#SHAREorg

SHARE

IBM System z & Storage Synergy

Lisa Gundy
IBM Corporation
(lisat@us.ibm.com)

Corey Smeaton
IBM Corporation
(csmeato@us.ibm.com)

August 15, 2013
Session Number 14139

SHARE in Boston

Agenda

- DS8870 Overview
- DS8870 + z/OS Synergy
  - Business Continuity
  - Recent Synergy Enhancements
- Summary

DS8870 -> 5th Generation DS8000 Disk System

- Building on a market-proven, reliable code base!
- 94% of the same proven microcode

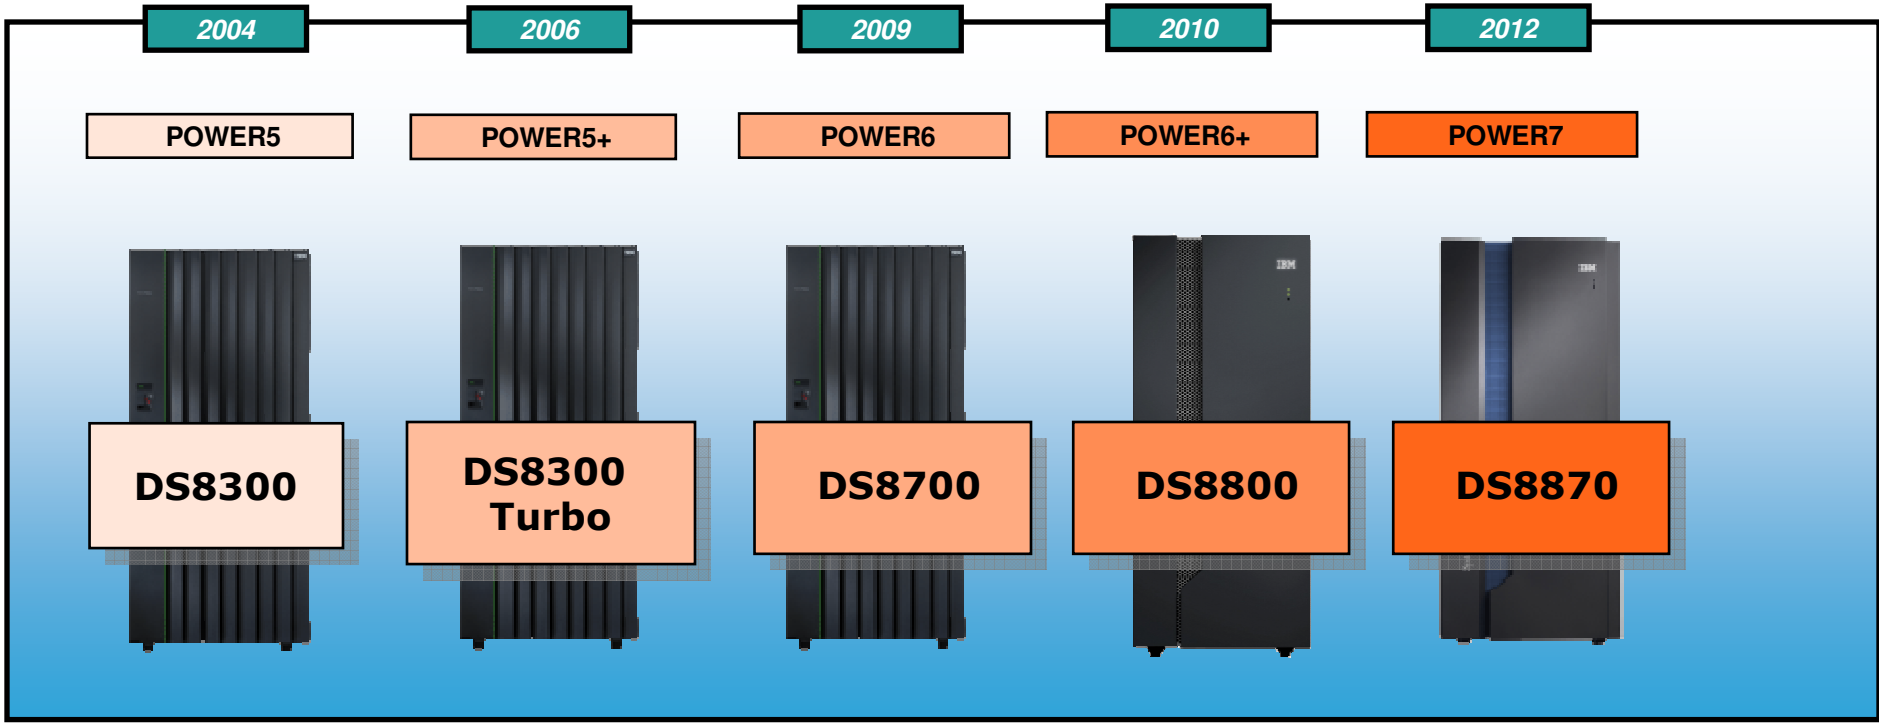

- *Designed for Enterprise environments with over 5-9's availability natively*
- *Designed for Enterprise environments with over 6-9's availability when DS8000 with Metro Mirror is combined with GDPS/PPRC HyperSwap*

DS8000 Enterprise Disk Systems – Hardware Evolution

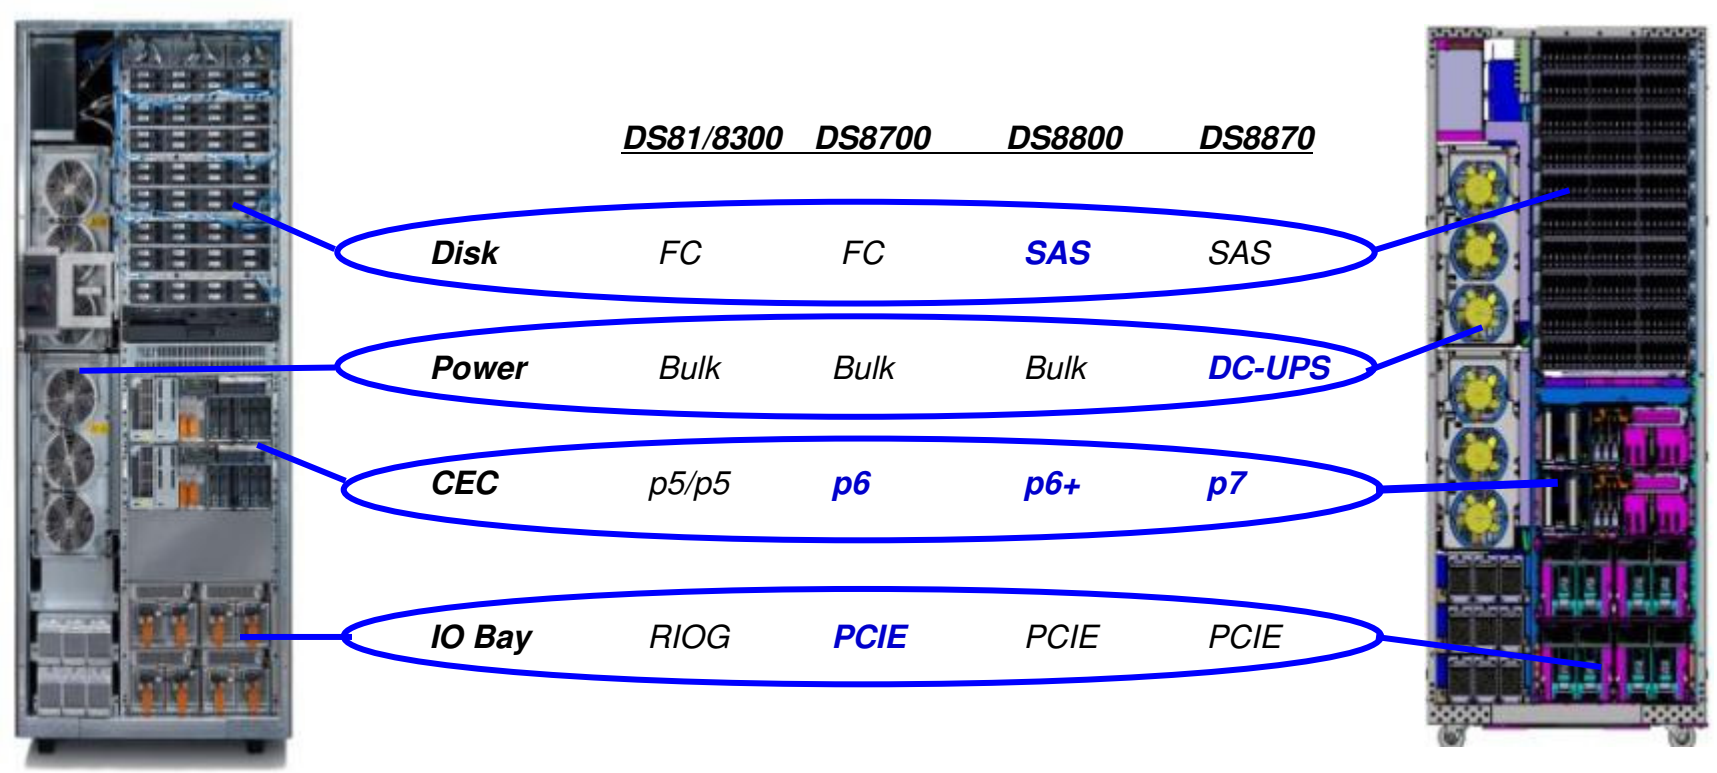

- Incremental changes between versions maximizes quality

IBM System Storage DS8870 — Features That Continue

Easy to use GUI

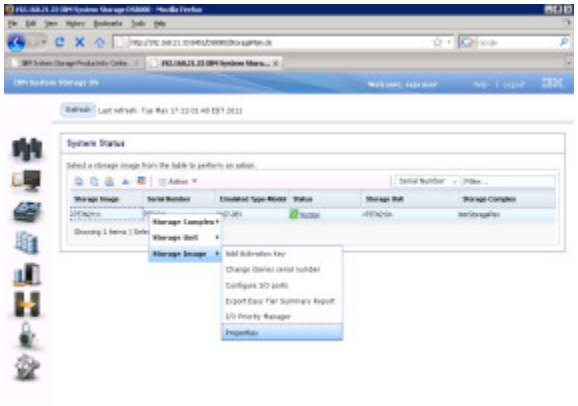

Storage Pool Striping

Storage Pool Striping is an algorithm choice for volume creation which allows for better backend disk utilization

Volumes are created by allocating one Extent from available Ranks in an Extent Pool, in a round-robin fashion

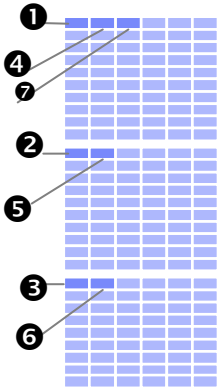

"Quick Init"

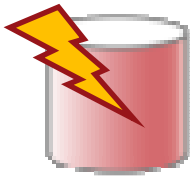

Quick Initialization provides volume initialization that is up to 2.6 times faster and therefore allows the creation of devices and making them available as soon as the command completes.

Align Risk and Performance

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

TIER 7: RPO=near zero, RTO <1min, Automatic Continuous Availability

TIER 6: RPO=Near Zero, RTO <1Hr. to 4 hours, Manual Disk or Tape Data Mirroring

TIER 4: RPO > 15 min. RTO= 4+ hours, Manual PiT or SW Data Replication.

I/O Priority Queuing

I/O Priority Manager attempts to make sure the most import I/O operations get serviced when a given rank is overloaded by the workload on the storage system

EASY TIER

Easy Tier 1 (DS8700 R5.1)

Automated cross-tier performance management for *SSD/HDD hybrid pools*

Easy Tier 2 (R6.1)

Automated cross-tier performance or storage economics management for hybrid pools with *any 2 tiers* (SSD/ENT, SSD/NL or ENT/NL)

Easy Tier 3 (R6.2)

Automated cross-tier performance and storage economics management for hybrid pools with *3 tiers* (SSD/ENT/NL)

Easy Tier 4 (R6.3)

Support for *encryption capable environments*

WLM Support for I/O Priority Manager in DS8K Series

- WLM collaborates with the I/O Priority Manager in DS8K storage servers.
  - This feature is supported on IBM System Storage® DS8K series, and requires a DS8K licensed machine code
- WLM sends I/O Priority Manager information about the goal fulfillment and importance of z/OS workloads (service classes).
- Passing these performance parameters to the storage server enables the I/O Priority Manager to determine which I/O requests are more important than others and which I/O requests need to be processed faster to fulfill the performance goals defined for the corresponding workload in z/OS.
- Using the passed information from WLM, the I/O Priority Manager throttles I/O requests of workloads which exceed their goals to help I/O requests of workloads which do not fulfill their goals.
- New IEAOPT parameter STORAGESERVERMGT={YES|NO}

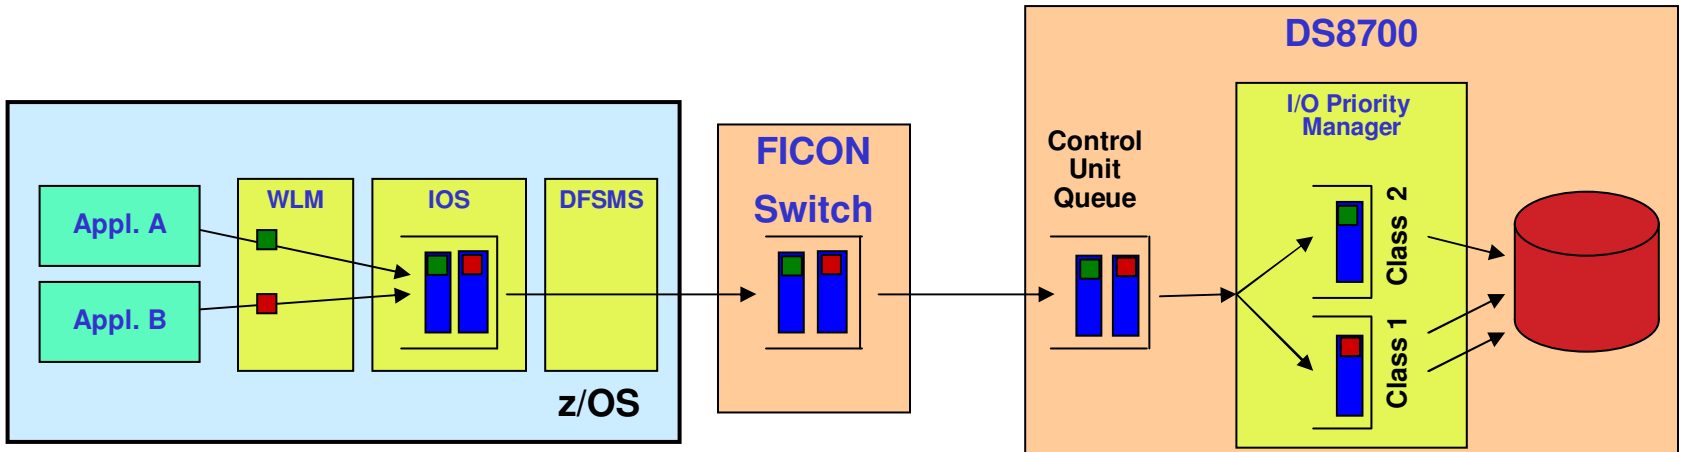

DS8870 IBM Easy Tier

Storage Tier *Optimization*

- Monitors performance of each ‘extent’ (1 GB, sub-volume level or 3390-1 equivalent) to determine the data ‘temperature’
- Creates extent migration plan for optimal data placement every 24 hours based on performance statistics
- Migrates extents within an extent pool according to plan over 24-hour period
- A limited number of extents are chosen for migration every 5 minutes

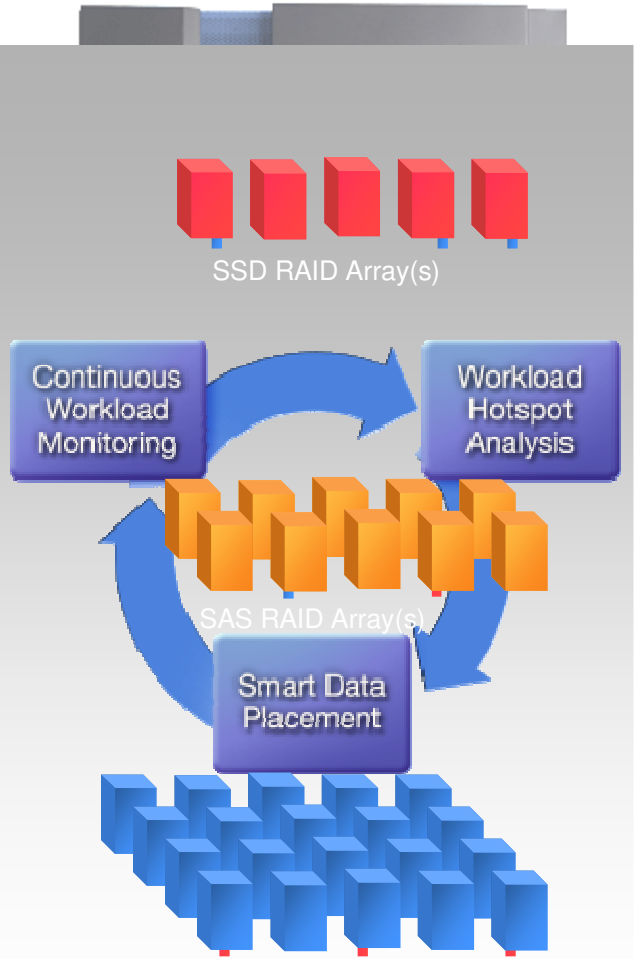

IBM Self-Optimizing Storage with Easy Tier

Easy Tier balances performance and cost automatically

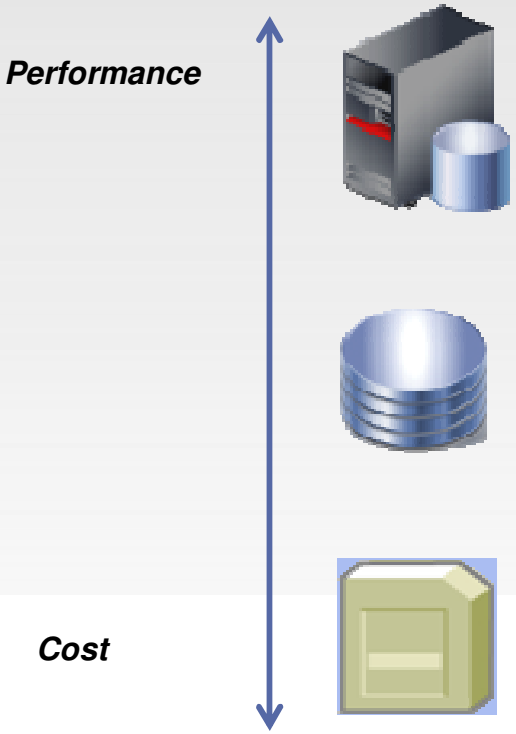

- Automatic movement of data to the right disk tier to balance cost and performance
- Continuous rebalancing ***within*** each tier to maintain peak performance across all drives
- Maximum benefit when Easy Tier extends beyond the disk system

DS8870 + z/OS Synergy
- Business Continuity

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** 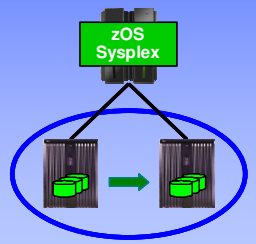 **RPO=0 & RTO=0**	**Multi-site workloads can withstand site and/or storage failures** 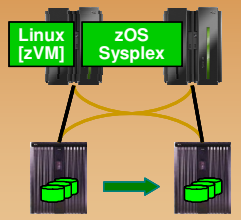 **A/S RPO=0 & RTO<1 hr or** **A/A RPO=0 & RTO mins**	**Two Data Centers** **Rapid Systems Disaster Recovery with "seconds" of Data Loss** **Disaster recovery for out of region interruptions** 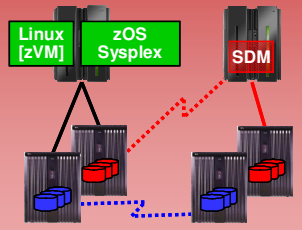 **RPO secs & RTO <1 hr**	**Three Data Centers** **High availability for site disasters** **Disaster recovery for regional disasters** 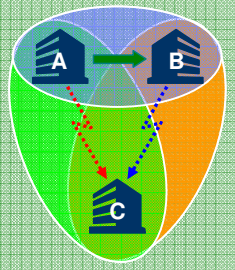 **A/S RPO=0 & RTO<1 hr or A/A RPO=0 & RTO mins and RPO secs & RTO <1 hr**	**Two or More Data Centers** **All sites active** 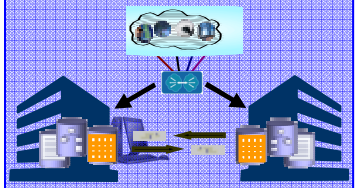 **RPO secs & RTO secs**
Components				
Tivoli – NV, SAz **STG** – System z, DS8K, PPRC **GTS** – GDPS code, Services	**Tivoli** – NV, SAz, SA MP, AppMan **STG** – System z, DS8K, PPRC, VTS **GTS** – GDPS code, Services	**Tivoli** – NV, SAz **STG** – System z, DS8K, GM, XRC **GTS** – GDPS control, Services	**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 Active / Active at Distance – Concept & Value

- Active/Active Sites is positioned as the next generation of GDPS
- Sites separated by unlimited 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
Failover Model	Failover Model	Near CA model
Recovery Time ≈ 2 min	Recovery Time < 1 hour	Recovery time < 1 minute
Distance < 20 km sites	Unlimited distance	Unlimited distance between

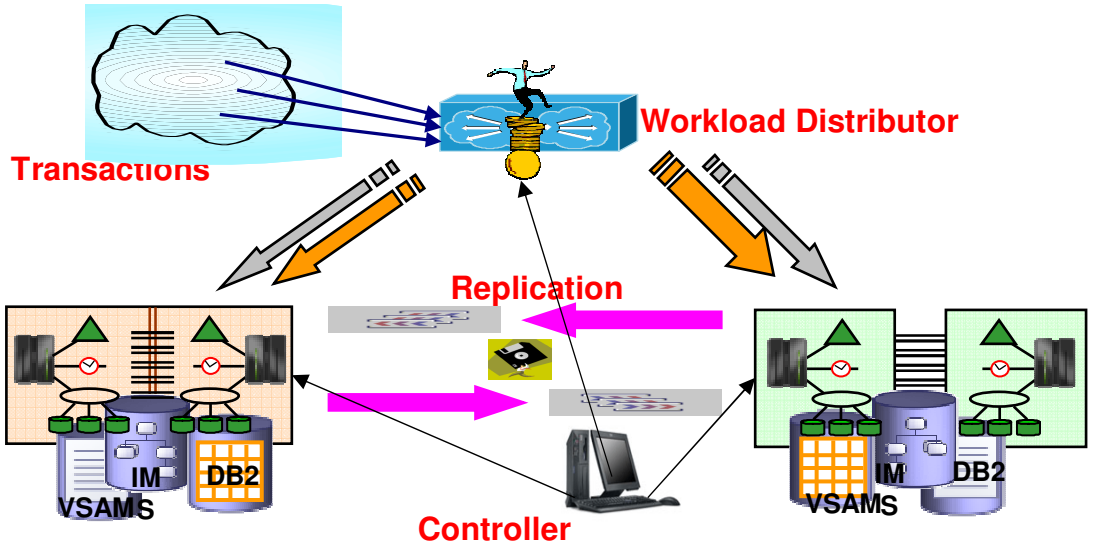

DS8870 + z/OS Synergy
- Recent Synergy Enhancements

GM Time Synch

Active/Active with Software Replication

- GDPS Active/Active (A/A)
  - Problem = Distributed workload not consistent with System batch or other replicated data in A/A configuration after disaster

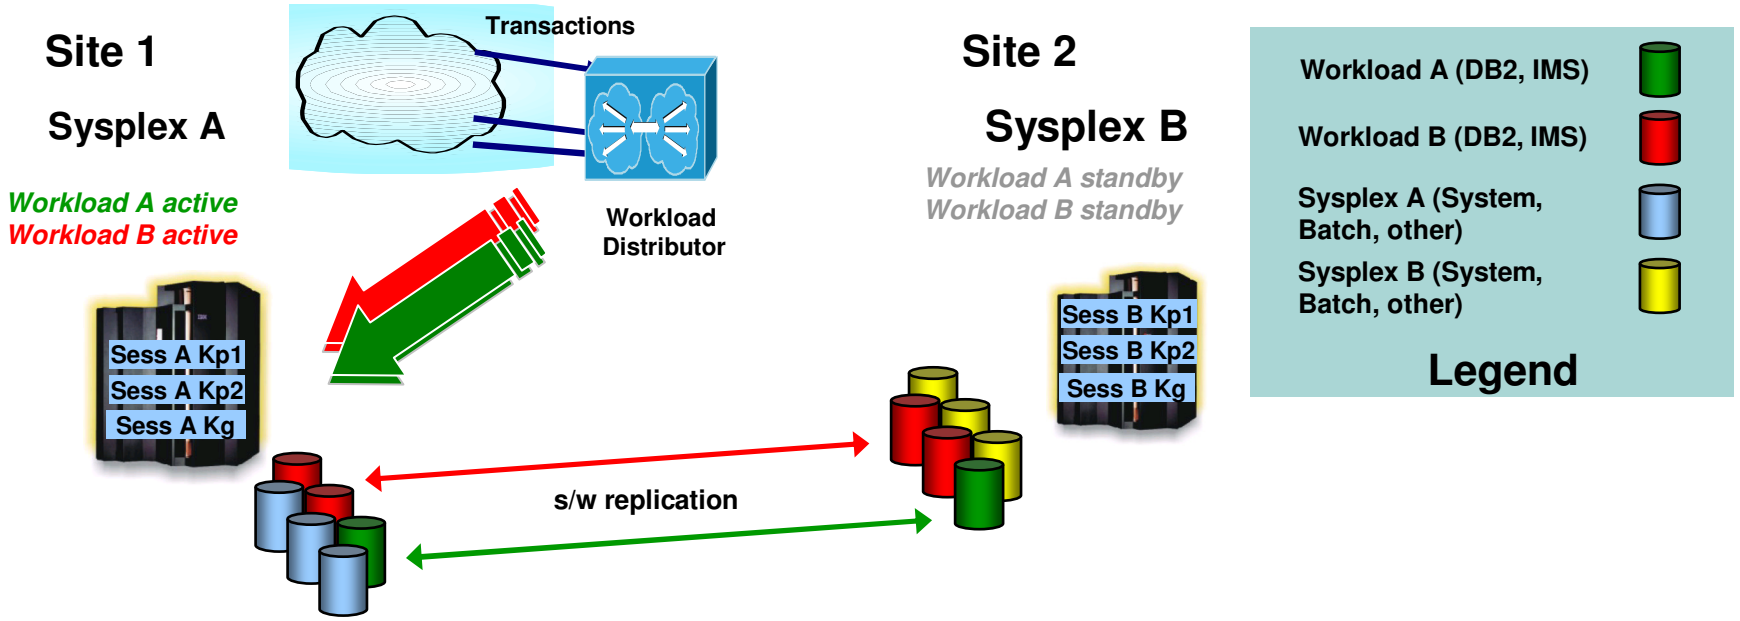

GM Time Synch

- GDPS Active/Active (A/A)
  - Solution = integrate disk replication into the A/A environment
  - Global Mirror (GM) integrated with software based replication

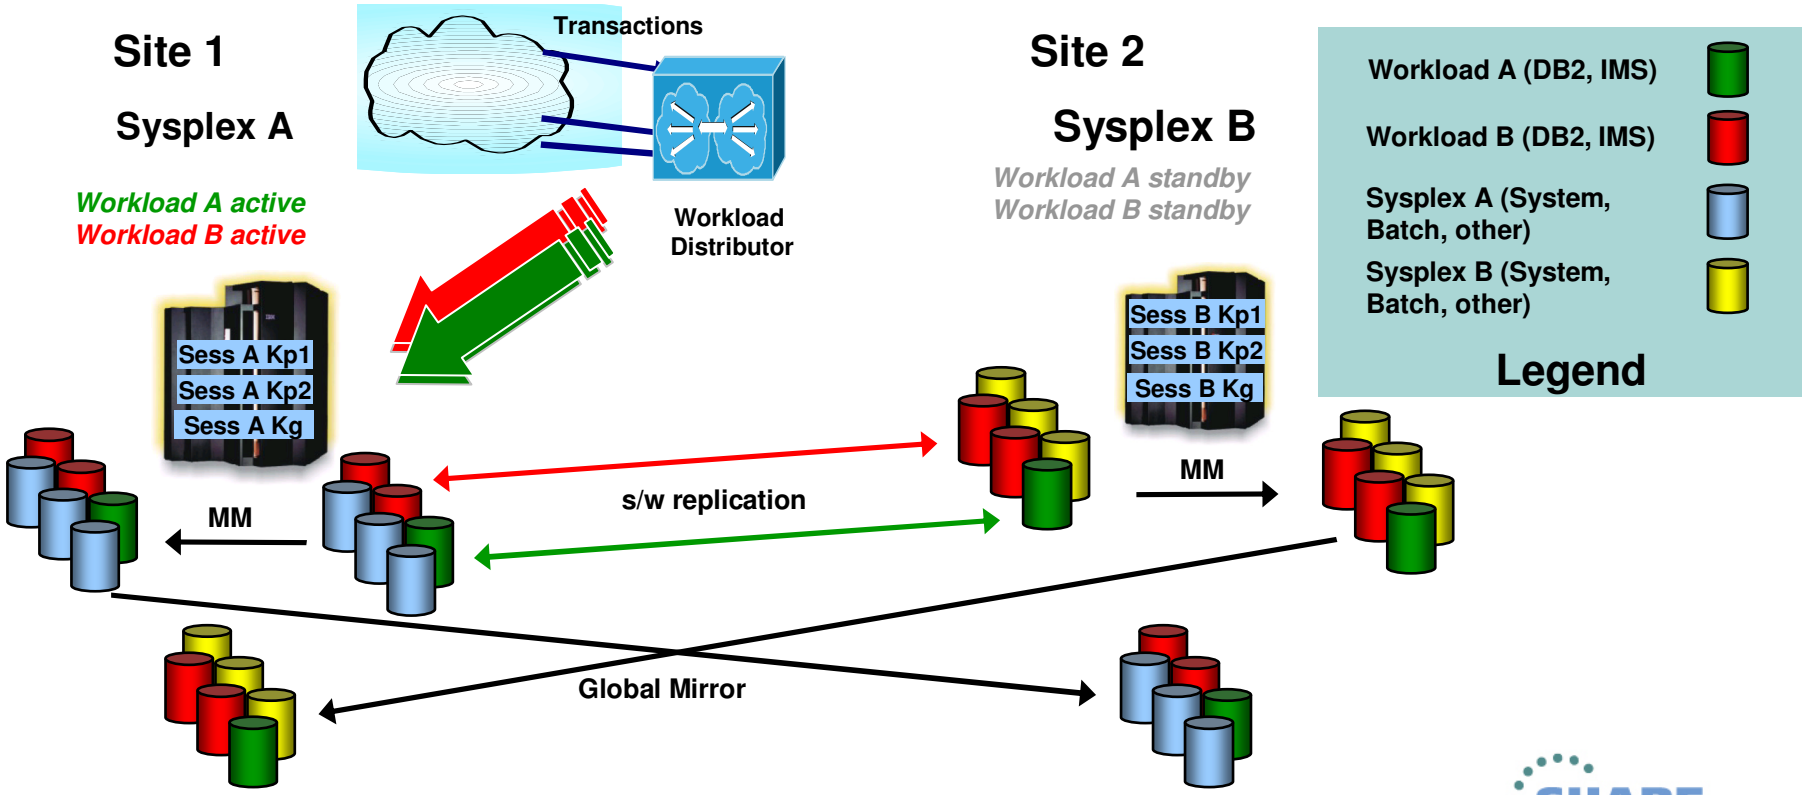

GM Time Synch

- Replication techniques use different clocks
  - Software replication time stamps = sysplex time reference
  - Global Mirror consistency group time = DS8K clock
  - Different mechanisms unaware of other time references

- Time synch function provides reference for GM to derive sysplex time
  - Leveraging existing behavior
    - Host timestamps read and write I/Os
  - I/O timestamps used to obtain the sysplex clock time
    - GM can report the consistency time based on the sysplex time

GM Time Synch

- Get sysplex name for system
  - Set System Characteristics (SSC) identifies sysplex name to SFI
  - GM session START/RESUME modified to allow specification of sysplex name
  - Timestamped I/O from the associated sysplex name used by GM master to calculate
    - *Current sysplex time*
    - *Delta between SFI and sysplex clocks*
    - *Drift*

GM Time Synch

- Consistency group (CG) formation
  - GM master estimates the current sysplex time
    - Applies delta to the current SFI clock value.
  - Calculated sysplex time used as FlashCopy sequence number during consistency group formation

- Disaster at the GM primary:
  - CG at remote site contains consistent copy of data at a known sysplex time
    - Used as PiT reference to recover with software based replication

GM Time Synch

```
RQUERY  ACTION(GMLSTAT)
RQUERY Output Devn(0F41) SCHSET(0) Action(GMLSTAT) Version(003)
*****************************************************************

SNbr GMLStat    GoodCg   Pct CrnBadCG TotBadCG LastGoodCGSCntlClock
--   ---------- -------- --- -------- -------- --------------------
33   Running    00000586  85 00000000 000000F1 26 Mar 2013 12:22:27

Master: Serial       SSID LSS   CGInt CGDrn CrdInt
        ------------ ---- --    ----- ----- -----
        0001075BDR31      0D        0    30    50

BadCGrpFormation: When  Serial       SSID LSS Reason     Activity
                  ----- ------------ ---- --- ---------- ----------
                  Last  0001075BDR31 ???? FF  InvaldSpTm RunInPrg
                  Prev  0001075BDR31 ???? FF  InvaldSpTm RunInPrg
                  First 0001075BDR31 ???? FF  InvaldSpTm RunInPrg

CurrentSCntlClock     CurrentGMTClock
--------------------  --------------------
26 Mar 2013 12:22:27  26 Mar 2013 12:23:03

CurrentSysplexClock  Drift LastGoodCGSysplexClock  SysplexName
-------------------- ---   --------------------    --------
26 Mar 2013 12:23:01   1   26 Mar 2013 12:23:01    LOCAL

*****************************************************************
```

GM Time Synch

- Available in DS8870 ucode release 7.1
- Enablement support available via APAR OA39733
- Exploitation in GDPS ACTIVE/ACTIVE with software based replication (QREP)

Non-Disruptive State Save Support (NDSS) / On Demand Dump (ODD)

Non-Disruptive State Save Support

- STATESAVE = Disk controller dump of internal structures for diagnostics
- Problem: Traditional statesave
  - Storage Facility Image (SFI) wide operation
  - Causes a warmstart of the SFI
    - Holds all I/O on the cluster for several (~ 5) seconds
  - Traditional statesaves typically avoided due impacts
    - Loss of FFDC
- Automatic statesave on error or via operator command
  - Statesave taken internally on certain XRC error conditions
    - Subject to XRC PARMLIB SCTRAP / SCTRAP2 settings
  - Operator command
    - F ANTAS000,SCTRAP ON
    - F ANTAS000,STATESAVE xxxx 0

Non-Disruptive State Save Support

- Solution: NDSS (ODD)
  - Captures dump of critical structures (statesave) without warmstart
  - XRC PARMLIB option: SCDUMPTYPE(NDSS)**
    - Now use NDSS for XRC error conditions
    - Only applies to XRC internally invoked statesaves
  - Operator command option added for NDSS**
    - F ANTAS000,STATESAVE xxxx 0 N
      - *N indicates NDSS*
      - *Without N, traditional statesave*

**Not subject to SCTRAP setting

  - New ANTRQST (API) request type STATESAVE / ANTTREXX command STATESAVE
    - GDPS exploiting at unplanned hyperswap
    - IOS exploiting via new SLIP command option
      - *in z/OS V2R1*
  - Allowed 1 every 5 minutes and 10 per 24 hour period

Non-Disruptive State Save Support

- ANTRQST and ANTTREXX new request types
  - ANTRQST:
    - ILK=ESSRVCS REQUEST=STATESAVE
  - ANTTREXX:
    - Using provided sample - ANTFREXX STATESAVE DEVN(xxxx)
  - Supports devices in alternate subchannel set using existing SUBCHSET keyword

- Only Required keyword: DEVN
- Optional Keywords
  - No validation
  - Allows specification of data to be stored in dump related to error or particular processing being done
  - SESSION
  - TYPE
  - LSS
  - CCA
  - FUNC
  - CALLER
  - SEQNO
  - DIAGRETC
  - DIAGREAS
  - TIME
  - TITLE
  - NDSS(YES | NO)
    - Allows for NDSS or traditional statesave

Non-Disruptive State Save Support

- NDSS = software/hardware term
  - also called On Demand Dump (ODD) by hardware
- Dump naming convention
  - cpssdump01_ODDdump
- Available in DS8870 ucode release 7.1
- Enablement support available via APAR OA38314
- Exploitation support in GDPS
- Exploitation support in SLIP command in z/OS V2.1

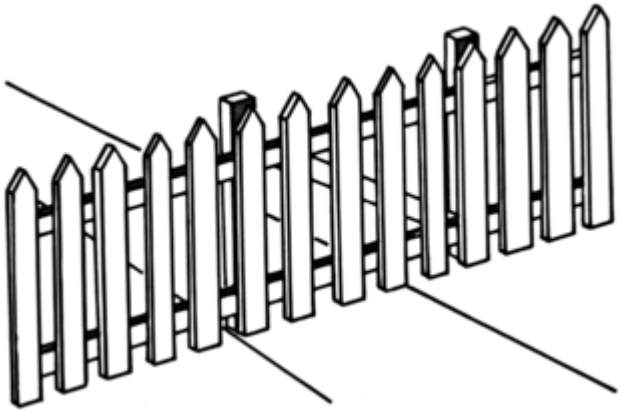

Soft Fence

Soft Fence

- Problem:
  - After a swap to alternate/secondary devices, not all devices at original production site likely have not failed, so some may still be accessible
    - Could be planned or unplanned
  - There have been instances where user has accidentally accessed an ‘old’ primary device, resulting in a data integrity issue
  - Accessing down level data = data integrity issue

Soft Fence

- Solution:
  - New function to prevent unintended IO to devices
  - Putting a device in a soft fenced state will prevent most subsequent I/Os
    - All reads/writes, most ‘active’ commands
    - Queries allowed
- Primary use case:
  - Isolate all of the old PPRC primary volumes after a HyperSwap or site failover
- Soft Fence state reset automatically on Failback

Soft Fence

- **New ANTRQST and ANTTREXX ILK=PPRC REQUEST=FENCE**
  - Required Keywords
    - DEVN
    - ACTION(FENCE | UNFENCE)
  - Key Optional keywords
    - SCOPE(DEV | LSS | MASK)
    - MASK
- ICKDSF
  - parameters added to the CONTROL command to clear a soft fenced state

Soft Fence

- PPRC Query changes
  - TSO CQUERY

```
ANTP0090I CQUERY FORMATTED LVL x
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                            (PRIMARY) (SECONDARY) *
*                                          SSID CCA LSS SSID CCA LSS*
*DEVICE   LEVEL     STATE    PATH STATUS SERIAL#    SERIAL#         *
*------ --------- ---------- ----------- --------- ---------       *
* BD00  PRIMARY.  SUSP(4)... ACTIVE..     1360 00  60  2BE0 00  E0 *
* SCH(0) ..............      ........... 0000000AZFP1  0000000BHYR1*
* ...........                                                      *
*THIS DEVICE IS CURRENTLY IN A SOFT FENCE STATE.                   *
* PATHS PFCA SFCA STATUS: DESCRIPTION                              *
* ----- --------- ------  ------------------                       *
* 1     0032 00A4   13    PATH ESTABLISHED...                      *
* ----  ----        00 NO PATH............                         *
* ----  ----        00 NO PATH............                         *
* ----  ----        00 NO PATH............                         *
* SUBSYSTEM    WWNN                       LIC LEVEL               *
* ----------- ----------------           -----------             *
* PRIMARY.... 5005076304FFC0A7           x.x.xxx.xxx             *
* SECONDARY.1 5005076304FFC2F4                                     *
```

Soft Fence

- ICKDSF

```
CONTROL UNIT(9000) SCOPE(DEV) SERIAL(ZA951)
CLEARFENCE
ICK00700I DEVICE INFORMATION FOR 9000 IS CURRENTLY
AS FOLLOWS:
                PHYSICAL DEVICE = 3390
           STORAGE CONTROLLER = 3990
           STORAGE CONTROL DESCRIPTOR = E9
           DEVICE DESCRIPTOR = 0E
           ADDITIONAL DEVICE INFORMATION = 4A001F3C
           TRKS/CYL = 15, # PRIMARY CYLS = 3339
ICK04035I DEVICE IS IN A SOFT FENCED STATE
ICK04000I DEVICE IS IN SIMPLEX STATE
ICK00091I 9000 NED=002107.900.IBM.75.0000000ZA951
ICK00001I FUNCTION COMPLETED, HIGHEST CONDITION CODE
WAS 0
```

Soft Fence

- **Exploited by GDPS and TPC-R**
- **Available via APARs**
  - **OA40510 (SDM)**
  - PM76232 **(ICKDSF)**
- **Requires DS8870 R7.1 microcode**

Query Host Access (QHA)

Query Host Access (QHA)

- Problem:
  - Unless specifically requested via keyword, FlashCopy and PPRC establishes fail if target/secondary is grouped (online)
  - No easy way to tell where a particular device may have path groups established or reserved

Query Host Access (QHA)

- Solution:
  - Provide ability to query a device and obtain any currently established path groups or reserves
  - Query Host Access to Volume provides:
    - CKD Path Group information
    - SCSI Reservation information
    - Capability to query of any device in a cluster
      - *using access device in same cluster*
  - GDPS will use the new ANTRQST capability to provide monitoring
  - Several z/OS components exploiting capability in future function ‘under the covers’

Query Host Access (QHA)

- ANTRQST
  - New ILK=ESSRVCS REQUEST=QHA
- ANTTREXX support
  - ANTFREXX QHA
- Query results data mapped by new macro
    - hlq.MACLIB(ANTQHA)
- Ability to protect command via new RACF Facility class
    - STGADMIN.ANT.ESS.QHA
- ICKDSF
  - ANALYZE command

Query Host Access (QHA)

- **Example TREXX call**
  - ANTFREXX QHA DEVN(F60) QRYSIZE(128) QRYINFO()
- Output
  - TREXX calls can give dump formatted output

```
ANTR8810I OK QHA SUCCESSFUL
+00000000 C1D5E3D8 C8C14040 01010000 00000000  00000080 0F600000 040F240F F0F0F0F0  *ANTQHA  .............-......0000*
+00000020 F0F0F0C2 C4D9F3F1 00000000 00010044  00000000 00000000 00000000 00200002  *000BDR31.......................*
+00000040 50880005 B9472827 CAECE0A6 00000000  00000000 00000000 0000FFF0 00000000  *&h.........w.............0....*
+00000060 50880112 66D02097 CA53DE70 00000000  00000000 00000000 00120936 00000000  *&h.....p.......................*
```

Query Host Access (QHA)

- New optional parameters have been added to the ANALYZE command to obtain host access information.
- `ANALYZE UNIT(9000) NODRIVE NOSCAN HOSTACCESS(ALL) DEVADDR(X'01',X'07')`

```
                HOST ACCESS INFORMATION LSS=01 CCA=07
+-----------------------------+----+--------+------+--------+---------+
|       PATH GROUP ID         |    |        |      |        | MAXIMUM |
|------+------+----+--------+  |    |        |      | DEVICE |NUMBER OF|
|      |      |CPU |CPU TIME|PATH|SYSPLEX |      |RESERVED|CYLINDERS|
|  ID  |SERIAL|TYPE|  STAMP |MODE|  NAME  |ONLINE|  TIME  |SUPPORTED|
+------+------+----+--------+----+--------+------+--------+---------+
|800002|B947  |2827|CA78BC17|S   |N/A     |NO    | ------ |120936   |
+------+------+----+--------+----+--------+------+--------+---------+
|800009|B947  |2827|CAC684B9|S   |PLEXM1  |NO    | ------ |120936   |
+------+------+----+--------+----+--------+------+--------+---------+
|800001|B947  |2827|CAC65DFD|S   |LOCAL   |NO    | ------ |120936   |
+------+------+----+--------+----+--------+------+--------+---------+
PATH MODE :
     S = SINGLE PATH
     M = MULTI PATH
SYSPLEX NAME :
     N/A = NOT AVAILABLE
```

Global Miror (GM) Pause With Consistency

GM Pause on Consistency Group Boundary

- Problem:
  - Typical Global Mirror Pause function:
    - Last consistency group (CG) is formed
      - *Further CG formation stops*
    - PPRC Global Copy relationships (in the GM session) continue transferring data to secondaries
      - *Global Copy secondaries considered 'dirty copies'*
      - *Consistent data resides on FlashCopy targets taken during CG formation*
    - D/R copies (on FlashCopy targets volumes) continue to age
    - Extra copy cannot be created using FlashCopy because the consistent data resides on a FlashCopy target

GM Pause on Consistency Group Boundary

- Problem continued:
  - In order to take a consistent set of practice or backup copies from a GM secondary:
    1. Quiesce applications
    2. Pause GM
       - *data drained to secondaries during CG formation*
    3. Suspend Global Copy while in a consistent state
    4. Resume applications
       - *Updates will result in out of sync bitmap tracking*
    5. FlashCopy Global Copy secondaries
    6. Resync the suspended Global Copy pairs
    7. Resume GM
  - Time to get practice/backup copies ~ 45 minutes for ~ 2k pairs

GM Pause on Consistency Group Boundary

- Solution:
  - GM Pause with consistency:
    1. Issue a pause with consistency to the GM session
       - *After CG formation, GM mechanism will suspend all global copy pairs with secondaries in consistent state*
    2. Take the point in time copy from Global Copy secondary
    3. Resume the Global Mirror session
       - *During RESUME, GM mechanism will unsuspend all global copy pairs suspended during PAUSE*
  - Time to get practice/backup copies ~ 15 seconds for ~ 2k pairs

GM Pause on Consistency Group Boundary

```
RQUERY DEVN(1F50) ACTION(GMLSTAT) SNBR(02)
RQUERY Output Devn(1F50) SCHSET(0) Action(GMLSTAT) Version(003)
SNbr GMLStat    GoodCg   Pct CrnBadCG TotBadCG LastGoodCGSCntlClock
--   ---------- -------- --- -------- -------- -------------------
02   CGPaused   00001B39 100 00000000 00000000 23 Jul 2012 19:38:14
.                                                                  .
Master: Serial       SSID LSS   CGInt CGDrn CrdInt
        ------------ ---- --    ----- ----- -----
        0001075TN141      06        0    30    50
.                                                                  .
CurrentSCntlClock     CurrentGMTClock
--------------------  --------------------
23 Jul 2012 19:41:40  23 Jul 2012 19:11:16
.                                                                  .
********************************************************************
```

GM Pause on Consistency Group Boundary

- New Status on TSO CQUERY

```
ANTP0090I CQUERY FORMATTED LVL 5
VOLUME REPORT
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                        SSID CCA LSS SSID CCA LSS*
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ---------- ---------    ---------  *
* 1F50  .........  SUSPEND.CG   ACTIVE..   3006 10 06   3007 10 07 *
* SCH(0) CRIT(NO).......      CGRPLB(NO). 0000000TN141 0000000TN141*
*       INCRES(NO).           AUTORESYNC(YES)                      *
* PATHS PFCA SFCA STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   1   0234 0304   13    PATH ESTABLISHED...                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*  SUBSYSTEM         WWNN                    LIC LEVEL            *
* -----------  ----------------           -----------             *
* PRIMARY....  500507630AFFC16D           7.6.31.66                *
* SECONDARY.1  500507630AFFC16D                                    *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 1F50. COMPLETION CODE: 00
```

GM Pause on Consistency Group Boundary

- DS8k RPQ in July 2012
- Generally available in DS8k 7.1
  - Available via APARs
  - OA42410 (SDM)
  - PM81469 (ICKDSF)
- Exploitation support provided by GDPS

Summary

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, *and only if* the disk and OS architects design solutions as a team

Key Points

1. IBM z/OS – DS8000 synergy increases infrastructure effectiveness and efficiency
2. IBM is uniquely positioned to provide improved communications between storage and operating system
3. IBM z/OS – DS8000 synergies delivers value to customers *today*

Innovation waits for no one. Do not forget that IBM is always staging for their next big enhancement and it has more stakes in this game as the provider of Mainframe servers, storage, and software than those who only supply the storage for the Mainframe. In addition, the importance of an obscure feature may be well known to IBM (internally) but not necessarily yet visible to or understood by competing vendors.

Before you make your next Mainframe storage procurement, take a close look at IBM's storage offerings. If you are looking to optimize your infrastructure, you'll be glad you did.

----- *The Mainframe and Its Storage — The Search for Optimized Infrastructure*, The Clipper Group, Report #TCG2010007LI

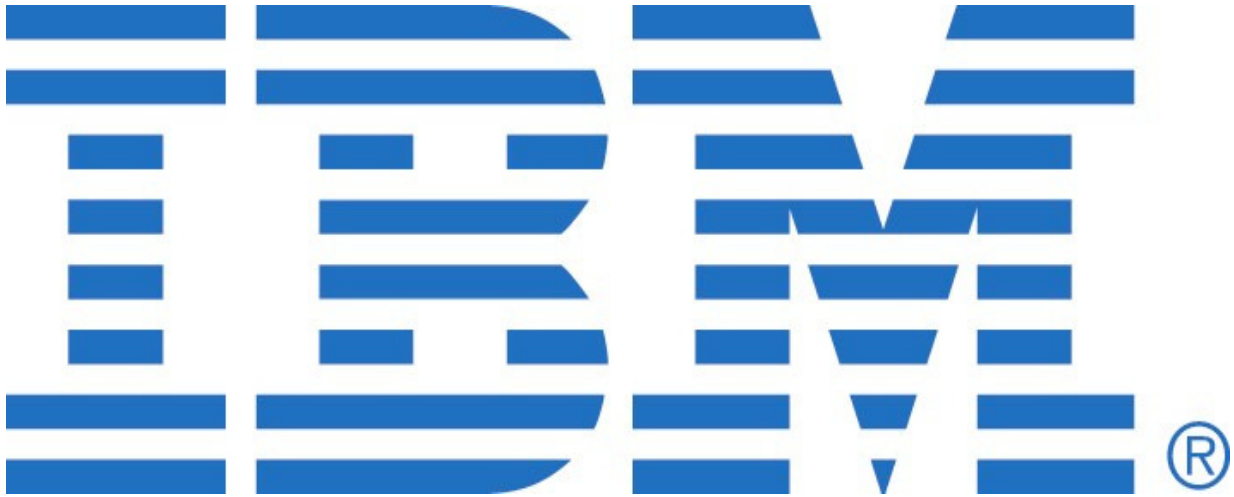

Additional Information, References, Disclaimers and Trademarks etc.

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
- Techdocs White Paper: 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.

Disclaimers

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice.

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

References in this document to IBM products, programs, or services does not imply that IBM intends to make such such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any on-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.


Thank You!
QUESTIONS?