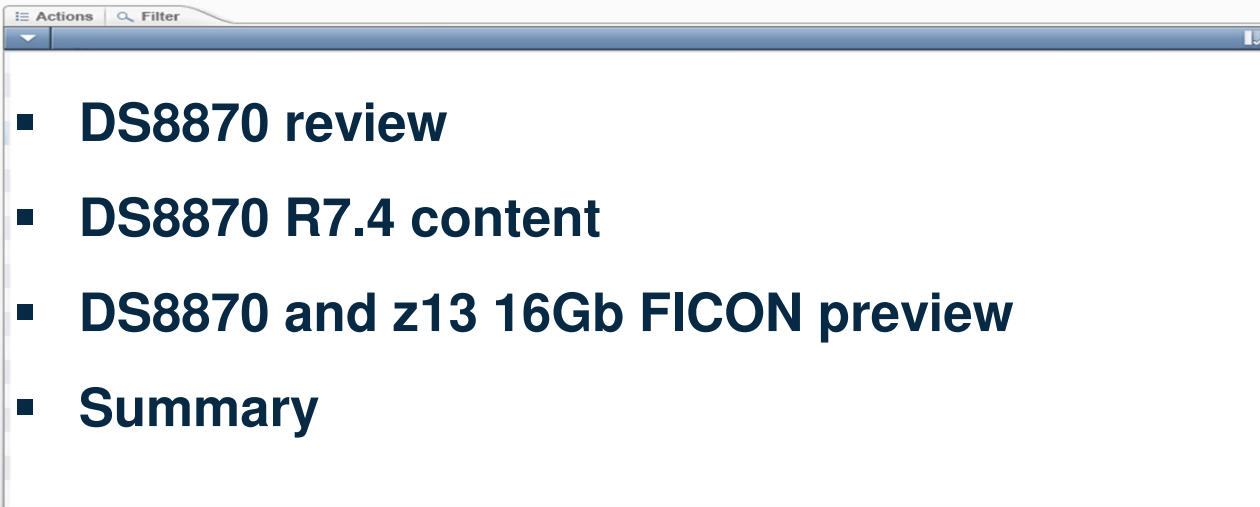


DS8870 and z Systems: Improve performance and resiliency for mainframe environments

January 2015 preview announcement

- 
- FICON Dynamic Routing** Reduce costs with improved and persistent performance for supporting I/O devices
 - 16Gb host adapters** Improve network performance with 2x faster FC and FICON adapters; minimize latency for DB2 log writes with zHyperWrite
 - Forward Error Correction** Preserve data integrity with more redundancy on the information transmitted via 16Gb adapters
 - zHPF Extended Distance II** Increase remote data speed with 50% better IO performance for remote mirror
 - Fabric Priority** Improved resiliency capabilities while enhancing the value of FICON Dynamic Routing

Agenda

- 
- A screenshot of a presentation slide with a blue header bar containing 'Actions' and 'Filter' icons. The slide content is a bulleted list of agenda items.
- **DS8870 review**
 - **DS8870 R7.4 content**
 - **DS8870 and z13 16Gb FICON preview**
 - **Summary**

