Topic: What's New with the DS8800 ?

Session # 12040

August, 2012

bobkern@us.ibm.com

SHARE
in Anaheim
2012

Agenda

- **DS8800 Overview**

- **Performance -> End to End Focus**
  - zHPF Improvements
  - z/OS DB Exploitation of zHPF + DS8K
  - DS8K Easy Tier
  - zWLM + DS8K I/O Priority Manager

- **Management & Growth**
  - Saving Device Addresses
  - Larger Volumes EAVs
  - HyperPavs
  - Alternate Subchannel Exploitation

- **Business Continuity -> Availability & Disaster Recovery**
  - DS8700/DS8800 Data Replication Technology
    - **Eliminating DB2 & IMS Backup Windows**
      - **zCDP for DB2 (DB2 System Level Backup & Restore)**
      - **zCDP for IMS**
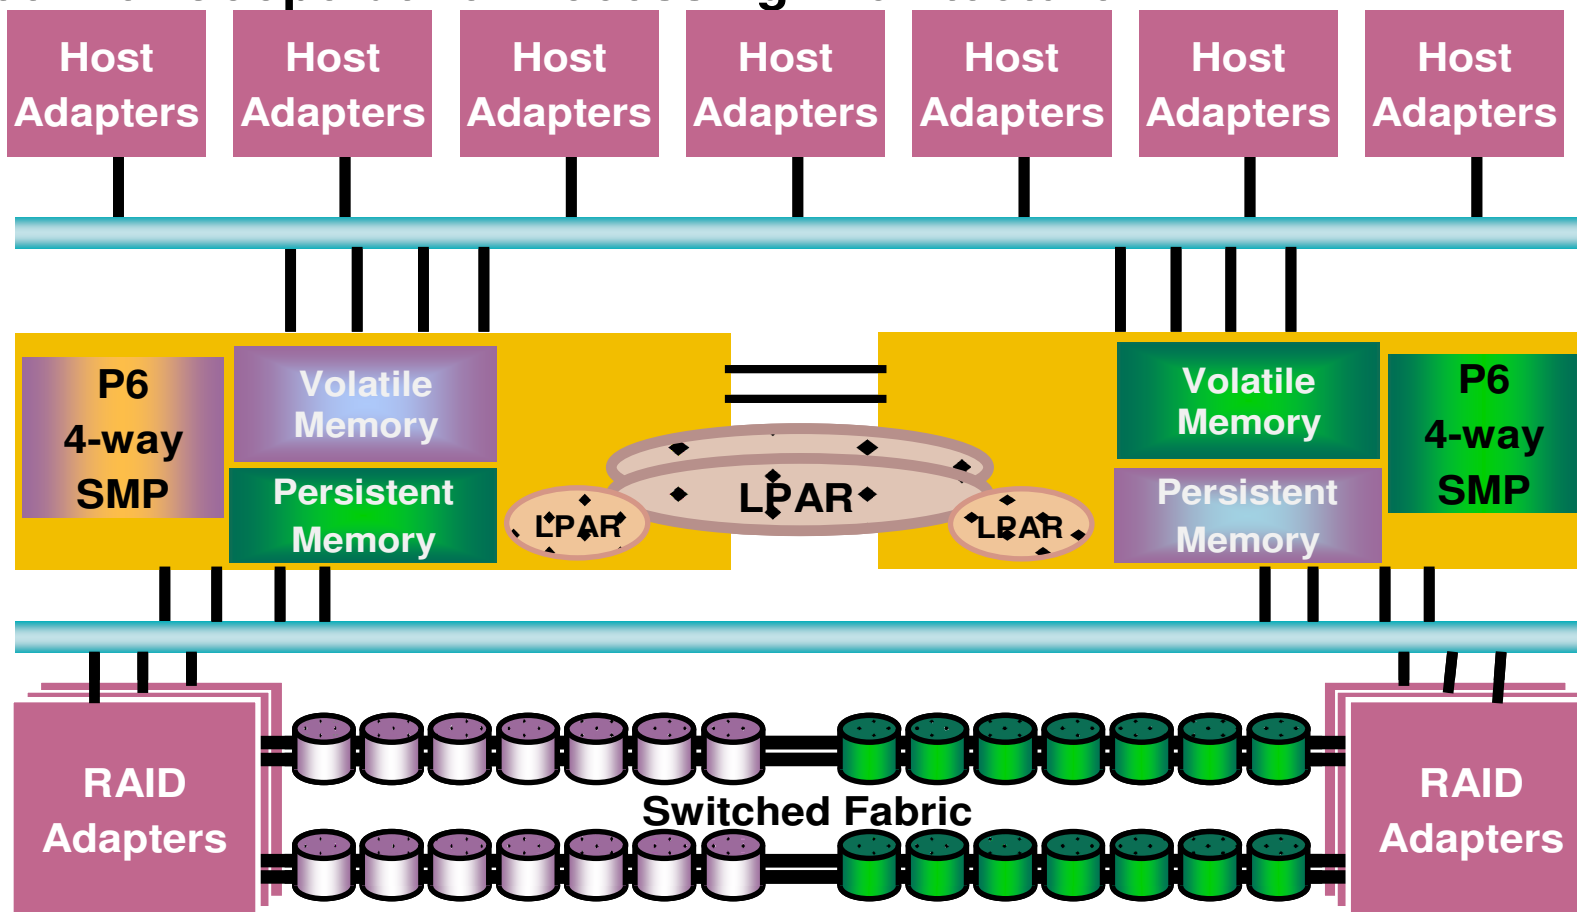 - GDPS High Availability & D/R Protection

- **Summary**

DS8000 - General Hardware Structure –

Three-Tier Cooperative Processing Architecture

IBM DS8800 Disk Storage Subsystem

Up to over 40% increase in performance and almost 90% more drives in same single-frame footprint

System-wide hardware upgrade and IBM POWER Server synergy

- ✓ **New POWER6+ controller**
- ✓ **New 8Gb/s host and device adapters**
- ✓ **New 2.5" 6Gb/s SAS (SAS-2) drives and drive enclosures**

Features / Business Value

- ✓ **New processors and adapters boost performance with faster access to host servers and disk drives**
- ✓ **New small-form-factor drives offer faster performance, higher availability, lower energy consumption, and allow denser footprint for more effective scalability**
- ✓ **Easy Tier – Dynamic Storage Hierarchy – SDD/FC/SATA**
- ✓ **IBM Storage Common GUI**
- ✓ **Synergy with System z**
- ✓ **Storage Based Data Replication Functionality**
- ✓ **GDPS High Availability & Disaster Recovery**

Client Benefits

- ✓ **System-wide upgrades enable faster performance and higher storage capacity within the same footprint**
- ✓ **Smaller form factor drives offer faster performance, higher availability, smaller footprint, and lower energy consumption**

DS8800	DS8800
DDMs	16-1056
DDM Interface	6Gbps SAS-2
Enterprise (FC/SAS) DDM Types	SAS - 146, 450, 600 GB
SATA DDM Types	-
SSD DDM Types	300 GB
RAID Types	RAID 5, 6, 10
Max Usable Capacity	471 TiB
Max Sequential Bandwidth (MB/s)	11.8GB/s
Max Number of LUNs / CKD volumes	64K total
Max N-Port Logins/Port	510
Max Process Logins	2K
Max Logical Paths / CU	512
Max LUN Size	16 TB
Dynamic Provisioning	Add / Del / Depopulate rank
Cache / NVS	32-384GB / 1-12GB
Processor	P6+ 5.0GHz 2 or 4-way
Host Adapters	8 Gb FC x 4 or 8 ports per adapter
Host Adapter Slots	16
Max Host Adapter Ports	128
Single DA Throughput	1,600MB/s+
DA Slots	16

SHARE in Anaheim
2012

System Storage and System z A unique synergy

- System z and System Storage Synergy Advantages:
  - ► Collaborate on architecture and design
  - ► End to End Focus – Place Function in correct place in HW/SW Stack !
    - ► System z Hardware/Channel/DS8800
    - ► DB2/z/OS DFSMS, IOS/zHPF/DS8800
    - ► zWLM can manage workload end to end
    - ► Performance, Availability, Management & Growth
  - ► Conduct early, rigorous and comprehensive stress testing In System z labs for every enterprise storage release. Tests are designed to push the limits of functionality and robustness.
  - ► Share skilled support resources with enterprise class experience and expertise

- This helps System Storage and System z development to:
  - ► Better design products that work well together with more robust interlocking
  - ► Bypass potential pitfalls with fixes implemented before customers are ever impacted
  - ► Implement streamlined, efficient, integrated solution offerings
  - ► Ensure that products that have passed stringent testing together

System z DS8800 Synergy Items
FICON Express2 MIDAW
Support for 64K cylinder 3390s (about 19 3390-3s)
AMP: Adaptive Multi-stream Prefetching in a Shared Cache *
Maximum concurrent FICON requests per FICON port
HyperPAV
C.U.I.R. (control unit initiated reconfiguration))
I/O Priority Queuing
Full FlashCopy compatibility
GDPS /Global Mirror
XRC (a.k.a. z/OS Global Mirror)
z/OS Global Mirror (XRC) write pacing
z/OS Global Mirror (XRC) suspend rather than long busy
Dynamic volume expansion *
z/OS Global Mirror (XRC) multiple reader
z/OS Global Mirror (XRC) ext. distance FICON
z/OS M/GM HyperSwap with incremental resync
z/OS Basic HyperSwap
3390s larger than 64K cylinders

SHARE
in Anaheim
2012

Performance

Link Protocol Comparison for a 4KB READ

FICON

zHPF

OPEN EXCHANGE,
PREFIX CMD & DATA

READ COMMAND

CMR

4K of DATA

STATUS

OPEN EXCHANGE, send a Transport Command IU

4K OF DATA

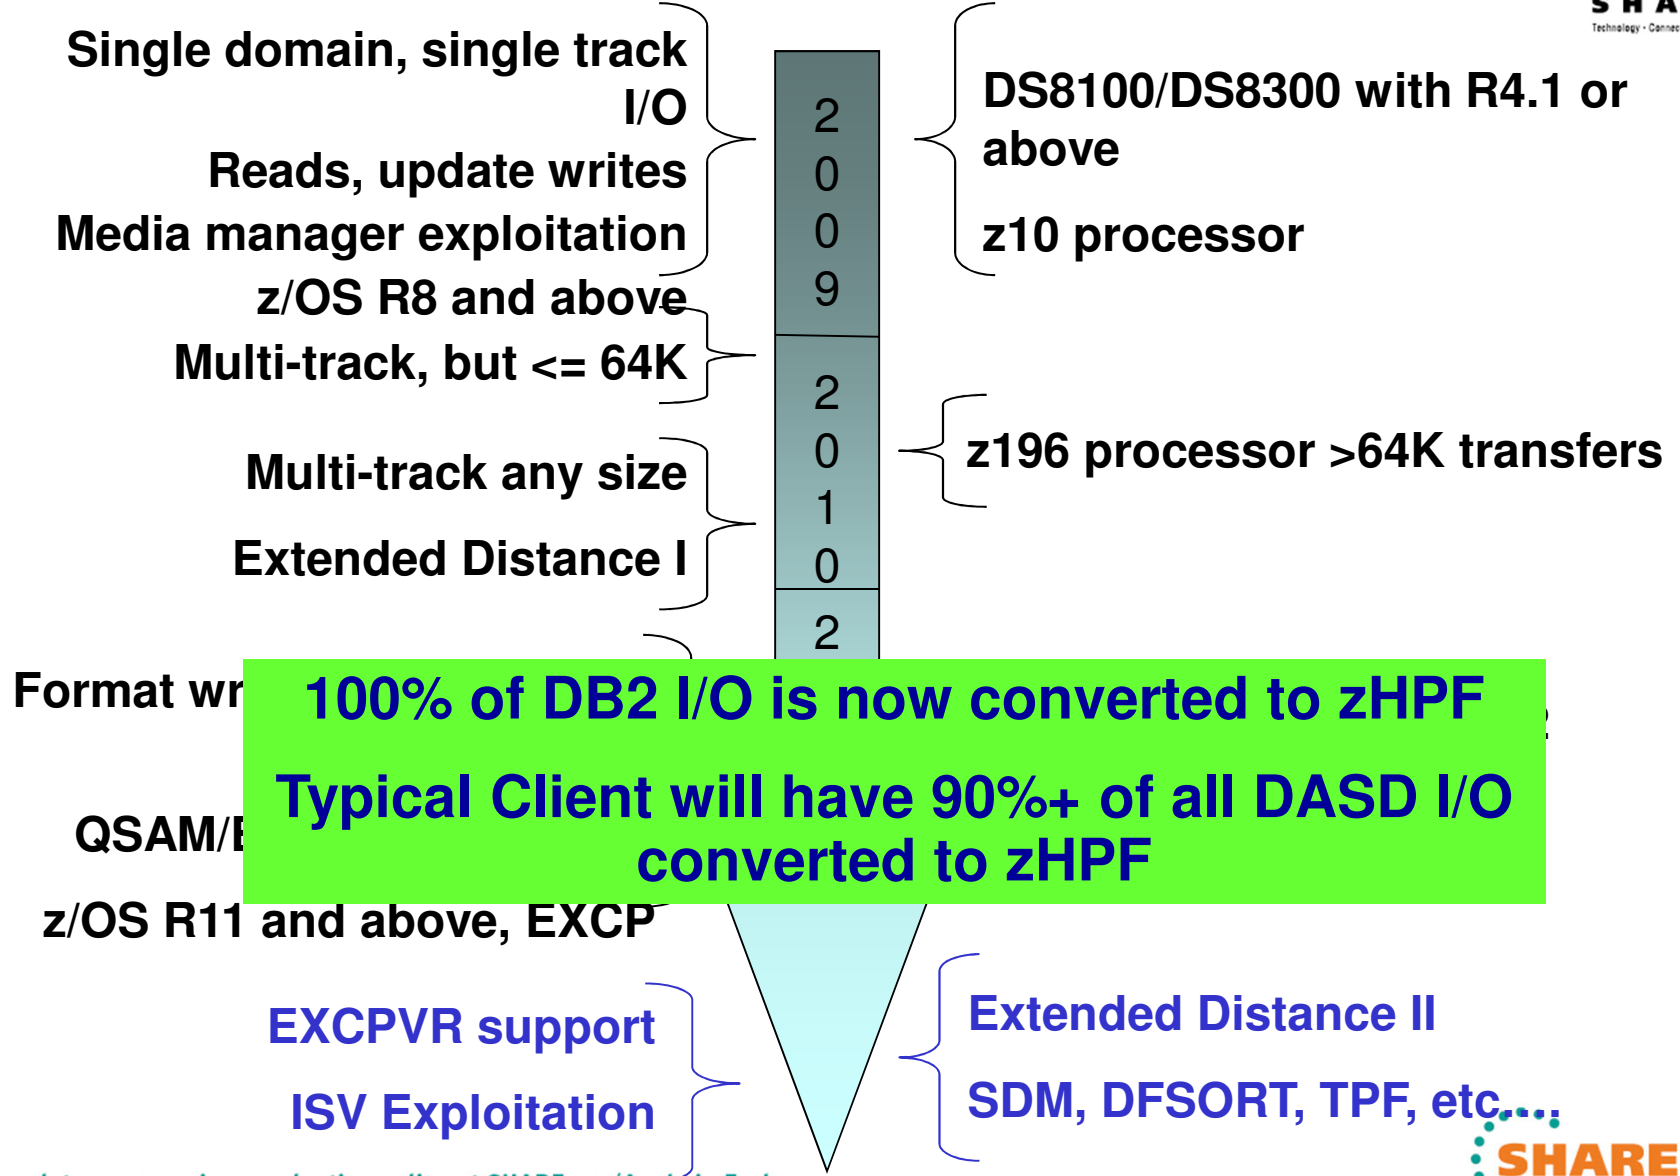**Send Transport Response IU CLOSE EXCHANGE**
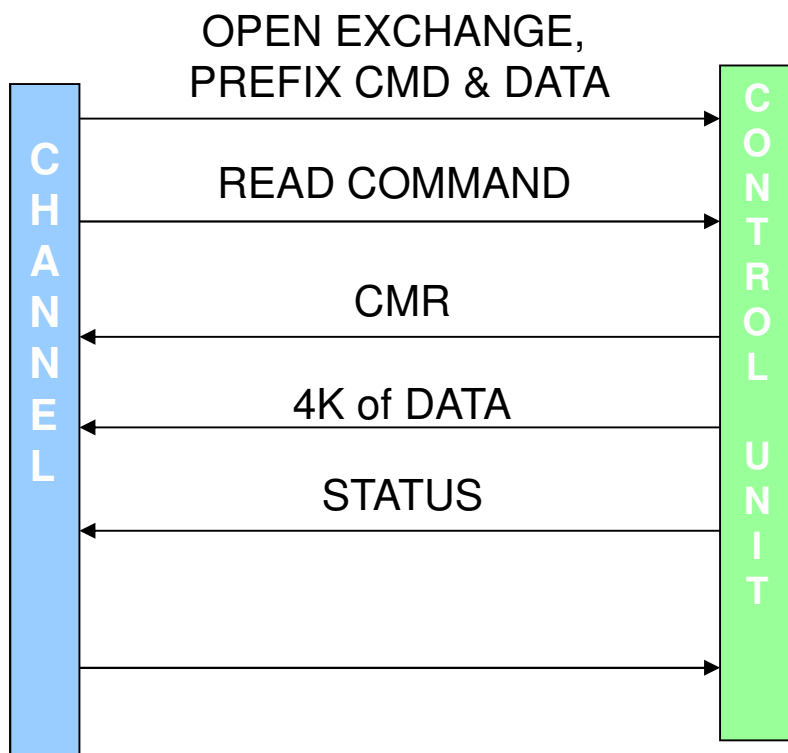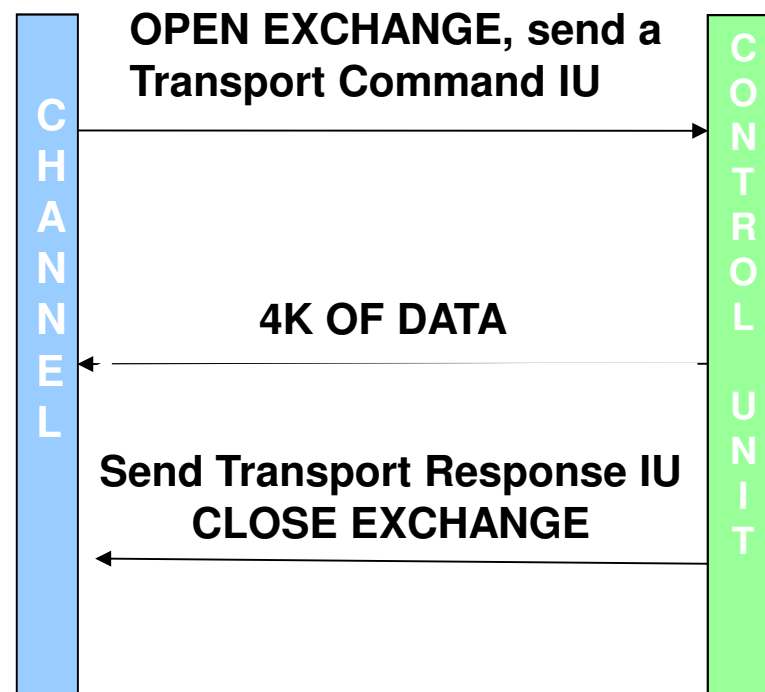
IU – Information Unit

zHPF provides a more streamlined link protocol than FICON

Complete your sessions evaluation online at SHARE.org/AnaheimEval

zHPF Evolution

Single domain, single track I/O

Reads, update writes

Media manager exploitation

z/OS R8 and above

Multi-track, but <= 64K

Multi-track any size

Extended Distance I

2009

2010

2

DS8100/DS8300 with R4.1 or above

z10 processor

z196 processor >64K transfers

100% of DB2 I/O is now converted to zHPF

Typical Client will have 90%+ of all DASD I/O converted to zHPF

Format wr

QSAM/E

z/OS R11 and above, EXCP

EXCPVR support

ISV Exploitation

Extended Distance II

SDM, DFSORT, TPF, etc....

SHARE
in Anaheim
2012

z/OS Storage Leadership with DS8700/DS8800 R6.2
- *High Performance FICON (zHPF) extensions for optimal system performance*

- **Shorter elapsed times for I/O intensive batch reduces batch window**

  - **New DS8000 I/O commands optimize QSAM, BPAM, and BSAM access methods for exploiting zHPF**
    - **Benefits non-extended format sequential and basic and large format sequential data sets**
  - **I/O service times are improved by up to 30%**

- **Complete conversion of DB2 I/O to zHPF maximizes resource utilization and performance -> Significant Throughput gains in many areas.**

  - **Format write throughput increases up to 52% (4K pages)**
    - **Applies to Load, Reorg writes to shadow, Restores**
    - **DB2 V8 increase only 30%, DB2 9/10 increase up to 52%.**
  - **Pre-formatting throughput increases up to 100%**
  - **Sequential pre-fetch throughput increases up to 19%**
  - **Dynamic pre-fetch throughput increases up to 23% (40% with SSD)**
  - **Disorganized index scans yield throughput increases up to 111% on DB2 10 (more with 8K pages)**
    - **DB2 V9 throughput increases up to 43%**
    - **Together, DB2 10 and zHPF is up to 11 times faster**
    - **Synchronous I/O cache hit response time decreases by up to 30%**
  - **Improvements in cache handling decrease response times**
    - **Skip sequential index-to-data access improves cache miss processing by 3x to 4x**
  - **Up to 50% reduction in the number of I/O operations for query and utility functions**
    - **New DS8700/DS8800 R6.2 Algorithm to handle DB2 List-Prefetch I/O feature with the z114, z196 GA2.**

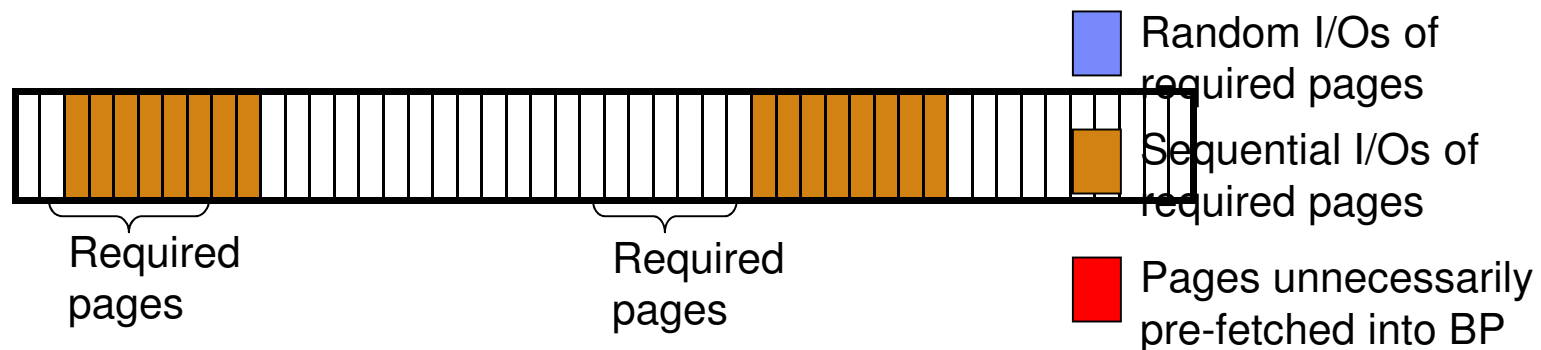Complete your sessions evaluation online at SHARE.org/AnaheimEval

SHARE in Anaheim
2012

Skip Sequential – Dynamic Pre-fetch

- **Data may be accessed in a skip sequential manner**

  - Gaps between required sequential pages
    - Dynamic pre-fetch learns into pre-fetch after 5 sequential pages
      - Triggering next pre-fetch quantity

  - What if application does not need pre-fetched pages?
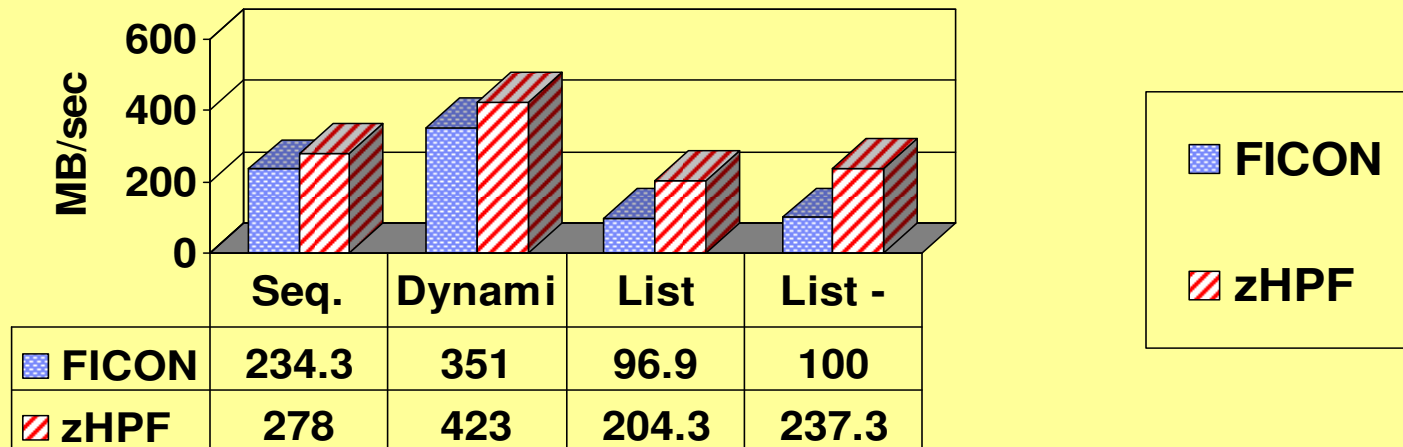
5 seq pages to learn into pre-fetch

Pre-fetched pages

5 seq pages to learn into pre-fetch

Pre-fetched pages

Random I/Os of required pages

Sequential I/Os of required pages

Pages unnecessarily pre-fetched into BP

Required pages
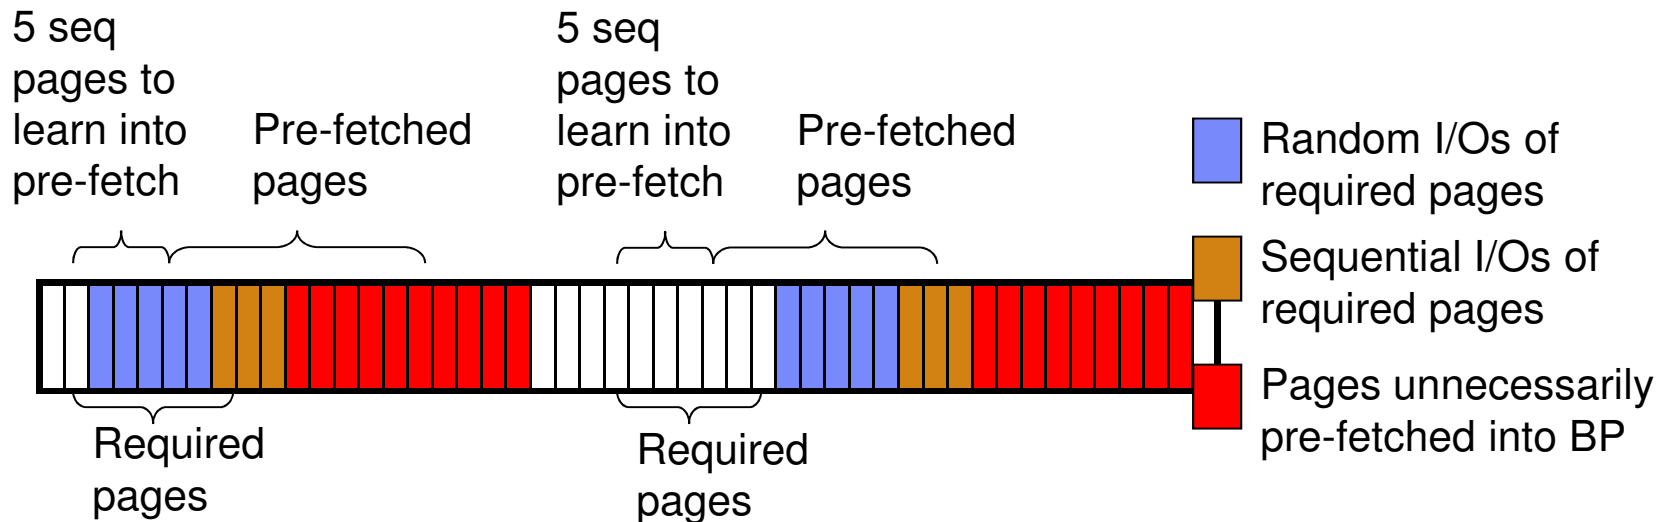
Required pages

Skip Sequential – DB2 List Pre-fetch

- **List-prefetch does not need to learn into sequential access**
  - **Only those pages required are read into the buffer pool**
- **zHPF List-Prefetch Support can yield 20x I/O Service Time Reduction compared to cache miss**
- **z/OS DB2 keeps 2 Read I/O Chains Outstanding at a time each reading 32 DB2 Pages -> 64 DB2 Pages read at any PiT by DS8800.**

Required pages

Required pages

Random I/Os of required pages

Sequential I/Os of required pages

Pages unnecessarily pre-fetched into BP

DB2 Prefetch

- **DB2 uses "sequential prefetch" for table scans**

- **DB2 uses "dynamic prefetch" for…**

  - **Index scan if the index is organized**

  - **Index-to-data access if the data pages are sequential**

- **DB2 uses list prefetch for…**

  - **Index scan if the index is disorganized (DB2 10 only)**

  - **Log apply (e.g. Online REORG and RECOVER)**

  - **Incremental COPY**

  - **Reading a fragmented LOB**

  - **RUNSTATS table sampling (DB2 10 only feature)**

Read 4K pages from cache
FICON Express 8 & 8S

	Seq.	Dynami	List	List -
FICON	234.3	351	96.9	100
zHPF	278	423	204.3	237.3

Legend: FICON, zHPF

- ❏ **zHPF increases sequential prefetch by 19%**

- ❏ **zHPF increases dynamic prefetch by 23%**

- ❏ **With FICON Express 8, zHPF increases list prefetch by 111%**
 - ▪ **FICON Express 8S adds another 16% (145% more than FEx8 FICON)**

- ❏ **With 8S, dynamic prefetch still has 78% more throughput than list prefetch**

SHARE in Anaheim
2012

List prefetch 32x4K pages from cache FICON Express 8 & 8S
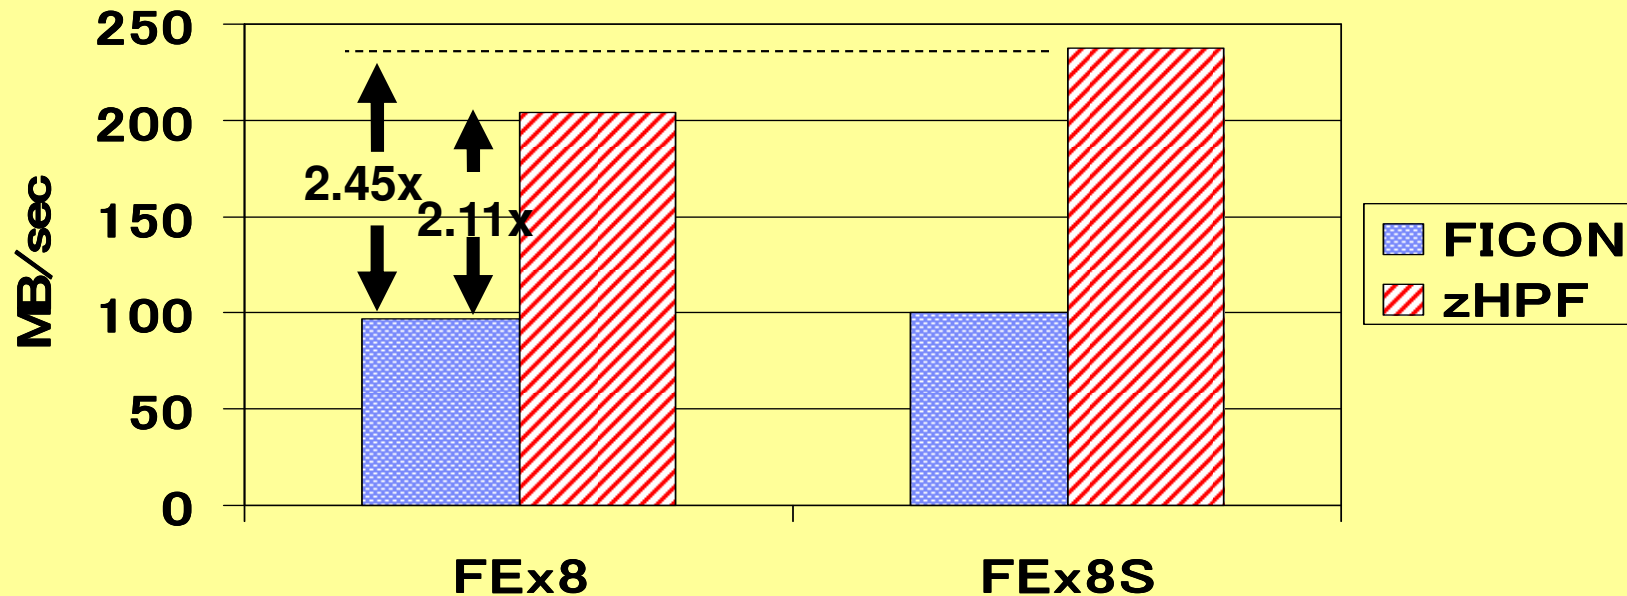
- With FICON Express 8, zHPF increases list prefetch by 111%

- FICON Express 8S adds another 16% (145% more than FEx8 FICON)

DS8800 R6.3 zHPF Enhancements

- ✓ **QSAM, BPAM, BSAM & EXCPVR access methods**

- ✓ **All Format Writes across all access methods now zHPF eligible**

- ✓ **Ficon Express 8S -> DB2 List PreFetch (QUERY & Utility Functions) – Further Performance improvements.**

- ✓ **SETSMS SAM_USE_HPF(YES/NO) – Turn function ON/Off.**

DS8700/DS8800 – I/O Priority Manager

Application Level Quality of Service

**- Managed by zWLM based on
   Service Class (QofS) specified.**

– **Dynamic workload control**

– **Host Management & DS8800. (I/O
   Priority & QofS Management
   controls.)**

– **Automatic performance
   adjustment to customer needs**

– **Customer needs: Applications
   (Volume/LUN groups) & workload
   priorities**

STORAGE CLIENTS

DS8K STORAGE SERVERS

2012
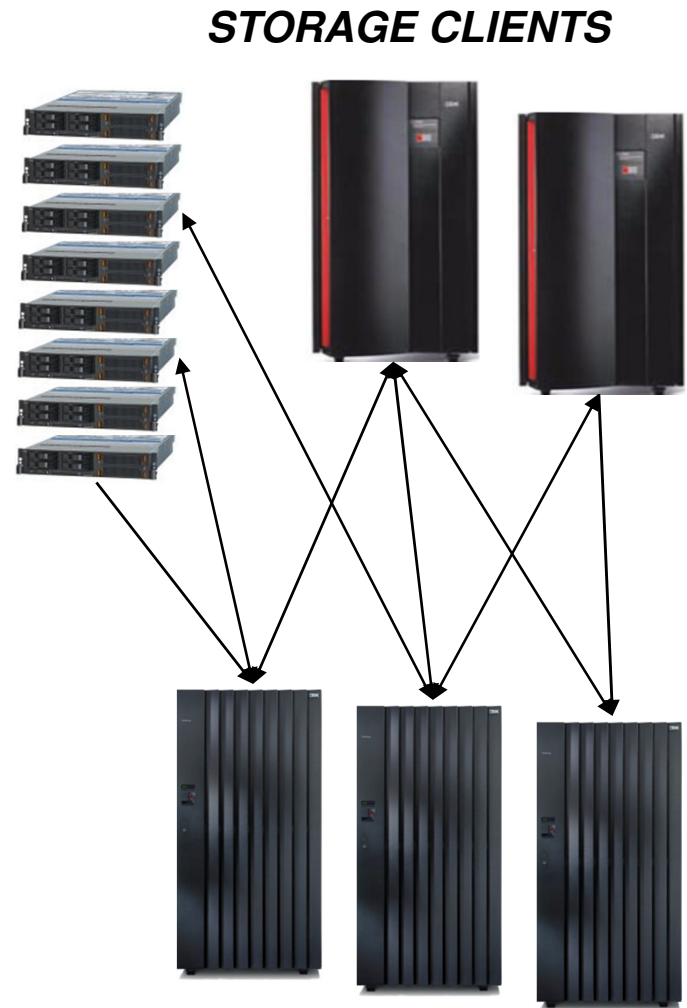
DS8K Easy Tier

- **Function** – **DS8K Automatic Lower Interface management. Monitor disk 'Hotspots' & dynamically moves extents to optimize performance.**

- **More Efficiency (R6.2)**
  - **Hybrid storage pools can be created between any two types of storage (SSD, FC, SATA)**
    - **Mix of FC and SATA for a cost optimized "tier 1.5" system**

- **More Flexibility**
  - **Automatic data relocation between all 3 storage types (Release 6.2)**
    - **For customers with sparsely accessed data sets**
    - **Coldest data migrates to SATA tier**
  - **Manual mode depopulation of an array or of a tier**

- **More Performance**
  - **Within a hybrid volume, data is arranged across multiple arrays for optimum performance**
  - **Release 6.2: Automatic performance rebalance within one tier**
    - **For customers with a single storage type (i.e. Enterprise disk), RAID hotspots are avoided by balancing hot extents across multiple arrays**

Flexibility

Performance

Efficiency

SHARE
in Anaheim
2012

IBM Self-Optimizing Storage with Easy Tier

Easy Tier balances performance and cost automatically

Performance

Cost

- Automatic movement of data to the right disk tier to balance cost and performance

- Continuous rebalancing **within** each tier to maintain peak performance across all drives

- Maximum benefit when Easy Tier extends beyond the disk system

IBM Easy Tier optimizes performance and costs across tiers

- Storage analyzes data I/O patterns and adapts to improve performance

- Makes efficient use of flash technology

- Analytics-based approach to identify data that can be hosted on a lower tier
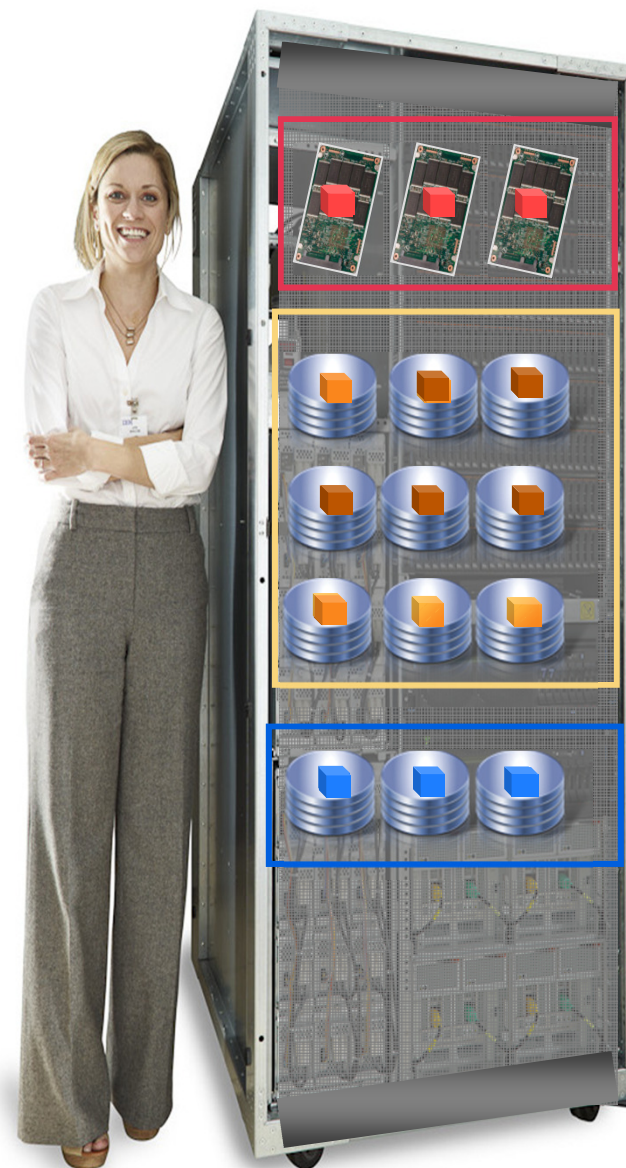
SSD optimization boosts application performance by 4x

Financial Brokerage Application
(Easy Tier with SSD + HDD)
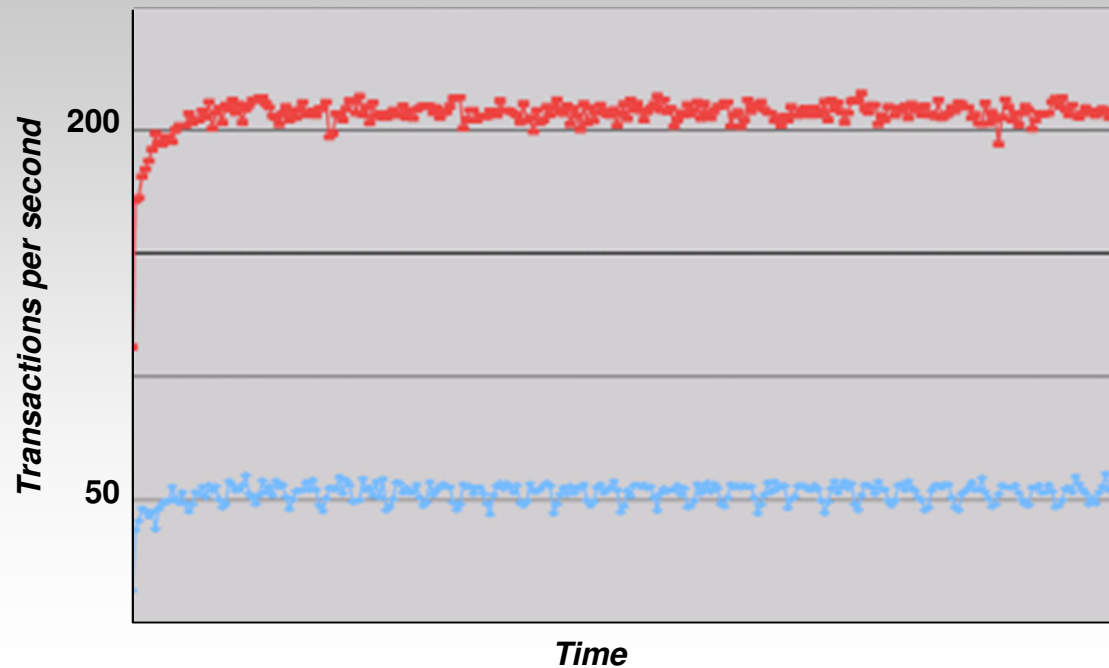
4x

throughput increase with only

10%

of a database onto SSD

Easy Tier also optimizes within each tier

- Continuous
  rebalancing keeps
  performance
  optimized
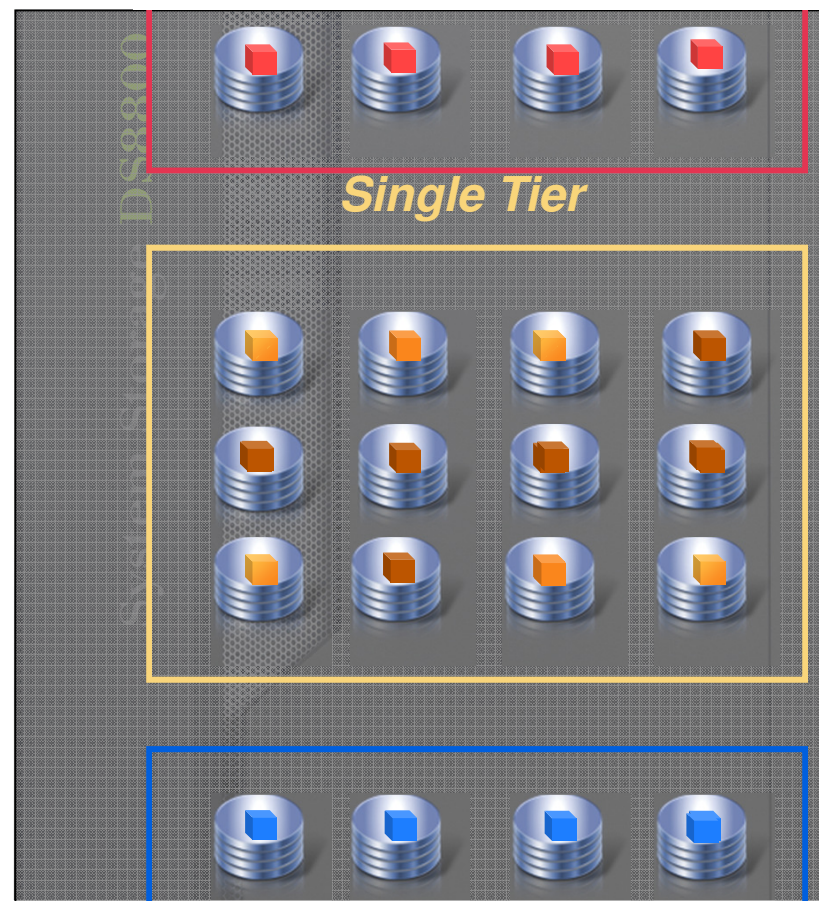
3x performance boost for single-tier optimization And better response times

OLTP

3X Improvement

Response Time (ms) vs Throughput (IOPs)

— 80 15K RPM drives (RAID-5), initially skewed
— 80 15K RPM drives (RAID-5), after auto-rebalance
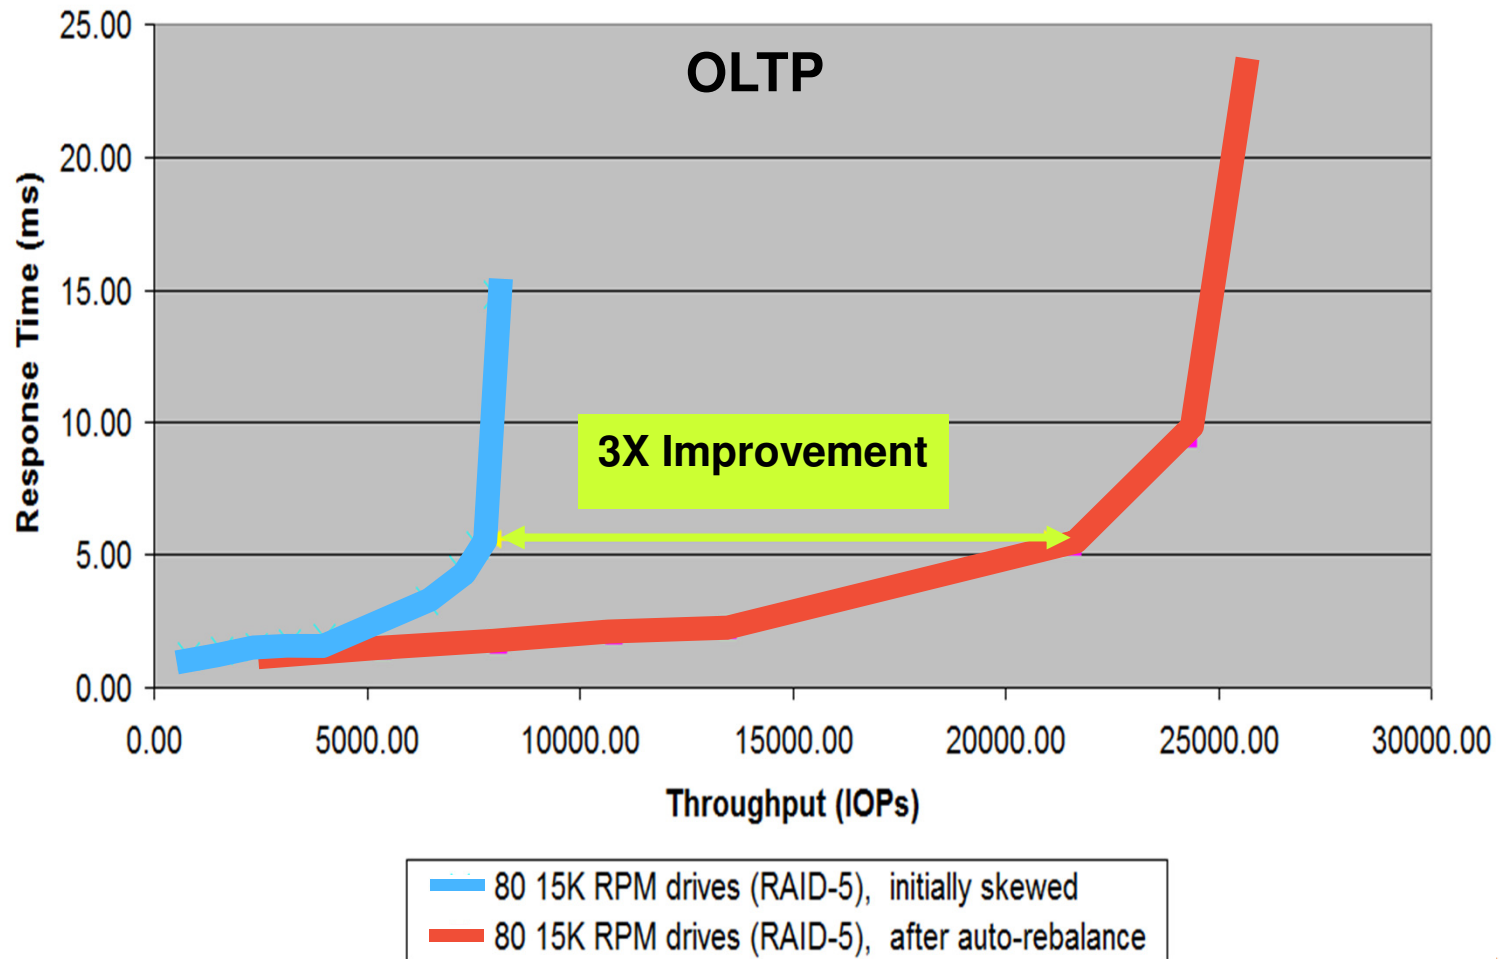
DS8000 Easy Tier June, 2012 SODs

1. DS8000 designed to leverage direct attached solid state storage on selected AIX and Linux servers. Easy Tier will Manage the solid state storage as a large and low latency cache for the hottest data, while preserving advanced disk functions such as RAID & remote mirroring.

2. API for Middle-ware and ISV programs to provide hints n data usage to Easy Tier.

3. A new high density flash storage module that will accelerate performance to another level with cost-effective, high density, SSD drives
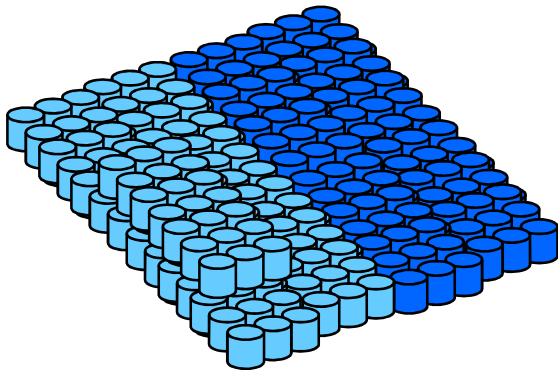
Management & Growth

Lets Review - "I am Running Out of UCBs"
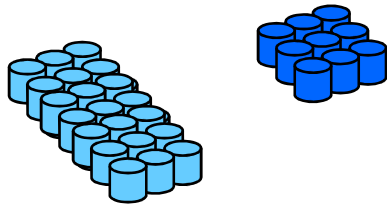
*Subchannel Set 1*

Subchannel Set 0

- **Alternate Subchannel Set Exploitation**
  - **50% reduction in device numbers**
- **HyperPAV**
  - **10x reduction in PAV devices by assign PAV aliases on demand as the work load requires**
  - **PAV-aliases virtualized across operating system images for more efficient use on addressing constructs**
- **Extended Addressability Volumes**
  - **223 GB Volume & Now 1TB volume Sizes**
  - **Reduce system resources and overhead with managing fewer resources**
- **Dynamic Volume Expansion**
  - **Without copy services intact**
- **MIDAWs**
  - **Performance enhancement for using larger datasets**
- **TDMF and zOSDMC (aka LDMF)**
  - **Non-disruptively consolidate data on to a single larger volume and consolidate device number ranges**
- **PPRC Secondary's in Alt Subchannel Set w/HyperSwap**

SHARE
in Anaheim
2012

DS8K R6.2 - Extended Addressability Volumes (EAV)
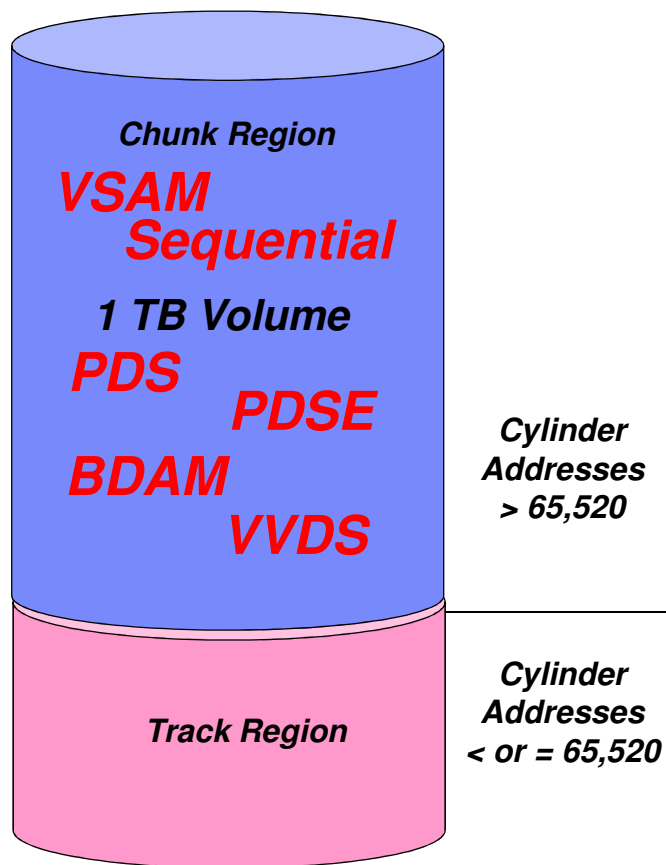
- **DS8700/DS8800 R6.2 (RQ 6785)   R6.1.5**
  - Support up to Mod 1062 = 1,182,006 Cylinders = 1.004+ TB (10^{12})
  - Prior limit was Mod 236
  - Configurable in increments of
    - Mod1 (1113 Cylinders) on volumes > 65520 cylinders (Mod 58+)
    - One cylinder on volumes with =< 65520 cylinders
  - Includes FlashCopy, PPRC, and XRC support
  - Allows 20x growth in data without the complexity in defining and managing additional devices
  - Provides for significant device number consolidation
  - Supported on z/OS V1R12, V1R13 and above
- **Continue Exploitation (z/OS 1.11 – z/OS 1.13)**
  - Non-VSAM Extended Format Datasets
  - Sequential Datasets
  - PDS
  - PDSE
  - BDAM
  - BCS/VVDS
  - 1 TB Volume Sizes
- **Dynamic Volume Expansion**
  - Not permitted with Copy Services
  - Automatic VTOC and Index Rebuild

Chunk Region

VSAM
Sequential

1 TB Volume

PDS
PDSE
BDAM
VVDS

Cylinder Addresses > 65,520

Track Region

Cylinder Addresses < or = 65,520

3390-9 EAV Volume

SHARE
in Anaheim
2012

IBM DS8000 HyperPAVs

- **Background – What are HyperPavs ?**
  - Heritage PAV design designed and implemented for ESCON attachments
    - **Relatively static binding of PAV-alias to PAV-base**
    - **Dynamically PAVs Tunable by zWLM on zWLM interval (5 min. default)**
    - **Complex rules of thumb for allocating the number of PAV-aliases**
    - **Many times PAVs are over allocated needlessly**
  - HyperPAV maintains a shared pool of PAV-aliases
    - **PAV-aliases bound to PAV-bases on an I/O request basis dynamically by DS8K on each I/O request**
    - **Up to 10X increase in efficiency**
- **Benefits**
  - 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 situations like '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 in the DS8000
  - Increases I/O Parallelism
  - On a HyperSwap, HyperPavs are effective immediately

Complete your sessions evaluation online at SHARE.org/AnaheimEval

SHARE
in Anaheim
2012

GDPS Alternate Subchannel Exploitation

- **GDPS/PPRC w/HyperSwap, GDPS PPRC HS Manager & TPC-R HyperSwap -> PPRC Targets placed in alternate Subchannel Set & Subchannel Set switch occurs on HyperSwap.**

- **GDPS/XRC**

  - **Removes restriction that all devices in a GDPS/XRC configuration be defined and identified with unique device addresses.**

  - **GDPS can bypass checking for unique device addresses**

    - **Allows duplicate device addresses to be used across SDM systems**
    - **Increases the number of devices that can be managed in a GDPS/XRC configuration.**

  - **Additional UCB constraint relief provided with**

    - **Allows the FlashCopy target devices to be defined only in the GDPS Controlling system but not to be defined in the SDM systems.**
    - **Support for a 'no UCB FlashCopy' where the FlashCopy targets need not be defined to any systems in the GDPS.**

- **Largest GDPS/MzGM w/HyperSwap environment is now approaching 1.2PB+ of primary data. (4.8pbs total). Largest GDPS/MGM w/HyperSwap customer has 450+TBs Primary Data. (2.25pbs total)**
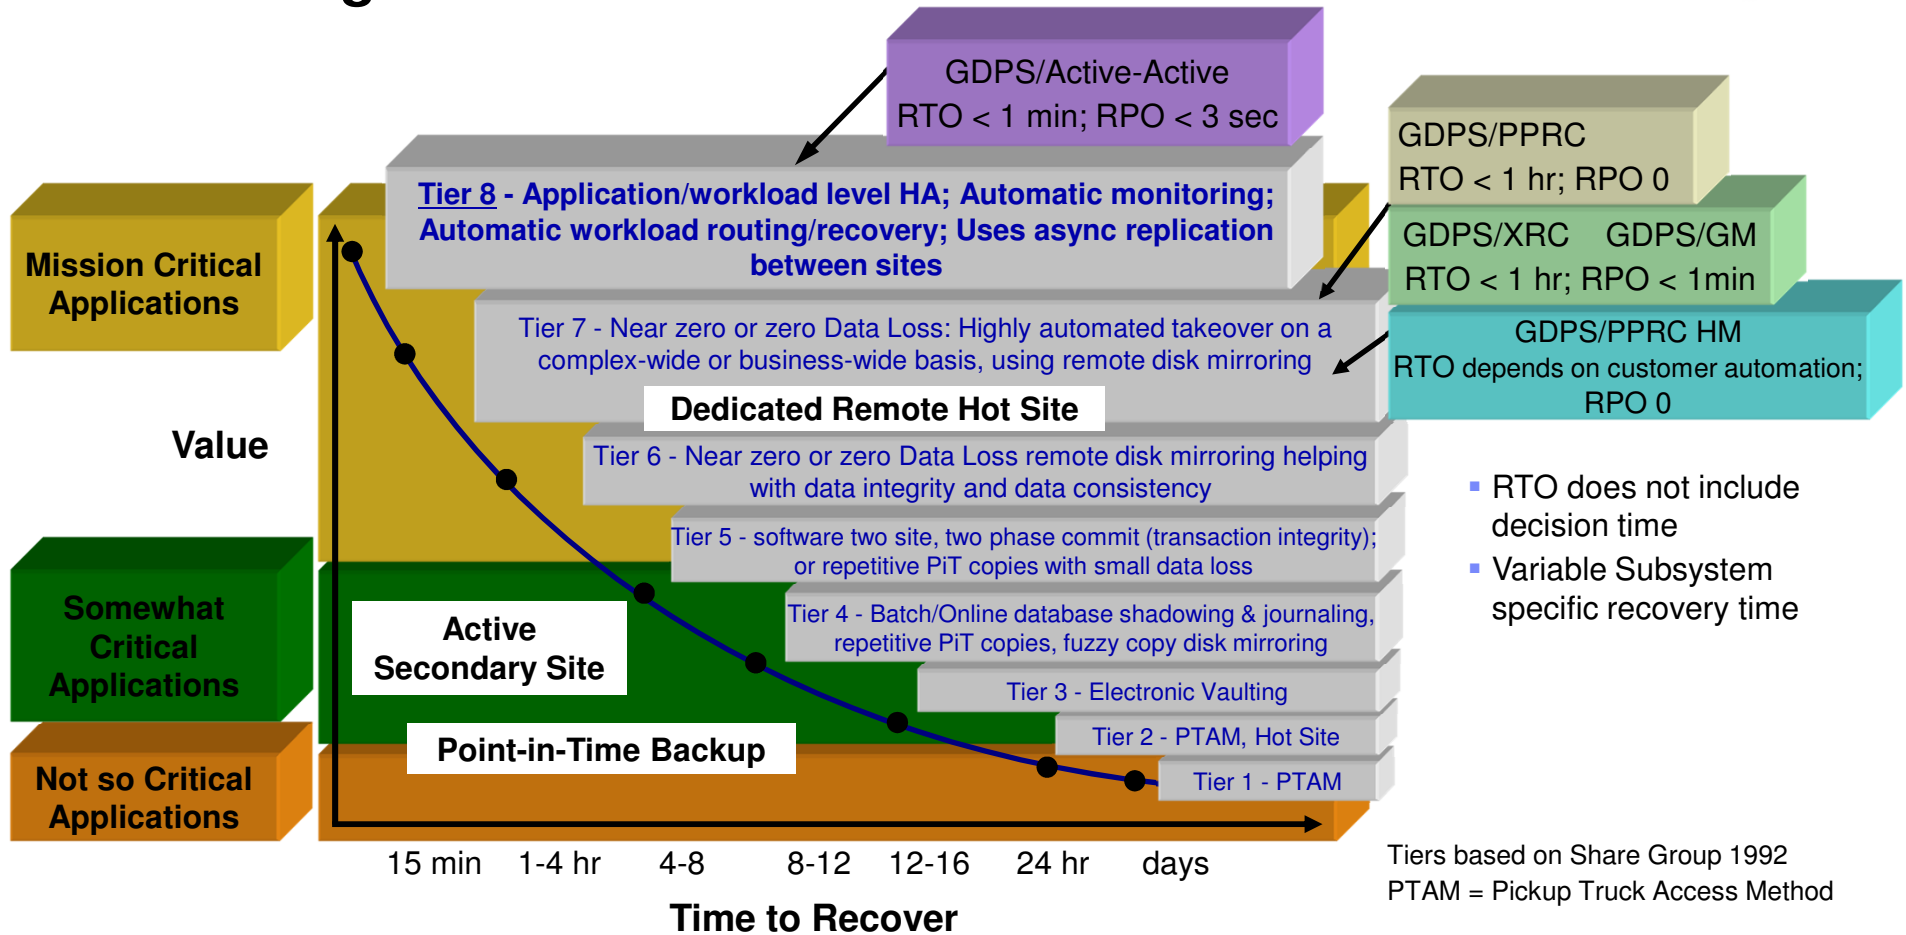
Business Continuity – Availability & Disaster Protection

Tiers of Disaster Recovery
Level Setting GDPS/Active-Active

SHARE
Technology · Connections · Results

GDPS/Active-Active
RTO < 1 min; RPO < 3 sec

GDPS/PPRC
RTO < 1 hr; RPO 0

GDPS/XRC   GDPS/GM
RTO < 1 hr; RPO < 1min

Tier 8 - Application/workload level HA; Automatic monitoring; Automatic workload routing/recovery; Uses async replication between sites

GDPS/PPRC HM
RTO depends on customer automation; RPO 0

Tier 7 - Near zero or zero Data Loss: Highly automated takeover on a complex-wide or business-wide basis, using remote disk mirroring

Mission Critical Applications

Value

Dedicated Remote Hot Site

Tier 6 - Near zero or zero Data Loss remote disk mirroring helping with data integrity and data consistency

Tier 5 - software two site, two phase commit (transaction integrity); or repetitive PiT copies with small data loss

- RTO does not include decision time
- Variable Subsystem specific recovery time

Somewhat Critical Applications

Active Secondary Site

Tier 4 - Batch/Online database shadowing & journaling, repetitive PiT copies, fuzzy copy disk mirroring

Tier 3 - Electronic Vaulting

Point-in-Time Backup

Tier 2 - PTAM, Hot Site

Not so Critical Applications

Tier 1 - PTAM

15 min   1-4 hr   4-8   8-12   12-16   24 hr   days

Time to Recover

Tiers based on Share Group 1992
PTAM = Pickup Truck Access Method

Failover models can only achieve so much in improving RTO

RPO = Recovery Point Objective - how much data to recreate?

RTO = Recovery Time Objective - how long being without service?

SHARE
in Anaheim
2012

Storage Based Data Replication Enabling Core Technologies

> **FlashCopy**
>> **Internal Copy**
>> **Available on:**
>> **DS6000, DS8000, ESS**
>> **SAN Volume Controller**
>> **DS4000, DS400, DS300, XIV**

PiT Incremental FlashCopy

+ Metro Mirror

Available on:

DS6000, DS8000, ESS

Time

Source Target

Write **Read**

Read and write to both source and copy possible

Optional background copy

When copy is complete, relationship between source and target ends

Primary

Prod

PiTCopies

PiT Incremental FlashCopy

Global Copy (Asynchronous)

WAN

REMOTE

PPRC-XD Secondary

FlashCopy

PiT Incremental Copy
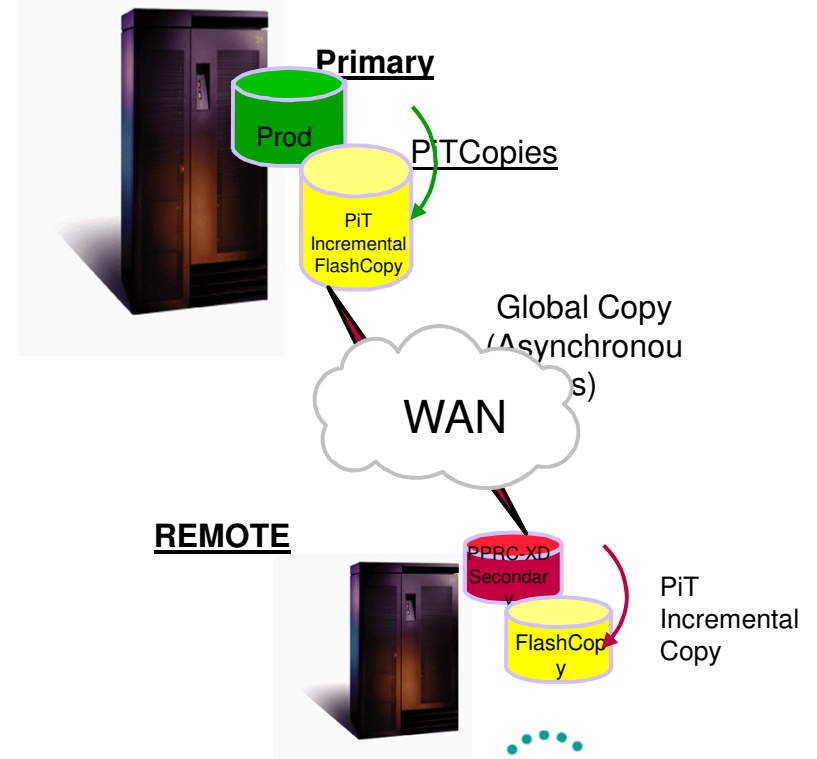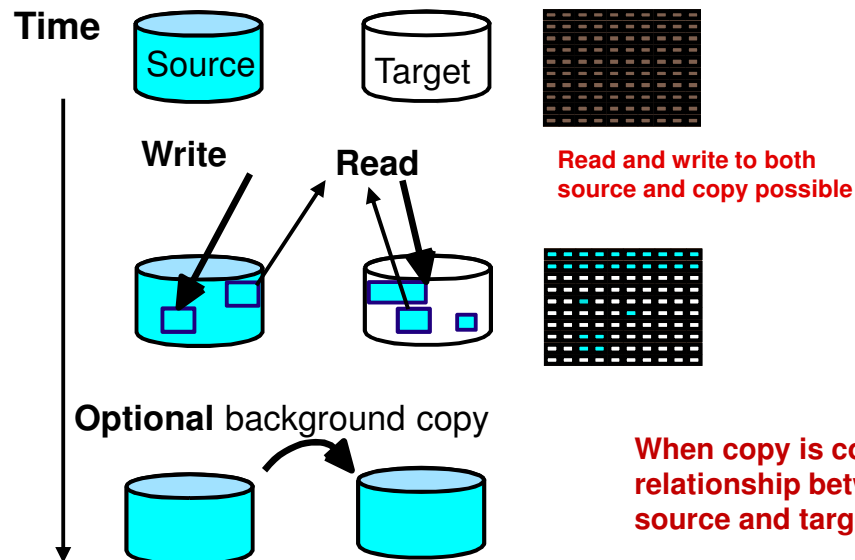
FlashCopy Client Scenarios

- **FC Targets are a GC Source (Send to remote site)**

- **zCDP for DB2, zCDP for IMS**

- **Fast Batch Restart After Batch Failure**

- **FlashCopies for Backups/Clones etc**

  - **Consistent FlashCopies**

  - **Backups -> PiT FlashCopy off MM, GM, zGM Targets**

    - **Data Freeze MM -> PiT FlashCopy Target Devices**
    - **Remote Pair FlashCopy**
    - **Pause GM -> FC target Devices**
    - **zGM (XRC) Zero Suspend FC**

- **FlashCopy as an Acceleration Function**

  - **Fast Defrag**

  - **DB Utilities**

zCDP for DB2 & zCDP for IMS – Eliminate Backup Windows

DB2 & IMS System Level Backup & System Level Restore
 - Backup calls HSM with DB Tables, HSM FlashCopy to SMS Copy Pool, then DB Logs.
 - DB2 & IMS Maintain Cross Volume Data Consistency. No Quiesce of DB required.

DFSMShsm function that manages Point-in-Time copies

• Combined with DB2 BACKUP SYSTEM, provides non-disruptive backup and recovery to any point in time for DB2 & IMS databases and subsystems (SAP)

Copy Pool Application

CP Backup Storage Group

FlashCopy

Onsite   Offsite

Dump to Tape

Multiple Disk Copies

• Up to 5 copies and 50 Versions of DB2 & IMS Image Copies, managed by MC.
• Automatic Expiration
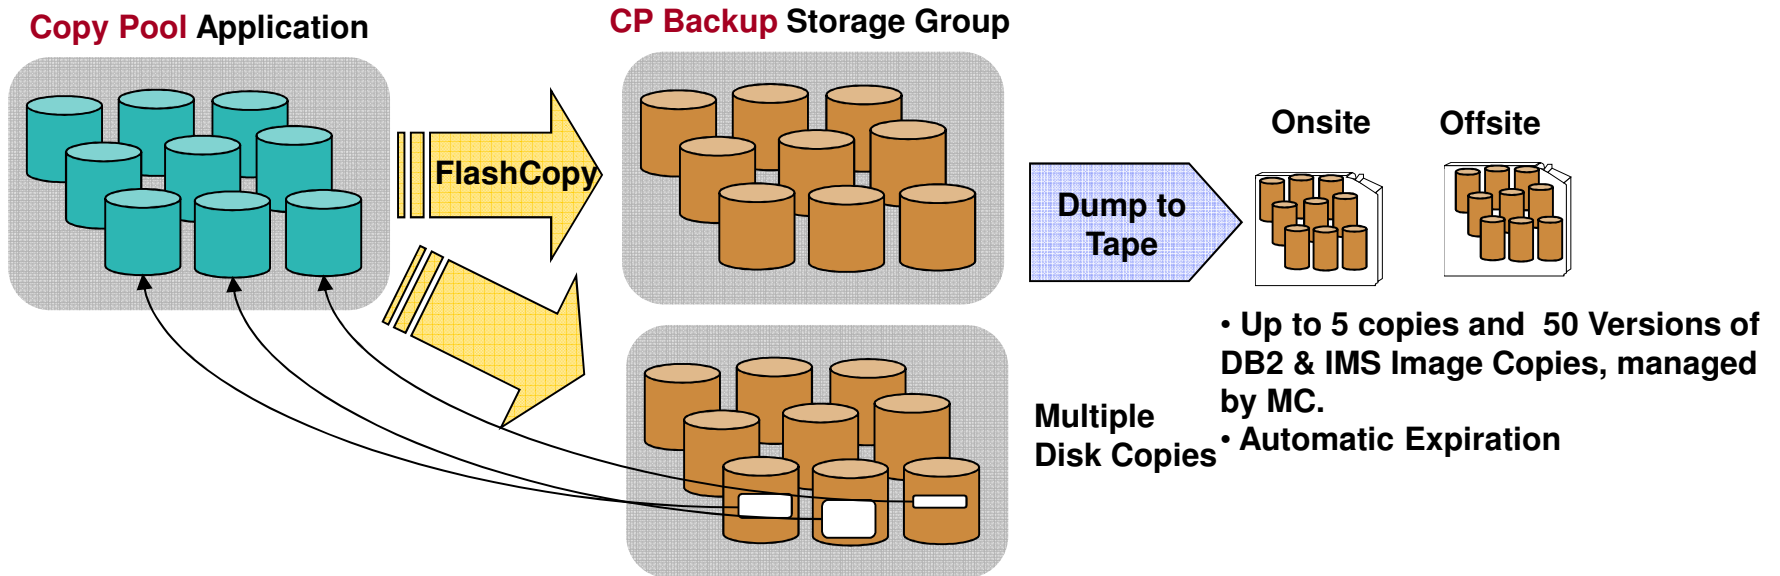
★**Recovery at all levels from either disk or tape!**

• Entire copy pool, individual volumes and …
• Individual data sets

SHARE
in Anaheim
2012

Storage Based Data Replication Enabling Core Technologies

➤ Metro Mirror

- ➤ Synchronous mirroring
- ➤ Available on:
  - ➤ DS6000, DS8000, ESS
  - ➤ SAN Volume Controller
  - ➤ DS4000, DS400, DS300,
  - ➤ XIV

➤ Global Mirror

- ➤ Asynchronous mirroring
- ➤ Available on:
  - ➤ DS6000, DS8000, ESS
  - ➤ DS4000, DS5000
  - ➤ XIV

➤ Metro Global Mirror

- ➤ Three site synchronous and asynchronous mirroring
- ➤ DS6000, DS8000

Primary Site A — Metro Site B

MM Primary Volume → MM Secondary Volume

Up to 300 km

Primary Site A — GM Primary Volume

Out of Region Site B — GM Secondary Volume — FlashCopy Journal Volume

Unlimited Distance

Primary Site A — MM Primary Volume → Metro Site B — MM Secondary GM Primary Volume

Out of Region Site C — GM Secondary Volume — FlashCopy/Journal Volume

Unlimited Distance
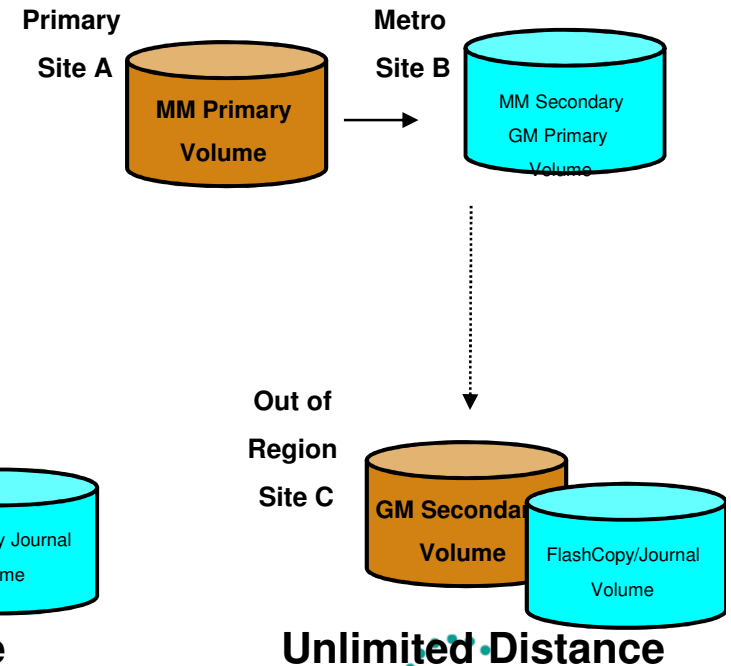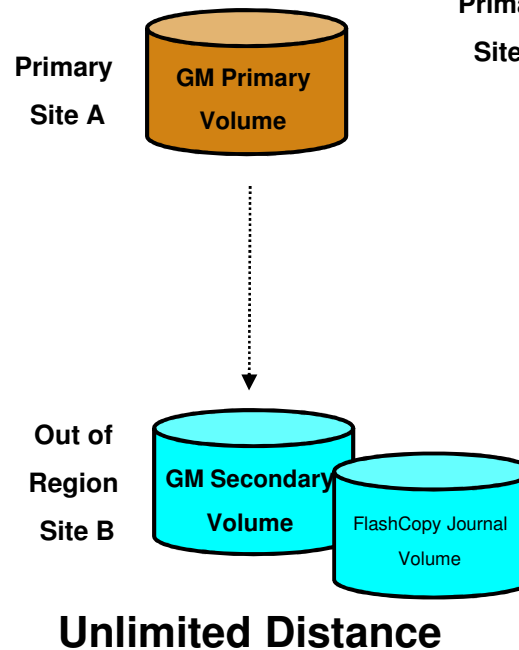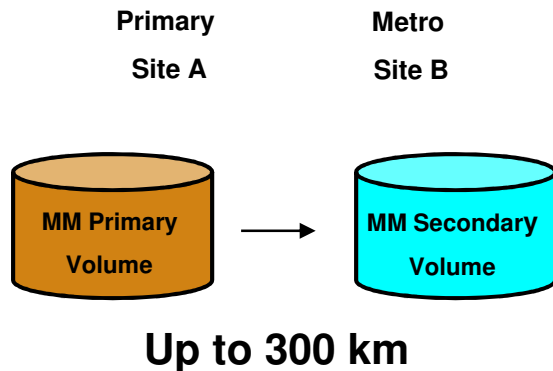
DS8700/DS8800 System z - Storage Based Data Replication Enabling CoreTechnologies

- **z/OS Global Mirror (XRC)**
  - **Asynchronous mirroring – System z ONLY)**
  - **Available on DS8000 & ESS**

- **z/OS Metro Global Mirror (MzGM)**
  - **Asynchronous mirroring – System z ONLY)**
  - **Available on DS8000**
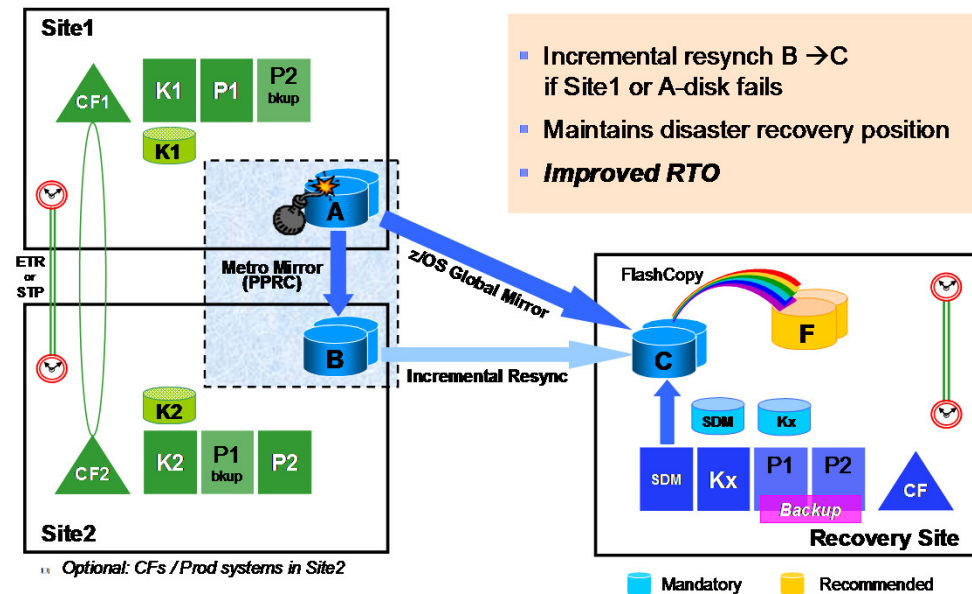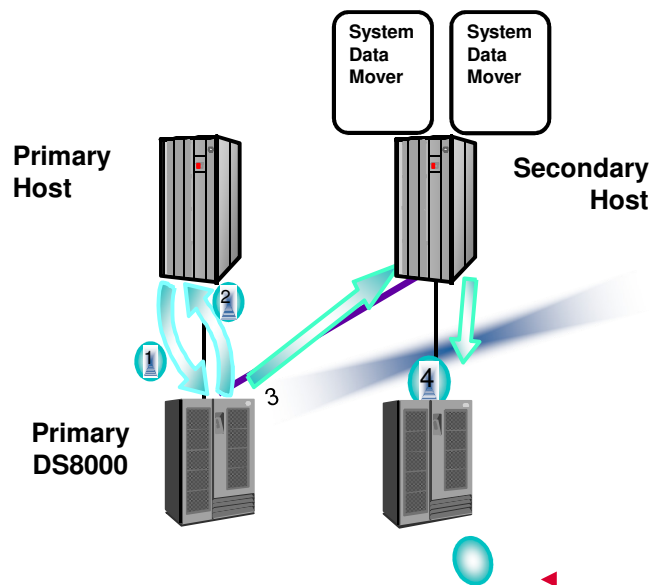
IBM HyperSwap Technology: Near Continuous Data Availability
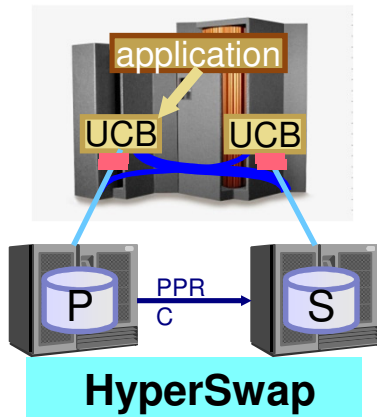
HyperSwap

DS8700/DS8800 R6.2 New Function:
Hyperswap Synergy zOS Resilience (Health Status Message)

Improves detection of when to trigger a Hyperswap. Provides a mechanism for the DS8K to alert the host of specific events that may cause an adverse effect to host operations. These messages indicate to GDPS that action may need to be taken in order to preemptively avoid problems on this controller and to swap to the secondary DS8K.

Link Summary Event Notification

Improves detection and handling of PPRC Suspends Allows for the DS8K to notify all hosts connected to a single LSS for a PPRC suspend event that spans across multiple volumes in that LSS through a single summary notification mechanism.

✓ **Designed to Provide Continuous Availability of Data for System z**
  ✓ **Facilitated by new PPRC microcode functionality and z/OS® IOS code**
    Also now available with AIX (OpenSwap) managed by TPC-R

✓ **HyperSwap™ is:**
  ✓ **Integration of very fast swapping of PPRC'd disk subsystems with z/OS, System z hardware, and GDPS or TPC-R.**
  ✓ **Switching to alternate copy of System z data can be accomplished in seconds to minutes**
  ✓ **Supported on Synchronous PPRC**
  ✓ **HyperSwap Options:**
    ✓ **z/OS Basic HyperSwap,**
    ✓ **TPC-R Full Function HyperSwap,**
    ✓ **GDPS HyperSwap Manager**
    ✓ **GDPS/PPRC w/HyperSwap Full Function**

✓ **Intended Benefits:**
  ✓ **Designed to offer continuous availability of data**
    ✓ **Disk Maintenance**
    ✓ **Site Maintenance**
    ✓ **Data Migration**
    ✓ **Disk Failure**
    ✓ **Site Failure**
  ✓ **Fast and Scalable System z Enterprise Data Center swap: scales to very large configurations**
  ✓ **Repeatable, reliable, confident recovery: No operator interaction,**

SHARE in Anaheim 2012

IBM DS8000 Feature - Remote Pair FlashCopy
(ex. GDPS/PPRC HS + FlashCopy)

Function: FlashCopy to PPRC Source, send FC command to target instead of all data (Duplex Pending). Maintains HyperSwap being Active.

(Ex. zCDP for DB2 Backups at both Sites w/HyperSwap.)

Scope & Restrictions:

▪Full volume and data set level operations

     - IBM Remote Copy FC can be combined with:

         - Incremental FlashCopy

         - Background copy or no background copy

         - Nocopy to copy

         - FlashCopy consistency groups

         - FlashCopy of open devices using a CKD access device

▪FlashCopy features not supported with Preserve Mirror function:
  –Commit, Revert, Fast Reverse Restore
  –Space Efficient FlashCopy (Source or Target)

▪Local target and remote target cannot be space efficient

▪Both PPRC pairs must be Metro Mirror pairs in full duplex

▪Withdraw of Preserve Mirror relationship will cause bit to be set for PPRC pair at target indicating not true mirror while still full duplex

▪Function will be provided with REQUIRED and PREFERRED options

Local Storage Server

Remote Storage Server

Local A — *full duplex* → Remote A

FlashCopy

FlashCopy   Metro Mirror   *FlashCopy*

Local B — *full duplex* → Remote B

SHARE
in Anaheim
2012
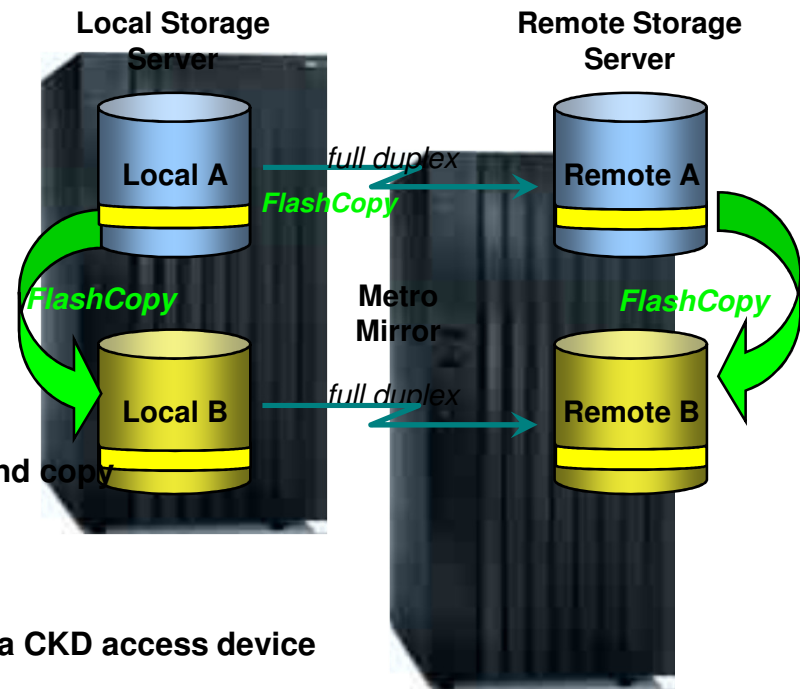
Data Replication Management Options
- **PPRC Manager (FlashCopy Manager)**
- **TPC-R**
- **GDPS/PPRC HyperSwap Manager**
- **GDPS/PPRC HyperSwap**

TPC Replication Manager

- Setup Copy Sessions
- Execute Copy Operations
- Monitor Copy Status
- Manage/Monitor Consistent Groups
- Alert Operations on Exceptions / Failures
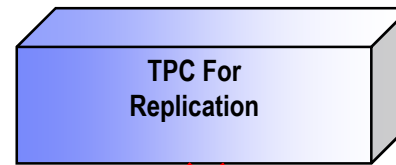
- DS6000, DS8000 support
- Global Mirror Support
- Replication Progression Monitoring
- High Availability
- Disaster Recovery Automation (failover, failback)

TPC For Replication

Primary/Source Site

DS8000
DS6000
ESS
SAN Volume Controller

Second/Target Site

DS8000
DS6000
ESS
SAN Volume Controller

- Automated copy services configuration
- Central operations for copy services
- Operational status on copy services operations
- Assistance with recovery on failures

METRO

SHARE in Anaheim
2012

GDPS Solutions

There are multiple GDPS service products under the GDPS solution umbrella to meet various customer requirements for Availability and Disaster Recovery

GDPS/PPRC HM	GDPS/PPRC	GDPS/GM & GDPS/XRC	GDPS/MGM & GDPS/MzGM	GDPS/Active-Active
Continuous Availability of Data within a Data Center	Continuous Availability / Disaster Recovery within a Metropolitan Region	Disaster Recovery at Extended Distance	Continuous Availability Regionally and Disaster Recovery Extended Distance	Continuous Availability, Disaster Recovery, and Cross-site Workload Balancing at Extended Distance
Single Data Center Applications remain active — Continuous access to data in the event of a storage subsystem outage — RPO=0 & RTO=0	**Multi-site workloads can withstand site and/or storage failures** — A/S RPO=0 & RTO<1 hr or A/A RPO=0 & RTO mins	**Two Data Centers** Rapid Systems Disaster Recovery with "seconds" of Data Loss — Disaster recovery for out of region interruptions — RPO secs & RTO <1 hr	**Three Data Centers** High availability for site disasters — Disaster recovery for regional disasters — A/S RPO=0 & RTO<1 hr or A/A RPO=0 & RTO mins and RPO secs & RTO <1 hr	**Two or More Data Centers** — All sites active — RPO secs & RTO secs
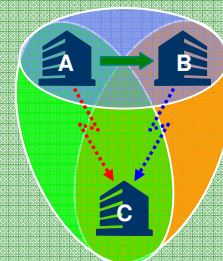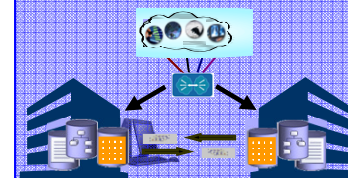
Components

Tivoli – NV, SAz **STG** – System z, DS8K, PPRC **GTS** – GDPS code, Services	**Tivoli** – NV, SAz, SA MP, AppMan **STG** – System z, DS8K, PPRC, VTS **GTS** – GDPS code, Services	**Tivoli** – NV, SAz **STG** – System z, DS8K, GM, XRC **GTS** – GDPS control, Services	**Tivoli** – NV, SAz **STG** – System z, DS8K, MGM, MzGM **GTS** – GDPS code, Services	**Tivoli** – NV, SAz **AIM** – Multi-site Workload Lifelife **IM** – DB2 & IMS replication **STG** – System z, DS8K, GC **GTS** – GDPS code, Services

z/OS GDPS Active/Standby at Distance

- Active/Standby Sites is positioned as the next generation of GDPS
- Sites separated by <u>unlimited</u> distances, running same applications and having the same data to provide cross-site Workload Balancing and Continuous Availability / Disaster Recovery
- Customer data at geographically dispersed sites kept in sync via replication
- Configurations: Active/Standby, Active/Query (SOD)

GDPS/PPRC	GDPS/XRC or GDPS/GM	Active/StandBy (Appl FO/FB)
Failover Model	Failover Model	Near CA model
Recovery Time ≈ 2 min	Recovery Time < 1 hour	Recovery time < 1 minute
Distance < 20 km	Unlimited distance	Unlimited distance
		DB2 & IMS Data
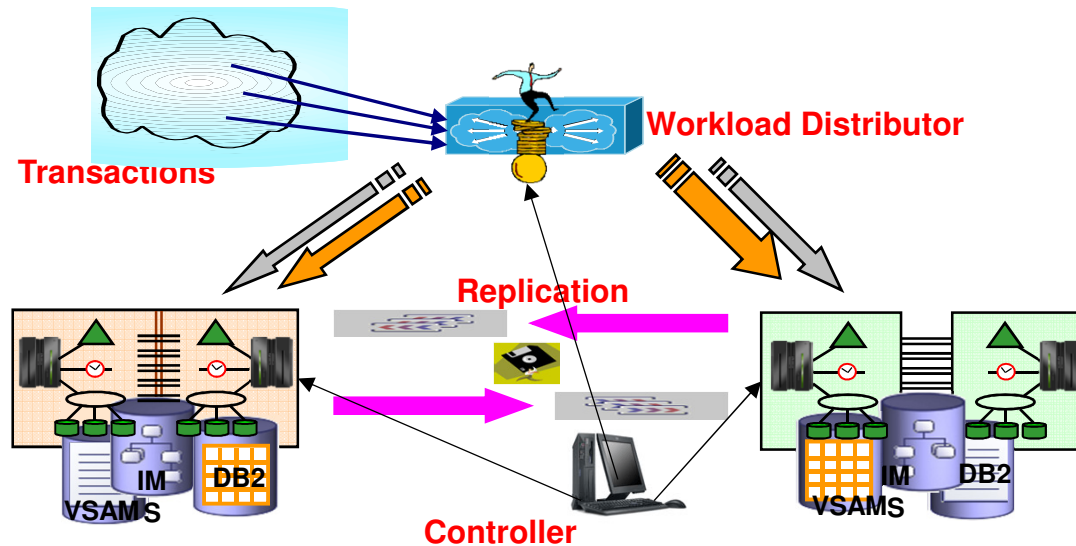
Transactions

Workload Distributor

Replication

Controller

IM DB2 VSAMS

SHARE
in Anaheim
2012

z/OS + DS8700/DS8800 Synergy Summary

Distributed Systems and DS8000 Synergy Items

- Pre-Deposit Write (Copy Services one protocol exchange & Keeps Pipe Full)
- Caching Algorithms – AMP, ARC, WOW, 4K Cache Blocking
- Easy Tier
- I/O Priority Manager

Performance

- HyperSwap (AIX)
- GDPS x/DR, GDPS DCM
- Power HA p, Power HA i
- Remote Pair FlashCopy & Enhancements
- FlashCopy & Functions like Fast Reverse Restore

Availability

- Dynamic Volume Expansion
- Thin Provisioning
- Space Efficient FlashCopy
- z/OS Distributed Data Backup
- Disk Encryption

Management/Growth

SHARE in Anaheim 2012

z/OS and DS8000 Synergy Items

- **zHPF Enhancements (now includes all z/OS DB2 I/O)**

  – **Extended Distance FICON**

  – **Caching Algorithms – AMP, ARC, WOW, 4K Cache Blocking**

  – **DFSMS Recognition of SSDs**

  – **Easy Tier**

  – **z/OS GM Multiple Reader Support**

  – **SSDs + DFSMS + zHPF + HyperPAV + DB2**

  – **I/O Priority over Ficon & within DS8K managed by zWLM Service Class**

  – **zWLM + DS8K I/O Priority Manager**

  – **HyperPAV**

  – **GDPS & GDOC Automation**

  – **GDPS/Active/Standby**

  – **HyperSwap Technology Improvements**

  – **Remote Pair FlashCopy & Enhancements**

  – **zCDP for DB2, zCDP for IMS – Eliminating Backup Windows**

  – **1 TB EAVs**

  – **Quick Init for CKD Volumes**

  – **Dynamic Volume Expansion**

  – **Space Efficient FlashCopy**

  – **z/OS Distributed Data Backup**

  – **System z Discovery & Automatic Configuration (zDAC)**

  – **Alt Subchannel Exploitation**

  – **Disk Encryption**

Performance

Availability

Management/Growth

RED – DS8700/DS8800 R6.2 LIC

SHARE in Anaheim

2012

Additional Information, References, Disclaimers and Trademarks etc.

References

- **TechDocs White Paper: IBM Handbook on Using DS8000 Data Replication for Data Migration -** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101716

- **TechDocs White Paper: IBM z/OS Multi-Site Business Continuity** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101635

- **TechDocs White Paper: IBM DS8800 Data Consolidation** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102100

- **TechDocs White Paper: IBM HyperSwap Technology April 2010** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101289

- **TechDocs White Paper: IBM System z and DS8000 z/OS Synergy** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101528

- **Title: IBM z/OS Data Corruption Trends & Directions** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101804

- **Redpaper: IBM Storage Infrastructure for Business Continuity** http://www.redbooks.ibm.com/abstracts/redp4605.html?Open

- **Redpaper: IBM System Storage DS8700 Easy Tier** http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/redp4667.html?Open

Additional Information

- Web sites:
  GDPS   www.ibm.com/systems/z/gdps
  Parallel Sysplex   www.ibm.com/systems/z/pso
  Bus Resiliency z   www.ibm.com/systems/z/resiliency
  Bus Resiliency   www.ibm.com/systems/business_resiliency
  System z www.ibm.com/systems/z/hardware
  Storage www.ibm.com/systems/storage

- **Redbooks®**   **GDPS Family: An Introduction to Concepts and Capabilities**
  www.redbooks.ibm.com/abstracts/sg246374.html?Open

- **GDPS Web Site White Papers and Presentations**
  - GDPS: The Ultimate e-business Availability Solution
  - IBM Implementation Services for GDPS/Global Mirror
  - GDPS Business Continuity Solutions
  - Consistency Groups in a Nutshell
  - DS6000™ / DS8000™ Data Replication
  - GDPS Solutions

- e-mail: gdps@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
LINUX is a registered trademark of Linux Torvalds
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
Intel is a registered trademark of Intel Corporation
* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Complete your sessions evaluation online at SHARE.org/AnaheimEval

Disclaimers