

# DFSMS Update

Storage's latest and greatest

Thomas Reed  
IBM DFSMS Support  
[TReed@us.ibm.com](mailto:TReed@us.ibm.com)



# Please note

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

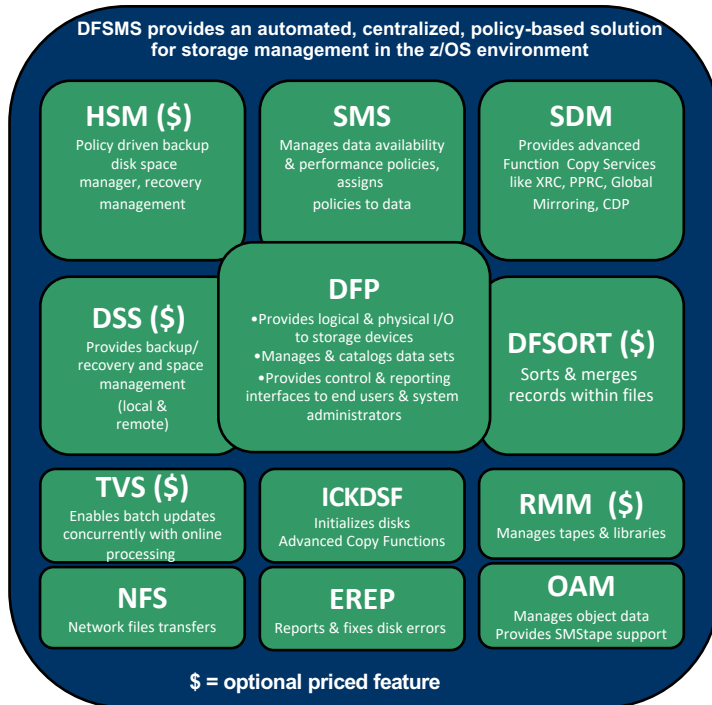
# DFSMS: Providing System Managed Storage on z/OS\*

*Storing, managing, protecting, and serving data on z Systems*



DFSMS is the standard methodology worldwide for managing enterprise data and storage on the z/OS platform

*Data is becoming the world's new natural resource, transforming industries and professions ...*



## DFSMS drives value as the data hub for System z:

- **Creates integrated solutions by exploiting new hardware features**
  - Enable better utilization and management of Storage HW
- **Maintains leadership in policy based storage management**
  - Improved storage administrator productivity and simplified management of the z/OS environment
- **Strengthens business resiliency by exploiting new opportunities and advancements in data protection solutions**
  - Point-in-time copy, fast replication, and continuous data mirroring functions while preserving consistency
- **Supports growing businesses and mission critical workloads by providing continuous availability, scalability/performance and flexibility of storage and data**
  - Increased data storage capacity and scalability to cope with explosive growth of data volumes and database sizes
  - High Availability with simpler, faster, and more reliable recovery operations
  - Ability to cope with increased security and compliance requirements
- **Enables cross platform data and storage**
  - Data availability at all levels of the storage hierarchy

# Session Objectives

- **Statements of Direction**
- **Access Methods**
  - Encryption
  - BSAM Support for zHyperWrite™
  - VSAM RLS AIX Lock Enhancements
  - CA Reclaim Enabled by Default
  - VSAM Performance Enhancements
  - VSAM zHyperLink Exploitation
- **HSM / DSS**
  - Encryption Enhancement
  - Transparent Cloud Tiering
  - DSS Increased Limits for Logical Backup
- **DFSORT**
  - Block Record Transfer
  - CPU Cache Performance Enhancements
  - Unicode Support
- **IDCAMS and SMS**
  - zFS Enhancements
- **OAM**
  - Multiple OAM Address Spaces
- **DADSM / CVAF**
  - VTOC Safe Update Interface
  - SMF 42-27 Record Enhancements
- **RMM**
  - Use of SMS Management Classes
  - WHILECATALOG Enhancements
  - UXTABLE Enhancements
- **Storage Synergy (backup)**
  - Cascading FlashCopy
  - Thin Provisioning
  - Read Only Volumes
  - zHyperLink™

# Statement of Directions

- **Statement of Direction, dated January 10, 2017**

- *IBM intends to deliver field upgradable support for zHyperLink™ on existing IBM System Storage DS8880. zHyperLink is a short distance mainframe attach link designed for up to 10x lower latency than High Performance FICON. zHyperLink is intended to speed DB2 for z/OS transaction processing and improve active log throughput.*
  - Lowering database transactional latency is critical to accommodate the use of new data sources and an increase in transaction volumes while enabling traditional and new mobile workloads to meet their SLA's.
  - Intended to make z/OS a more attractive platform for deploying new workloads, for growing heritage workloads with improved scalability, and leveraging reuse of existing storage assets.



- **\*\*NEW\*\* z/OS V2.3 GA Announcement, dated July 17, 2017**

- *IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express. zHyperLink Express is a short distance mainframe attach link designed for up to 10x lower latency than High Performance FICON.*

See [reference links](#) for recent blogs

For more check out: **Storage Performance with zHyperlink Session: DK**

# z/OS DFSMS Highlights

- **Storage Synergy**

- **Cascading FlashCopy**

- Thin Provisioning

- Read Only Volumes

- zHyperLink™

In the Backup  
Section




# Storage Synergy

## • Cascading FlashCopy

- **New SPE Enhancement:** Cascading FlashCopy enables the target of a FlashCopy to be the source of a subsequent FlashCopy (forward cascading) or the source of a FlashCopy to be the target of a subsequent FlashCopy (backward cascading)
- Use Cases:
  - Restore a Full Volume FlashCopy while maintaining other FlashCopies
  - Dataset FlashCopy combined with Full Volume FlashCopy
  - Improved dataset FlashCopy flexibility
    - Perform another FlashCopy immediately from a FlashCopy target
    - Create many dataset FlashCopy targets from a single gold source (don't wait for background copy)

? **Why it Matters:** Removes existing FlashCopy restrictions and provides improvements and simplification to numerous FlashCopy disaster recovery scenarios.



# z/OS DFSMS Highlights

## • Access Methods

- Encryption
- BSAM Support for zHyperWrite
- VSAM RLS AIX Lock Enhancements
- CA Reclaim Enabled by Default
- VSAM Performance Enhancements
- VSAM zHyperLink Exploitation





# Access Methods

- **Encryption**

- Various factors are driving the need for clients to adopt extensive use of encryption across their enterprises, including compliance mandates and the threat of data breaches.
- ***New enhancement:*** Provide users the ability to encrypt their data sets, using either SAF or SMS policies, ***without*** changing their application programs
  - Data set level granularity
  - Supports separation of access control for data set and encryption key label
  - Enabled through RACF and/or SMS policy
  - Provides mechanisms for audit readiness

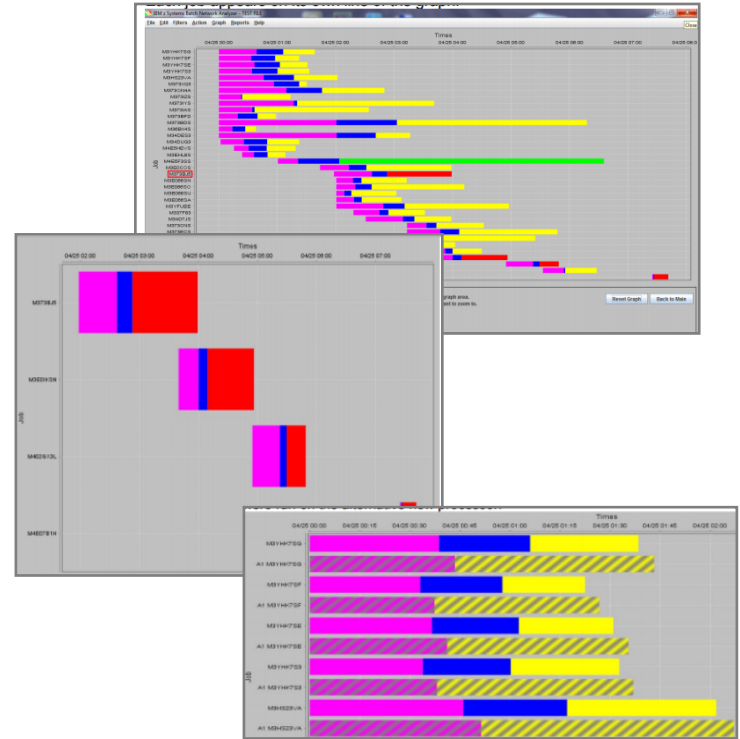


# z Systems Batch Network Analyzer (zBNA) Tool

## • zBNA Enhanced for Encryption

- zBNA is an as-is PC-based analysis tool designed to analyze batch windows
- Uses SMF workload data and generates graphical and text based reports
- Previously enhanced for zEDC to identify & evaluate BSAM / QSAM compression candidates
- **Enhanced for Encryption to help clients estimate the CPU impact of enabling encryption**
- zBNA V1.8.1 can be used by IBMers, BPs and clients
- zBNA tool:  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132>

- **NEW!** Support added for z/OS data set encryption and coupling facility encryption



# Access Methods

## • Encryption Requirements

- The minimum processor hardware is z196 or higher processor with CEX3 or later.
  - z196/z114 require CEX3 (feature 0864)
  - zEC12/zBC12 require CEX3 (feature 0864) or CEX4 (feature 0865)
  - z13 - CEX5 (feature 0890)
- DFSMS uses feature 3863, Central Processor Assist for Cryptographic Functions (CPACF) to encrypt and decrypt



# Access Methods

- Supported data set types include **extended format** (version 2 only) sequential BSAM and QSAM data sets and all types of **extended format** VSAM data sets
  - Covers DB2, IMS, zFS, Middleware, Logs, Batch, & ISV Solutions<sup>1</sup>
- The following data sets cannot be extended format:
  - Temporary data sets
  - SORTWK data sets
- Data set types that do not support encryption include
  - Basic and Large format sequential
  - PDS/PDSE
  - BDAM



<sup>1</sup> Any applications or middleware making use of VSAM, QSAM, BSAM access methods. Refer to individual product documentation to confirm support of z/OS data set encryption.

# Access Methods

- Encryption Requirements
- Operating system requirements:
  - z/OS V2.3
  - z/OS V2.2 with new function PTFs (MAIN APAR OA50569)
  - z/OS V2.1 with Coexistence PTFs (Main APAR OA50569)
    - Includes the ability to access existing encrypted data sets but not create new encrypted data sets on V2.1.
    - DFSMSdss will still be able to restore as an encrypted data set on a lower release.
    - DFSMSshm will be able to recall and recover an encrypted data set on a lower release.
  - ICSF
    - HCR77C0 or HCR77A0 through HCR77B1 with APAR OA50450
  - Review **fixcat** to obtain all the latest maintenance.



Category	Description	Keyword
IBM.Function.DataSetEncryption	Fixes to enable and support the z/OS Data Set Encryption function.	DSENCRYPT/K

# Access Methods


## • Encryption Restrictions

- System data sets (such as Catalogs, RLS SHCDS, HSM control data sets) must not be encrypted, unless otherwise specified.
  - Data sets used during IPL must not be encrypted
- Encrypted data sets only supported on 3390 device types
- Sequential (non-compressed) extended format data sets with a block size of less than 16 bytes cannot be encrypted
- DFSMSdss REBLOCK keyword is ignored on COPY and RESTORE functions.
  - DFSMSdss ADRREBLK installation exit will not be called for encrypted data sets.
- DFSMSdss does not support VALIDATE processing when backing up encrypted indexed VSAM data sets. VALIDATE will be ignored.



# Compression and Encryption

- **Encrypted data does not compress!**
- Any compression downstream from encryption will be ineffective
- Where possible compress first, and then encrypt
- Less data to encrypt means lower encryption costs
- **z/OS data set encryption**
- Recommended to compress first (generic only for VSAM) then encrypt
- Data sets will remain encrypted during HSM and DSS migration and backup processing
  - Currently cannot use the zEDC feature to compress VSAM encrypted datasets during HSM or DSS processing
- Data sets will remain encrypted during hardware based data replication services



# Access Methods

## • How to Request Encryption

- A sequential or VSAM data set would be defined as 'encrypted' when a **key label** is supplied on allocation of a new sequential or VSAM data set
- A key label supplied in any of the following (using order of precedence as follows):
  - RACF Data set profile
  - JCL, Dynamic Allocation, TSO Allocate
  - IDCAMS DEFINE
  - SMS Construct: Data Class
- **Note:** Once a data set is created with a key label, there will be no way to override or replace the key label.



? **Why it matters:** Clients who are required to protect customer data can leverage the z Systems hardware encryption for data at rest through existing policy management... *without application changes.*

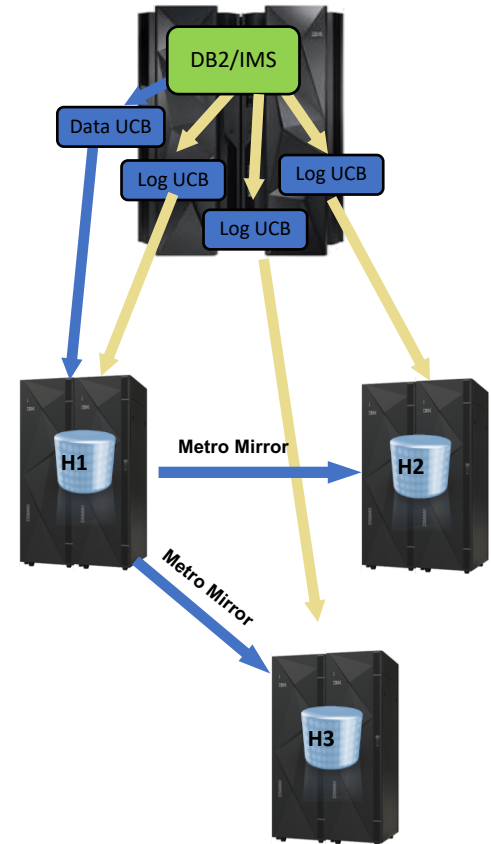


# Access Methods

- **BSAM Support for zHyperWrite**

- zHyperWrite storage technology was first introduced in December 2014 and combined DS8000 and z/OS enhancements.
- ***New SPE enhancement:*** APAR OA51385 introduces additional zHyperWrite exploitation in DFSMS BSAM
  - Allows authorized callers performing BSAM writes to an extended format data set to exploit zHyperWrite via a new DCB Extension (DCBE) option.
  - Write I/Os will be issued to primary and secondary volumes in parallel.
- IMS Release 15 is planned to exploit z High Performance FICON (zHPF™) and zHyperWrite to reduce latency and increase logging speed.
  - See the [Session reference links](#) for additional details

? **Why it matters:** Provides performance benefits and reduced latency when using Metro Mirror in a HyperSwap®-managed environment with GDPS® or IBM Copy Services Manager.



# VSAM RLS

## • AIX Lock Enhancements

- Today, an upgrade lock is held on the sphere to keep a VSAM RLS AIX upgrade set and the base cluster in sync.
  - In effect this makes the updates to the sphere single-threaded, which is unacceptably slow, especially with large records with many alternate keys.
- ***New enhancement:*** RLS replaces existing alternate index (AIX) upgrade lock with record locks and redo processing
  - Designed to keep the upgrade set and the base cluster in sync for update requests without forcing the updates to be single threaded.
    - No longer locks entire SPHERE during AIX upgrade
    - If an upgrade finds that the AIX CI has changed, the buffer is refreshed and activity restarted (REDO)
  - Pre-V2R3 releases need to have the toleration APAR OA48980 applied before running applications that have spheres shared between pre-V2R3 and V2R3 systems.



- **AIX Lock Enhancements**

- Performance measurements have show significant improvements in both CPU time and elapsed time.
  - Jobs doing inserts, erases, and updates saw *up to 48% CPU time improvements and up to 30% elapsed time improvements\**

? **Why it matters:** Allowing concurrent AIX updates potentially improves performance, most notably when updating numerous large records with many alternate keys.



*\* Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.*

# CA Reclaim

- Prior to V2.3, CA Reclaim can be enabled/disabled on a system level or on a data set level:
  - z/OS 1.12 and above CA Reclaim is **disabled** by default on a system level basis.
    - SYS1.PARMLIB(IGDSMSxx): CA\_RECLAIM (**NONE**/DATACLAS)
    - SETSMS CA\_RECLAIM() or SET SMS=xx
      - Dynamically set the CA Reclaim option on the system
  - z/OS 1.12 and above CA Reclaim is **enabled** by default for individual KSDSs.
    - DATACLAS parameter: CA\_Reclaim(**YES**/NO).
      - The value of YES or NO in the SMS DATACLASS is only used when a dataclass is assigned to a data set during define processing.
    - IDCAMS ALTER command: RECLAIMCA/NORECLAIMCA.
      - Alters data set level CA Reclaim attribute



- **CA Reclaim Enabled by Default**

- ***New enhancement:*** IBM recommends running with CA reclaim enabled
  - Provides improved DASD space usage and performance enhancements.
  - ***New*** healthcheck added to Health Checker to determine if VSAM CA Reclaim is enabled or not.
    - CHECK(IBMVSAM,VSAM\_CA\_RECLAIM)
    - Check is invoked during initialization and when CA Reclaim status changes
    - Health checker message idahc302E reports CA-reclaim is disabled on the system
  - Available on z/OS V2.1 and V2.2 with OA51002
- Refer to info apar II14640 for a Q&A type discussion on enabling CA Reclaim.
- **? Why it matters:** Customers who have allowed the system CA-reclaim setting in IGDSMSxx to default, and who desire no reclaims performed on the system, will need to set CA-reclaim to NONE in V2R3.



# Access Methods



- **VSAM Performance Enhancements**

- As customers consolidate workloads, and DB2 encourages customers to put one DB2 table per data set, the number of data sets needed goes up dramatically.
  - Many VSAM control blocks still reside in below the bar storage
  - Internally VSAM has large control block chains that must be processed during OPEN.
- ***New enhancement:*** Changes to VSAM to utilize less below the bar storage and enhancements in VSAM Open to more efficiently process data sets opened by DB2.
- ? **Why it matters:** Potential for improved performance and scalability for DB2 workloads by allowing the number of concurrent open data sets in a single address space to grow and by improving performance of data set open processing.

## • VSAM zHyperLink Exploitation

- ***New SPE SOD:*** When [zHyperLink](#) is enabled on an LPAR, VSAM (physical) read I/O will request zHyperLink support, if enabled.
  - Designed to be valid for all types of VSAM access (NSR,LSR,RLS)
  - Designed to be easy to enable
- See the following blog for more details:
  - [IBM DS8880 zHyperLinks™, the Gift That Keeps on Giving](#)



**?Why it matters:** Depending on system conditions\*, a zHyperlink request can have a much improved elapse time, allowing VSAM applications to run faster with less latency.

\* Disclaimer: *Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.*

# z/OS DFSMS Highlights

- **HSM / DSS**

- Encryption Enhancement
- Transparent Cloud Tiering
- DSS Increased Limits for Logical Backup
- Backup
  - [HSM Common Dump/Recover Queue](#)





## HSM Encryption Enhancement

- Prior to this enhancement, HSM didn't store in its inventory during backup or migration the encryption information for a data set other than an indication that the data set is encrypted.
  - New enhancement: The encryption information for a data set is recorded at the time of migration in a new record type added to the Migration Control Data Set (MCDS)
    - The key for this new record type will be '03'X and will be referred to as the MCDX record.
    - In addition, the encryption information for a data set at the time of backup is being recorded in a new extension of the MCC (Backup Control Data Set Backup Version Record) record type in the Backup Control Data Set (BCDS).
    - This encryption information will also be maintained in the DCOLLECT records created at the time a data is migrated or backed up and displayed by the HSM LIST command.
    - Available on z/OS V2.2 and z/OS V2.3 with OA52810.
- ? Why it Matters: Allows products like zSecure to create from the HSM control records an inventory of key labels needed to recover or recall data sets that have been encrypted by the access methods.




# Storage Synergy

- **IBM DS8K Transparent Cloud Tiering – Client Value**

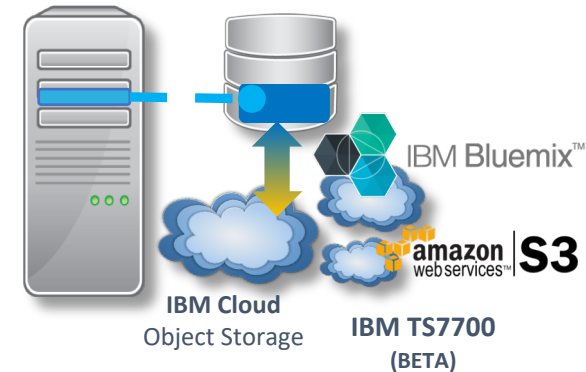
- TCT for DS8000 and DFSMSHsm saves z/OS [CPU utilization](#) by eliminating constraints that are tied to original Tape methodologies
- Direct data movement from DS8000 to cloud object storage without data going through the host
- Transparency via full integration with DFSMSHsm for migrate/recall of z/OS datasets

## Migrate with Tape

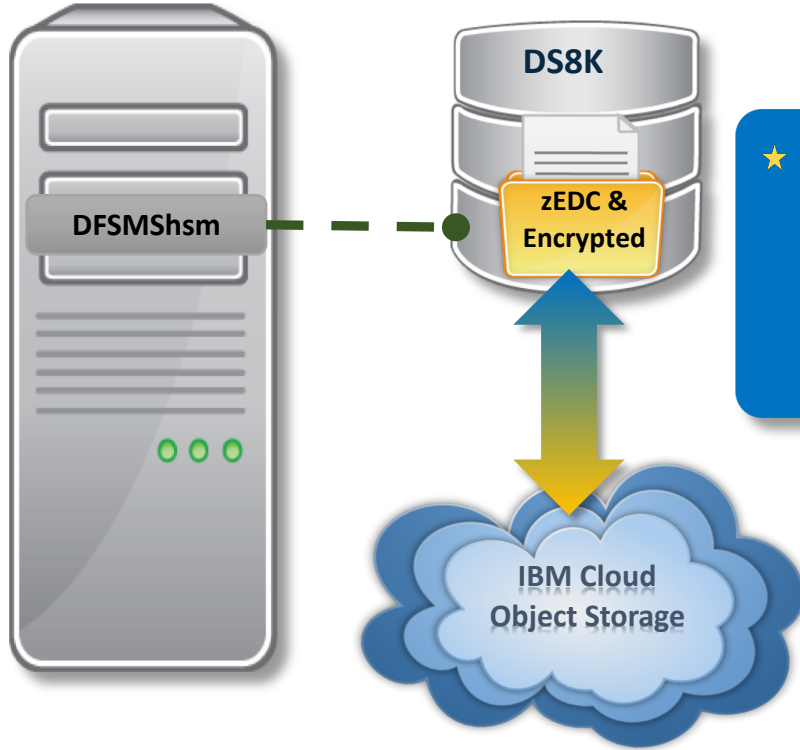


- ✓ **Reduce CPU Utilization**
  - ✗ 16K Blocksizes
  - ✗ Dual data movement
  - ✗ Recycle
  - ✗ Serial Access to Tape
- ✓ **Co-location**
- ✓ **HSM inventory**  
*(Eliminates OCDS)*

## Migrate with TCT & Cloud Storage



# Primary Use Case for Initial Delivery & Enhancements



★ Migration of large data sets for which a slow recall time is acceptable



**Compressed with zEDC**



**Data Set level Encryption**

- Ensures security across all storage tiers
- Enables low cost, off-premise Object Storage

*Large, historical data sets stored encrypted in Public Cloud*

# Storage Synergy



- **Transparent Cloud Tiering**

- ***New SPE enhancement 2Q17:*** Store and retrieve data in cloud object storage without the data going through the z/OS system
  - Does not require additional DS8000 hardware as it uses the existing ethernet ports on the DS8000 servers.
  - Object storage data can be provided on premises with IBM Cloud Object Storage or can use other private or public clouds which support the Swift object storage standard.
  - Exploited on z/OS V2.3 by DFSMSHsm and other software and provides the ability to use cloud storage for archive or backup data.
    - Management is done with new SMS constructs defined through ISMF panels.
    - Available on z/OS V2.1 with APAR OA48363 and on z/OS V2.2 with APAR OA50661.

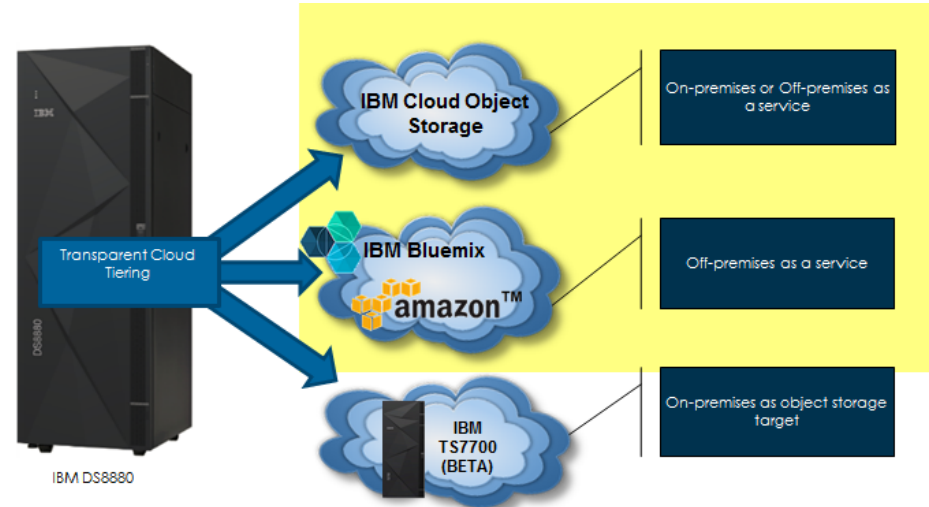
**? Why it Matters:** Allows for server-less data movement and potentially enables efficiencies in DFSMS which reduce CPU processing and free up MIPS for use by other applications and workloads.

# Storage Synergy

## • Transparent Cloud Tiering

- Initial release of TCT only supported Openstack Swift API to connect to object storage systems
- ***New SPE enhancement 3Q17:*** Support for S3 and IBM COS using S3 API

? **Why it Matters:** Expands server-less data movement support by extending options for Cloud Storage

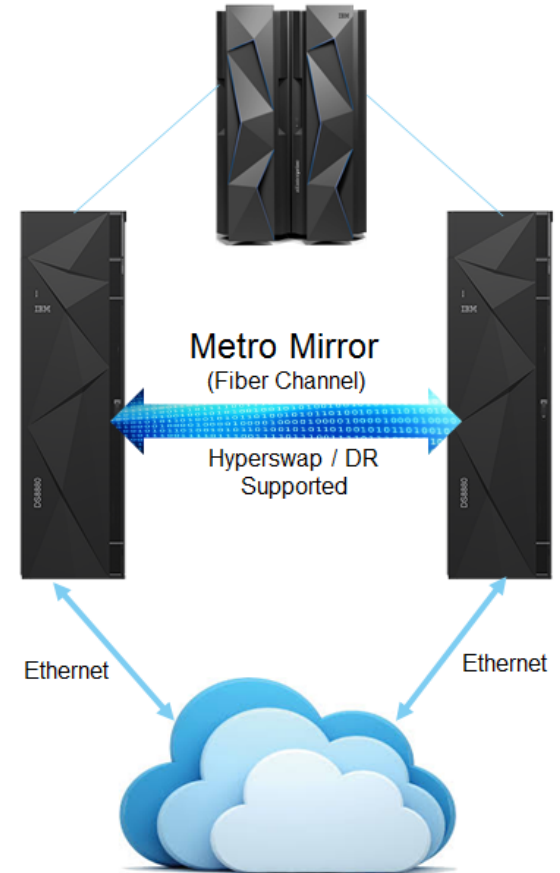


# Storage Synergy

## • Transparent Cloud Tiering

- Initial release of TCT restricted recall of data from cloud object storage to only Simplex volumes
- **New SPE enhancement 4Q17:** Migrate and recall of data to volumes in both Simplex and 2-Site Metro Mirror relationships
  - When data is recalled to a volume in a Metro Mirror relationship, it will automatically be synchronized to the MM Secondary
  - Supports Hyperswap (Planned/Unplanned) and PPRC Failover (DR)
    - Both DS8880s must be connected to the same cloud object storage
  - Flashcopy, Global Mirror, XRC continue to be restricted
  - Exploited on z/OS V2.3 and available on z/OS V2.1 and V2.2 with PTFs.

? **Why it Matters:** Expands server-less data movement support and extends the use case to include certain copy service functions



- **Transparent Cloud Tiering**

- Initial release of TCT only supported command based migration and recall using HSM
- ***New SPE enhancement 4Q17:*** HSM modified to allow data set migration copies to automatically be placed and retrieved from Cloud storage.
  - Define unique space management window for just Cloud migrations
  - Provide automated migration (Primary Space Management, Interval migration, and On-demand migration) based on SMS policies
    - Assign cloud policy to defined data groups via management class
  - Enables size-based tier selection (ie No Action, ML1, ML2, Cloud, Transition)

? **Why it Matters:** Allows HSM Primary Space Management to migrate unreferenced data sets to Cloud storage based on policies so users still have access to the data, but it does not take up space on system DASD or Tape.

- **DSS Increased Limits for Logical Backup**

- Today, users of DSS Logical Dataset Dump are limited to processing up to 131,070 data sets that pass their INCLUDE/EXCLUDE filter criteria.
  - If the user exceeds this limit, message ADR865E is issued indicating they must narrow the scope of their filter criteria, and the dump is not processed.
- ***New enhancement:*** DSS increases the limit of the number of data sets that can *potentially* reside on a logical data set backup to 2,147,483,392
  - Effectively increasing the INCLUDE/EXCLUDE filter capacity.
  - The ability to perform a logical data set RESTORE from a backup containing >131,070 data sets is supported on z/OS V2.1 and above with toleration PTFs for OA51382 installed.
    - Any attempt to restore from such a backup on a system that does not have the toleration PTFs installed will result in processing errors.

? **Why it Matters:** Potentially increase the backup list capacity to a point where it is no longer a processing constraint.





# z/OS DFSMS Highlights

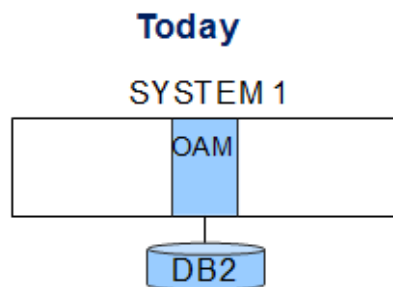
- **OAM**
  - Multiple OAM Address Spaces

- **Multiple OAM Address Spaces**

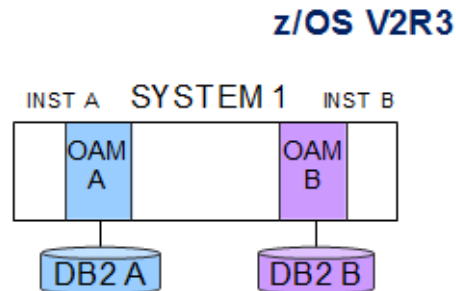

- ***New enhancement:*** Optional OAM configuration that allows a separate OAM tape library address space and up to **two** OAM object support address spaces
  - Each OAM object support address space can be connected to a unique DB2 subsystem, and each can optionally participate in different OAMplexes.
  - This will allow the ability for customers to deploy “production” and “test” capability for OAM or two separate “production” instances on the same system.
  - Application developers can exploit this functionality by directing OAM OSREQ application requests to a specific OAM (object support) instance.
    - Coexistence APARs are required.
  - Allows backward compatibility by retaining the existing Classic OAM configuration for customers that do not want to exploit the new Multiple functionality.

? **Why it Matters:** Provides OAM object users more flexible and scalable OAM configurations.

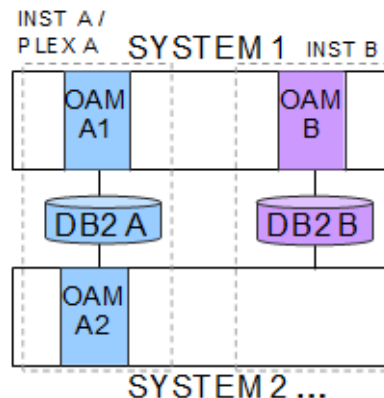
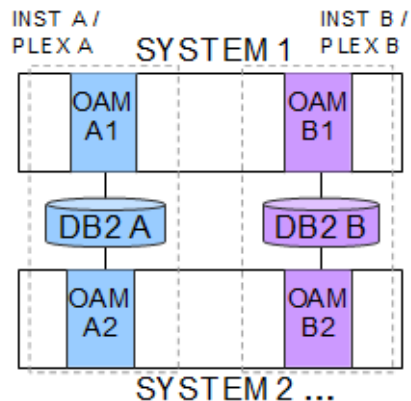
# OAM



- OR -



- OR -



- ❑ Maximum of two OAM Object 'instances' per system (1 subsystem and 1 address space per 'instance')
- ❑ An optional tape library instance may also be used on each system for a total of 3 OAM instances per system

# z/OS DFSMS Highlights

- **DFSORT**

- Block Record Transfer
- CPU Cache Performance Enhancements
- Support for Unicode

# DFSORT



- **Block Record Transfer**

- ***New SPE enhancement:*** DFSORT provides the ability to transfer blocks of records between DFSORT and the E15 and E35 user exits
  - Reduces the number of calls to the E15 and E35 user exits and the excessive transfer of records between user storage and DFSORT storage.
    - Can be used with COPY and SORT paths for both fixed and variable length records and coexists with transfer of single records already existing in E15/E35 exits.
    - Based on existing 64-bit support and requires use of 64-bit parameter list provided with z/OS V2.1.
    - Available on z/OS V2.1 and z/OS V2.2 with APAR PI47000. Exploited by DB2 R12.
      - [User Guide](http://www-01.ibm.com/support/docview.wss?rs=114&uid=isg3T7000719): <http://www-01.ibm.com/support/docview.wss?rs=114&uid=isg3T7000719>

? **Why it Matters:** Potential performance improvements due to the accelerated throughput of the records.

## • CPU Cache Performance Enhancements

- DFSORT was generating program code on the fly based on user parms for certain functions and wrote data directly into the instruction cache stream.
  - INCLUDE COND processing, OMIT COND processing, SORT statement with more than 5 keys
  - Can result in performance impacts (up to 7X CPU time\*).
- ***New SPE enhancement:*** DFSORT will avoid storing data into the instruction cache stream
  - The code generation programs changed to separate instruction areas and work areas used and avoid using the CPU instruction cache.
  - Available on z/OS V2.1 and V2.2 with APAR PI58848.
  - [White Paper](http://www-01.ibm.com/support/docview.wss?rs=114&uid=isg3T7000726): <http://www-01.ibm.com/support/docview.wss?rs=114&uid=isg3T7000726>



? **Why it Matters:** Potential provide performance improvements in both CPU and elapsed times for several DFSORT functions.

\* Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

- **Support for Unicode**

- ***New enhancement:*** With V2.3 DFSORT can now sort Unicode data in UTF8/UTF16/UTF32 encoding format
  - Enables users to SORT/MERGE Unicode Data with length of 1-450 Unicode characters.
    - The following are new DFSORT/ICETOOL reserved words which are no longer allowed as symbols: **UTF8, UTF16 and UTF32.**
    - In addition to the existing Abend codes, DFSORT will now display the reason codes from calling z/OS Unicode services.



? **Why it Matters:** Provide DFSORT users the added flexibility to SORT and MERGE Unicode data according to specific collation rules in the same manner that EBCDIC and ASCII data is today.

# z/OS DFSMS Highlights

- **IDCAMS and SMS**
  - zFS Enhancements



# IDCAMS and SMS

## • zFS Enhancements

- ***New enhancement:*** With V2.3 new IDCAMS keyword, **ZFS**, specifies that the cluster being defined is for linear data, and the linear data set is a Unix Z file system
  - When ZFS is specified, the linear data set is defined as extended addressable.
  - SMS provides a read only variable designation for ACS routines, **&RECOG, 'FS'**
    - Volume selection for DEFINE CLUSTER ZFS selects volumes indicated for zFS.
  - HSM provides new zFS file indication in Migration (MCD) and Backup (MCC) records.
    - zFS file indication is reflected in DCOLLECTs backup and migration records.
  - z/OS Unix will initialize the zFS file on the first mount.
  - Unused space in zFS files systems can be reclaimed using the new zFS administrative shrink command.

? **Why it Matters:** To aid in the migration of data from HFS file systems to zFS file systems, DFSMS provides additional zFS functionality that was previously only provided for HFS.

# z/OS DFSMS Highlights

- **DADSM / CVAF**
  - VTOC Safe Update Interface
  - SMF 42-27 Record Enhancements



- **VTOC Safe Update Interface**

- Today, there are many ways for callers to update VTOC DSCB records, but none provide any checking to insure the caller did not accidentally change fields that could corrupt the VTOC or cause the Index to be disabled.
- ***New enhancement:*** New CVAFDIR ACCESS=WRITE parameter added to indicate existing DSCB(s) are read and compared to the ones passed to insure essential fields are not being modified
  - **VALIDATE=(YES,NO)** allows the caller to update an existing format 1/8/9/3\* DSCB but not to modify **essential fields** in the record.
  - If any essential fields are modified, a [new CVSTAT code](#) will be passed back in the CVPL.
  - [Example and additional details](#) in the backup

? **Why it Matters:** Helps prevent accidental corruption to the VTOC and provides a safer way to update DSCB records in the VTOC.

# DADSM / CVAF – Essential Fields

- **The following fields are not allowed to be modified:**
- **Format 1/8 DSCB:**
  - DS1DSNAM Data set name.
  - DS1FMTID Format identifier (X'F1' or X'F8').
  - DS1NOEPV Number of extents on volume.
  - DS1EXNTS Three extent fields.
  - DS1PTRDS Pointer to first format 3 or format 9, or zero.
- **Format 9 DSCB:**
  - DS9KEYID Key identifier (X'09').
  - DS9SUBTY Subtype number for format 9 (currently always X'01').
  - DS9NUMF9 Number of format 9 DSCB's for this data set.
  - DS9FMTID Format identifier (X'F9').
  - DS9NUMF3 Number of format 3 pointers that follow.
  - DS9F3 Pointers to first to tenth format 3 DSCBs.
  - DS9PTRDS Pointer (CCHHR) to next format 9 DSCB, the first format 3 DSCB, or zero.
- **Format 3 DSCB:**
  - No fields in the format 3 DSCB are allowed to be modified

# DADSM / CVAF

- **SMF 42-27 Record Enhancements**

- z/OS V2.2 introduced a new SMF Record 42 subtype 27 for DASD VTOC Operations, which captures updates to the VTOC for IBM and Vendor-built channel programs.
- ***New enhancement: Enhanced SMF 42, subtype 27*** that provides all DSCBs that are affected by the write.
  - The following functions write the enhanced SMF record:
    - **DADSM: Create, Rename, Extend, Partial Release, Scratch**
    - **CVAFDIR WRITE (the CVCLID field can uniquely identify the writer)**
  - SMF42PSV (version number) field contains a value of 2, and all DSCBs affected by the activity are recorded.
    - There is a DSCB change section for the OLD and the NEW DSCBs.

? **Why it Matters:** Provides a VTOC update audit log to help diagnose problems when the VTOC is compromised - a comprehensive 'life of a dataset' footprint.



# New or enhanced SMF 42-27 fields

- **Changed fields** (changes in *italics*):
  - SMF4227R4 - Offset to *Old* DSCB section from start of record, including record descriptor word (RDW).
  - SMF4227R5 - Length of *Old* DSCB section
  - SMF4227R6 - Number of *Old* DSCB sections
  - SMF4227R7 - Offset to *New* DSCB section from start of record, including record descriptor word (RDW).
  - SMF4227R8 - Length of *New* DSCB section
  - SMF4227R9 - Number of *New* DSCB sections
- **New fields:**
  - SMF42RDSCB - Complete DSCB field added to address the entire DSCB instead of the key and the data fields having to be addressed separately.
  - SMF42RDSI - Data set indicators:

	Value	Meaning
SMF42RRSV	X'80'	Data set is erase on scratch

# z/OS DFSMS Highlights

- **RMM**

- Use of SMS Management Classes
- WHILECATALOG Enhancements
- UXTABLE Enhancements





- **Use of SMS Management Classes**

- Clients would like to manage both disk and tape data sets using the same set of policies; however, Management Classes used for disk data sets have limited influence on tape data.
  - Starting in V2.2, only RETPD and LASTREF could be assigned using the management class.
- ***New enhancement:*** Enhanced retention management by extending usage of SMS management classes.
  - A new SMS MC attribute, **Retention Method**, will allow the retention method (EXPDT or VRSEL), to be assigned to new tape volumes.
    - Several other SMS MC attributes are provided around volume set management as well as the ability to exclude tape datasets from VRSEL inventory management.

? **Why it Matters:** Helps simplify and consolidate RMM retention policies with SMS and potentially reduce inventory management overhead.



## – WHILECATALOG Enhancements

- In z/OS V2.2 RMM introduced retention by catalog to the EXPDT retention method.
  - A cataloged data set may prevent the volume from expiring on its expiration date if it has WHILECATALOG(ON).
  - Not all RMM subcommands allow to specify WHILECATALOG and expiration time.
- ***New enhancement:*** Add new **WHILECATALOG** and **Expiration Time** parameters to all applicable DFSMSrmm functions, including RMM reports.

? **Why it Matters:** Provides the user with a consistent interface and improved usability and flexibility.



- **UXTABLE Enhancements**

- Today, the UXTABLE is the only alternative to modifying SMS ACS routines and Management Classes if RMM users wanted to dynamically assign retention parameters to newly written tape data sets and volumes.
  - Existing UXTABLE is difficult to understand and manage and requires manual compilation and understanding of RMM exits.
  - Updates require the source code, and there is no way to check the contents of the currently loaded UXTABLE.
- ***New enhancement:*** Provide a new default Table which provides the same function as the UXTABLE but is simpler to maintain and manage
  - Defined in the EDGDEFxx PARMLIB member
  - Provided a new sample script, EDGRDEF, which can be used to convert an existing UXTABLE into the Default Table format.
- ? **Why it Matters:** Greatly simplifies how RMM users can assign default attributes based on DSNAMES, JOBNAME, or Expiration Date.

# Any Questions?

# Session reference links

- **zHyperLink**

- [IBM DS8880 zHyperLinks gives low latency access to storage](#)
- [DS8880 zHyperLinks™ Building a strong foundation](#)
- [IBM DS8880 zHyperLinks™, the Gift That Keeps on Giving](#)

- **zHyperWrite**

- [Improve your IMS Database Performance](#)
- IMS 15 [WADS Support for zHyperWrite](#)

- **DFSORT**

- Block Record Transfer [User Guide](#)
- CPU Cache Enhancement [Performance Whitepaper](#)

# Notices and disclaimers

- © 2018 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

# Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).
- .

Thank you!!



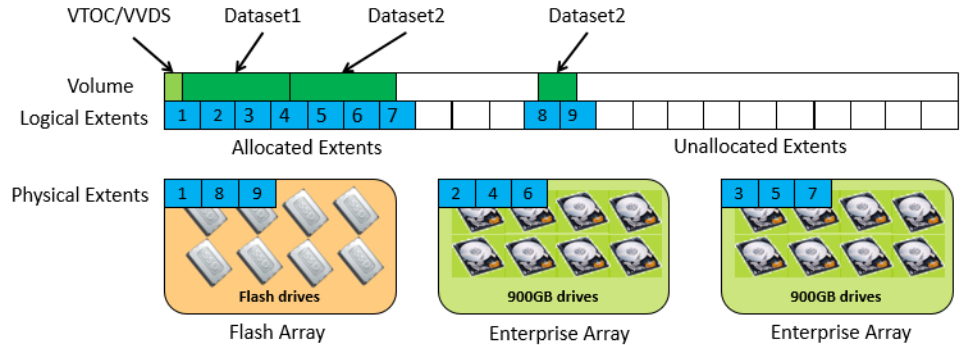
•Thomas Reed  
•z/OS DFSMS Support VSAM RLS/PDSE  
•TReed@us.ibm.com

# BACKUP



# Storage Synergy

- **Thin Provisioning**
- With DS8000 R8.1, CKD devices can be defined as thin provisioned / ESE devices
- **DFSMS Initial Support:**
  - Pool [utilization alerts](#) for extent pools
  - With z/OS 2.2, storage group utilization alerting, which can be helpful with thin provisioning
  - Enhanced [IDCAMS reports](#) to show thin provisioning statistics
  - DSS Move requests with FlashCopy pre-allocates extents to prevent data loss if the storage pool runs out of extents
  - Full volume space release with **ICKDSF INIT**.



Component	APAR
DSS	OA48707
SDM	OA48709
Device Support	OA48710, OA48723
IDCAMS	OA47811

# Storage Synergy

- Thin Provisioning

- *New SPE enhancement 4Q16:* Space Reclamation Tool

- DS8000 R8.2 provides for space release at an *extent level* for CKD volumes.
  - Exploited by DSS which performs the space release function.

- **SPACERel DDName (ddn) DYNam(volser,unit) STORGRP(groupname)**

- Determines free space extents on the designated ESE volumes and attempts to release the associated physical space to the extent pool.

- **Optional keywords**

- **CANCELERROR** specifies that the SPACEREL operation stops for the volume if a space release I/O related error occurs. Processing continues otherwise.
- **DEBUG** designates the level of progress information that DSS should provide for the SPACEREL operation.

- **Thin Provisioning**

- A new RACF FACILITY Class profile, **STGADMIN.ADR.SPACEREL**, is provided to protect the new command.
- Available on z/OS V2.1 and V2.2 with OA50675 and OA50677.

? **Why it Matters:** Provides a command to independently release space that is no longer needed to avoid out of space conditions, or to recover from out of space conditions after they occur.

# Storage Synergy

- **Thin Provisioning**

- Initial releases of thin provisioning allowed space release on a Metro Mirror Full Duplex Primary. However, space on the Secondary and a suspended Primary was not freed.
- ***New SPE enhancement:*** Expanded capability to release space for thinly provisioned volumes with copy services
  - Space will be released on the Secondary when released for a Metro Mirror Duplex Primary
  - Space can now be released on a suspended Primary. During resync, space will be released on the Secondary for any extents freed while the Primary was suspended.

**? Why it Matters:** Expands thin provisioning functionality and extends the use case to include certain copy service functions

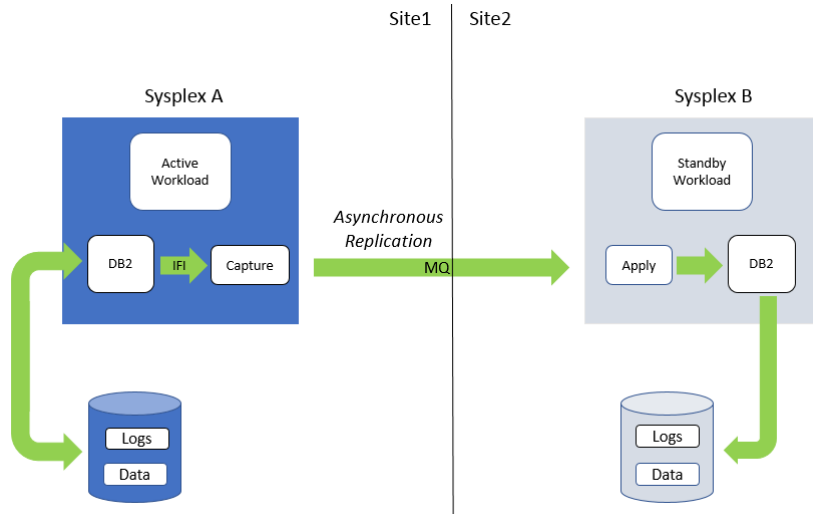
# Storage Synergy

## • Read Only Volumes

- GDPS Active-Active relies on asynchronous software replication to provide a long distance fast failover capability for HA and DR.
  - Potential for data loss if/when there is a disaster event.
- ***New SPE enhancement:*** Read Only Volumes
  - DS8000 enables z/OS to have read only access to a Metro Mirror secondary device to provide the basis for a Zero Data Loss option for GDPS/Active-Active.
  - The read-only attribute is known to z/OS, not the device.
    - The intent is to behave just like when z/OS runs as a guest on zVM with read-only mini-disks.
  - z/OS software support required:
    - z/OS V2R1 and V2R2 PTFs for APAR OA50068 will REQ all other required PTFs for enablement
    - Use HCD panels to configure PPRC Secondary devices as READ-ONLY
    - No PARMLIB keywords
    - New keywords for Allocation and Open to access data sets on read-only secondary devices

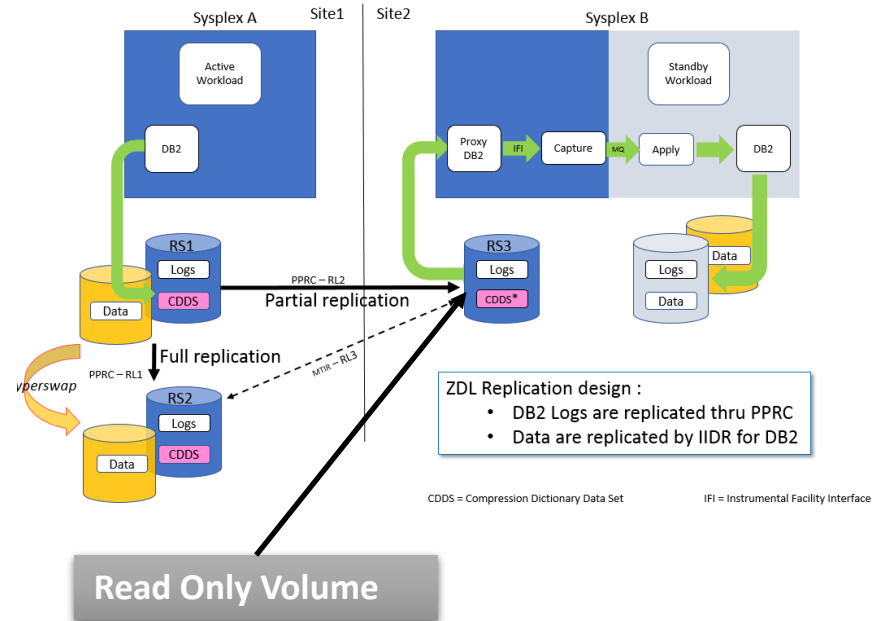
# Storage Synergy

- *GDPS Active / Active*



The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at IBM's sole discretion.

- *GDPS Active / Active with Zero Data Loss*



# Storage Synergy

- **Read Only Volumes**

- Application programs must specify new keywords to indicate that they understand the serialization provided by these devices.
    - Reads from PPRC secondary do not provide extent serialization
      - When a data set is accessed from a PPRC secondary, the application implicitly shares access with concurrent writes from the PPRC primary.
    - Loss of read integrity can occur if your records cross a track boundary
      - VSAM CI sizes that cross a track boundary are **not** supported.
  - No HyperSwap to system with read-only PPRC secondary devices.
- ? Why it Matters:** Helps to enable a zero data loss recovery in an GDPS Active/Active environment and provides improved TCO and HW efficiency by leveraging the secondary CU cache and IO capacity.

# Storage Synergy

- zHyperLink™

- New zHyperLink Technology accelerates transaction processing for z/OS
  - Lowering database transactional latency is critical to accommodate the use of new data sources and an increase in transaction volumes while enabling traditional and new mobile workloads to meet their SLA's.
- Critical components:
  - **DS8880**
    - zHyperWrite protocols built into zHyperLink protocols for acceleration of database logging with continuous availability
    - Investment protection for clients that already purchased the DS8880
    - New zHyperLinks compliment, but do not replace, FICON channels
  - **zNext**
    - New I/O paradigm transparent to client applications for extreme low latency I/O processing
  - **z/OS, DB2 and VSAM**
    - New approach to I/O processing. Allows reduction of I/O interrupts, context switching, L1/L2 cache disruption and reduced lock hold times typical in transaction processing work loads.

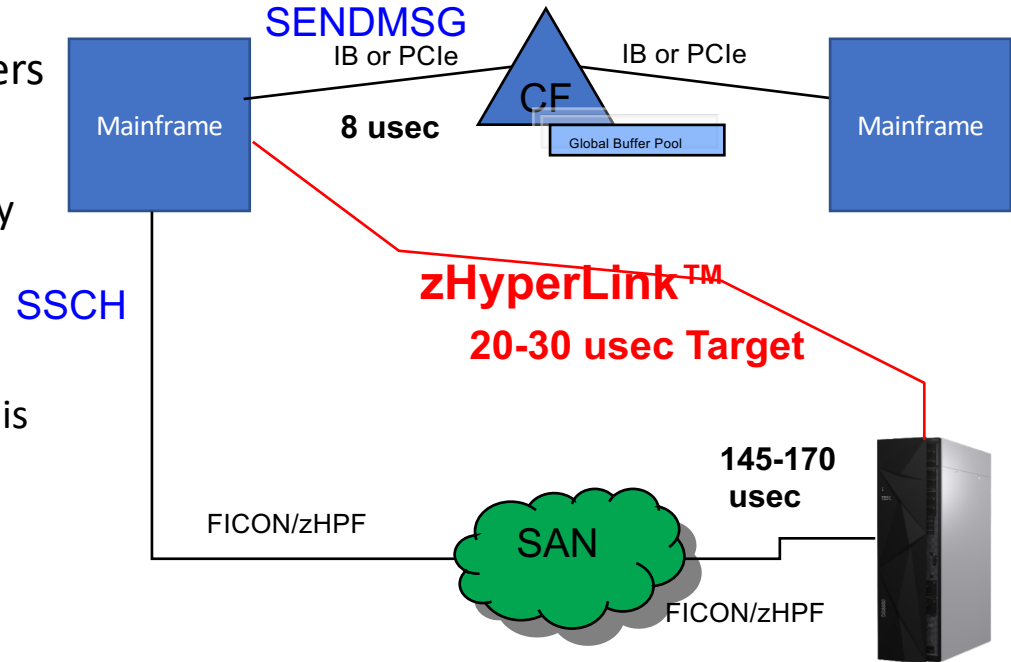
? **Why it Matters:** Provide a simple, transparent and consumable approach to enable extreme low latency I/O for the acceleration of transaction processing for DB2 on z/OS and VSAM applications, batch window reduction, and improved scalability while leveraging the reuse of existing storage assets.





# Storage Synergy - zHyperLink

- zHyperlink is fast enough so that the CPU can just wait for the data
- Point to point connection – 150 meters max distance from processor
- Connects z System’s Central Electronics Complexes (CECs) directly to the I/O Bay of the DS8880
- zHyperlink does not replace existing SSCH
  - Traditional I/O used when request is not eligible for zHyperlink or if zHyperlink request fails (e.g., read cache miss) and for device initialization
- Transparently Gives DB2 apps fundamentally better latency than apps on platforms without zHyperLink



# TCT - CPU Efficiency Estimator

- IBM has created a tool to estimate CPU savings
- HSM writes various statistics to SMF record specified by SETSYS SMF(smfid)
  - Recommended smfid is 240
- FSR records are written to smfid+1 (241)
  - FSRCPU records CPU time
  - Fields include dataset size and amount of data written
- ***NEW: Tool is now publically available:***
  - <ftp://public.dhe.ibm.com/eserver/zseries/zos/DFSMS/HSM/zTCT>

With a few days worth of SMF data, the estimator can determine:

1. Size of datasets to target for greatest cost savings
2. Estimated amount of CPU cycles saved by using Transparent Cloud Tiering



# Thin Provisioning Storage Pool Utilization Alerts

- IEA499E *dev,volser,epid,ssid,pcnt* EXTENT POOL CAPACITY THRESHOLD: AT *pcnt%* CAPACITY REMAINING
- IEA499E *dev,volser,epid,ssid,15%* EXTENT POOL CAPACITY WARNING: AT 15 % CAPACITY REMAINING
- 
- IEA499E *dev,volser,epid,ssid,pcnt* EXTENT POOL CAPACITY EXHAUSTED
- 



# Thin Provisioning IDCAMS LISTDATA output

```

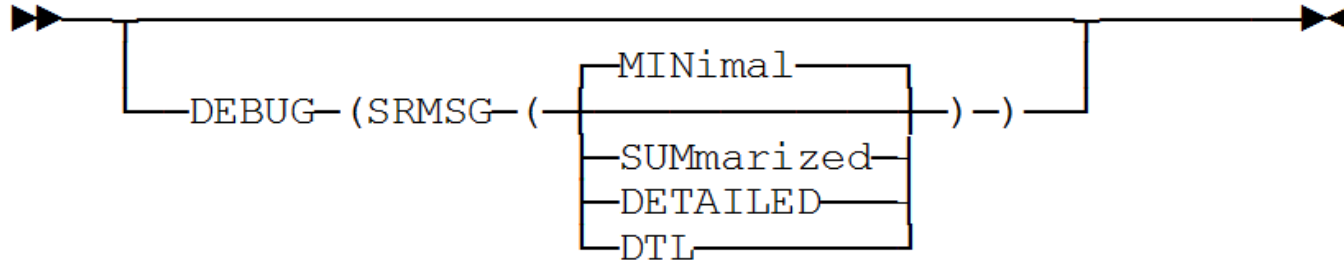
• LISTDATA VOLSPACE VOLUME(IN9029) UNIT(3390) ALL LEGEND
•
•           2107 STORAGE CONTROL
•           VOLUME SPACE REPORT
•           STORAGE FACILITY IMAGE ID 002107.961.IBM.75.0000000DKA61
•           SUBSYSTEM ID X'2400'
•           .....STATUS.....
•           CAPUSED      CAP      EXTENT
• DEVICE  VOLSER  (CYL)      (CYL)  POOL ID  SAM
• 900F    IN900F   3339      3339   0000    STD
• 902A    IN902A   2226      3339   0000    ESE
•           2107 STORAGE CONTROL
•           VOLUME SPACE REPORT
•           STORAGE FACILITY IMAGE ID 002107.961.IBM.75.0000000DKA61
•           SUBSYSTEM ID X'2403'
•           .....STATUS.....
•           CAPUSED      CAP      EXTENT
• DEVICE  VOLSER  (CYL)      (CYL)  POOL ID  SAM
• 9127    INF45   21         1113   0001    ESE
• 9129    INF49   21         3339   0001    ESE
• TOTAL NUMBER OF EXTENT SPACE EFFICIENT VOLUME(S) : 3
• TOTAL NUMBER OF STANDARD VOLUME(S) : 1
•

```



# SPACERel Optional Keywords

- **DEBUG**



- SRMSG subkeyword option designates the level of progress information that DFSMSdss should provide for SPACEREL operation.
- **SRMSG(MINIMAL)** specifies that DFSMSdss should issue a completion message and no additional information message will be provided. This is the default.
  - ADR006I (001)-STEND(01), 2016.299 15:12:30 EXECUTION BEGINS
  - ADR523I (001)-SRFP (01), SPACEREL FOR VOLUME SRE00A HAS COMPLETED

# SPACERel Optional Keywords

- **SRMSG(SUMMARIZED)** specifies that DFSMSdss is to issue ADR522I message when SPACEREL will begin to process a volume, and a completion message when the operation has completed. No free space extent information will be provided.
  - ADR006I (001)-STEND(01), 2016.299 15:12:30 EXECUTION BEGINS
  - ADR522I (001)-SRFP (01), THE FREE SPACE EXTENTS ON VOLUME SRE00A WILL BE PROCESSED BY SPACEREL
  - ADR523I (001)-SRFP (01), SPACEREL FOR VOLUME SRE00A HAS COMPLETED
  
- **SRMSG(DETAILED)** specifies that DFSMSdss is to issue ADR522I message when SPACEREL will begin to process a volume, and a completion message when the operation has completed. Free space extent ranges will be listed under ADR522I.
  - ADR006I (001)-STEND(01), 2016.299 15:12:30 EXECUTION BEGINS
  - ADR522I (001)-SRFP (01), THE FREE SPACE EXTENTS ON VOLUME SRE00A WILL BE PROCESSED BY SPACEREL
 

```

          SEQUENCE BEGIN C:H - END C:H
          00000001 00000001:1 00000001:5
          00000002 00000001:B 00000040:E
          00000003 00000056:0 00000063:E
          00000004 00000094:0 000000C7:E
          00000005 000000CE:0 00000458:E
          
```
  - ADR523I (001)-SRFP (01), SPACEREL FOR VOLUME SRE00A HAS COMPLETED

© Copyright IBM Corporation 2018. Technical

# SPACERel Optional Keywords

- **Examples:**

- SPACEREL DYNAM((SRCP03),      -  
                                  (SRCP04),      -  
                                  (SRCP05)) DEBUG(SRMSG(DTL))
  
- SPACEREL DYNAM(SRE00A) DEBUG(SRMSG(DTL))

- **Output for example #2:**

```
PAGE 0001      5695-DF175  DFSMSDSS V2R01.0 DATA SET SERVICES      2016.299 15:12
  SPACEREL DYNAM(SRE00A) DEBUG(SRMSG(DTL))
ADR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'SPACEREL '
  /*****/
  /*                STEP STEPT01                */
  /*****/
ADR109I (R/I)-RI01 (01), 2016.299 15:12:30 INITIAL SCAN OF USER CONTROL STATEMENTS COMPLETED
ADR016I (001)-PRIME(01), RACF LOGGING OPTION IN EFFECT FOR THIS TASK
ADR006I (001)-STEND(01), 2016.299 15:12:30 EXECUTION BEGINS
ADR522I (001)-SRFP (01), THE FOLLOWING FREE SPACE EXTENTS ON VOLUME SRE00A WILL BE PROCESSED BY SPACEREL
                       SEQUENCE BEGIN C:H - END C:H
                       00000001 00000001:1 00000001:5
                       00000002 00000001:B 00000040:E
                       00000003 00000056:0 00000063:E
                       00000004 00000094:0 000000C7:E
                       00000005 000000CE:0 00000458:E
ADR523I (001)-SRFP (01), SPACEREL FOR VOLUME SRE00A HAS COMPLETED
ADR006I (001)-STEND(02), 2016.299 15:12:30 EXECUTION ENDS
ADR013I (001)-CLTSK(01), 2016.299 15:12:30 TASK COMPLETED WITH RETURN CODE 0000
ADR012I (SCH)-DSSU (01), 2016.299 15:12:30 DFSMSDSS PROCESSING COMPLETE. HIGHEST RETURN CODE IS 0000
```



# DADSM / CVAF – VALIDATE Example

- Example of CVAF call with new VALIDATE parameter:

```
DIRWRITE DS 0H
```

```
    CVAFDIR ACCESS=WRITE,DEB=(R4),BUFLIST=BUFLHDR,MAPRCDS=YES, X
```

```
        DSN=DSNAME,MF=(E,CVAFDIR),VALIDATE=YES, X
```

```
        MULTIPLEDSCBS=YES,EADSCB=OK
```

```
CVAFDIR CVAFDIR MF=L
```

```
CVAFDIR MACRO PARM LIST
```





# DADSM / CVAF – Important Fields

- **Important CVPL, BFLE fields and new CVSTAT code:**
- CVPL fields (macro ICVAFPL)
  - CVCLID (4 character EBCDIC field, new with z/OS R2V2) – identifier provided by the caller of ACCESS=WRITE. This identifier is used in SMF 42 subtype 27 record field SMF42FACT (EBCDIC activity type).
    - CVFL4 DS XL1                      FOURTH FLAG BYTE
    - CV4VALID EQU X'03'              VALIDATE WAS SPECIFIED FOR CVAFDIR WRITE
- New CVSTAT code (passed back in the CVPL):
  - **88 - An essential field was attempted to be updated using CVAFDIR ACCESS=WRITE and VALIDATE=YES. No write occurred.**
- New flag byte in the BFLE (ICVAFBFL) to indicate which buffer had the invalid update
  - **BFLEVLER EQU X'40'      VALIDATION ERROR OCCURRED ON THIS BUFFER**




# Storage Synergy

## • Multi-Target PPRC (MT-PPRC)

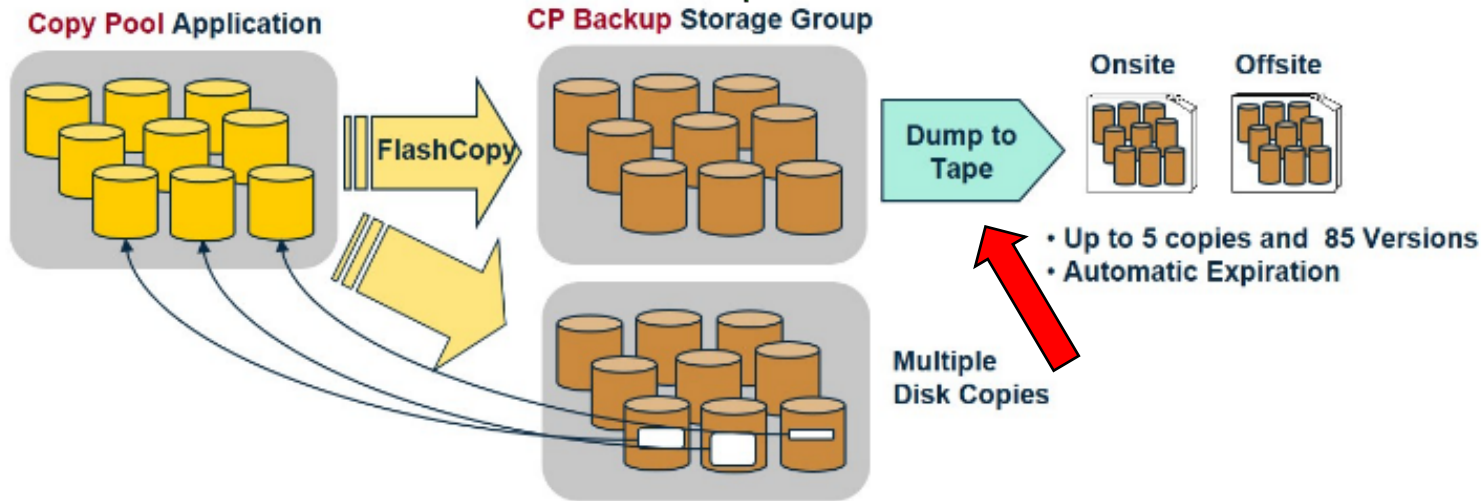
- Allow a single volume to be the source for more than one PPRC relationship.
- The capability of having multiple relationships allows for many different possible configurations.
- This function will also support an incremental resynchronization capability between the two target volumes in both planned and unplanned HyperSwap® situations.
- CSM and GDPS will provide support for MT-PPRC.

**? Why it Matters:** Having a second PPRC target and allowing for a quick incremental resync between the two targets prevents losing all Disaster Recover capability from a single outage.




# HSM

- HSM Common Dump / Recover Queue
- Dump commands (Command, Auto, Fast Replication) prior to V2.2 must be processed on the same HSM host that initiated the request.
  - z/OS V2.2 introduced the CDQ which supports distributed dump processing across multiple LPARs

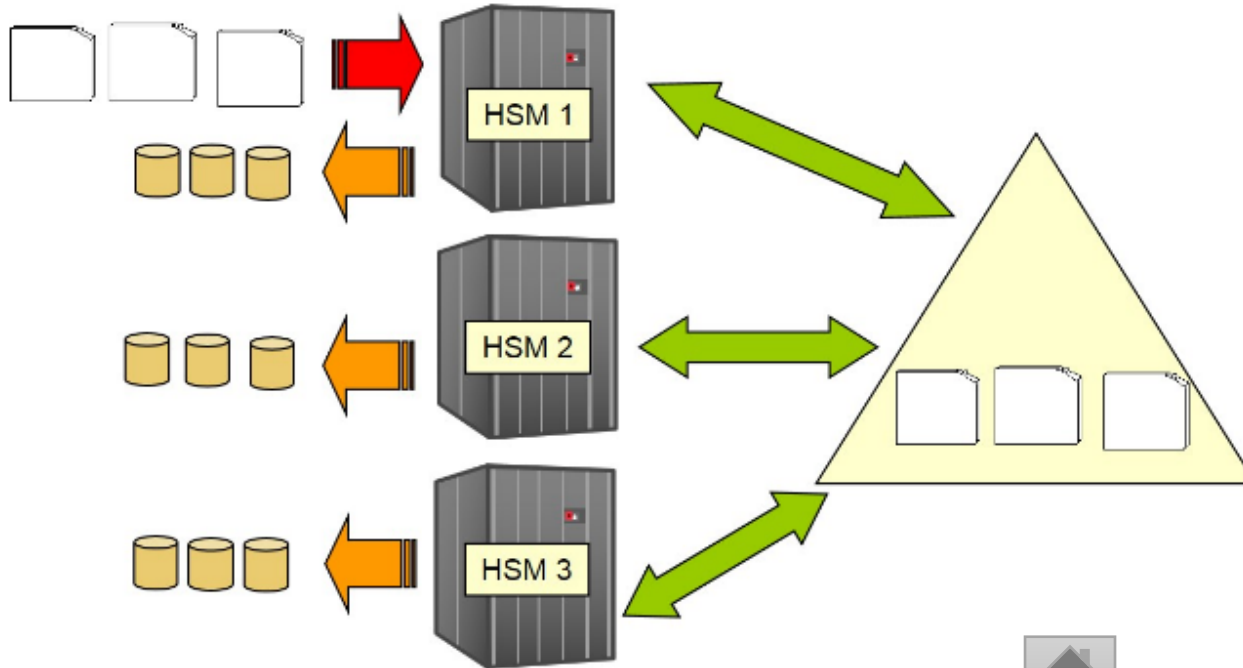


- HSM Common Dump Queue



- **HSM Common Recover Queue**
- zCDP for DB2 lacks the throughput capability to recover large DB2 environments from tape because the requests can only be performed from a single system.
  - The resources of a single host do not scale to the large copy pool sizes that need to be processed.
- ***New enhancement:*** Common Recover Queue (CRQ)
  - z/OS V2.2 introduced the CDQ which supports distributed dump processing across multiple LPARs. Function expanded to include recover processing.
  - Available on z/OS V2.2 with **OA47904**
- **? Why it Matters:** Improves overall throughput by distributing the workload across the sysplex instead of concentrating it on a single LPAR and helps reduce elapsed time when processing large DB2 copy pools.

- **\*\*NEW with OA47904\*\*** HSM Common Recover Queue



# Trademarks and Disclaimers

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](http://www.ibm.com/legal/copytrade.shtml):

CICS*	FICON*	Lotus*	System Storage
DB2*	FlashCopy*	MQSeries*	Tivoli*
DFSMS	GDDM*	Multiprise*	TotalStorage*
DFSMSdfp	GDPS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMSHsm	HiperSockets	Parallel Sysplex*	VM/ESA*
DFSMSrmm	HyperSwap	PR/SM	VSE/ESA
DFSORT	IBM*	QMF	VTAM*
DFSMS	IBM logo*	RACF*	WebSphere*
DS4000	ImagePlus*	Rational*	z/Architecture*
DS6000	IMS	RMF	z/OS*
DS8000	Intelligent Miner	System i	z/VM*
Enterprise Storage Server*	Language Environment*	System z	z/VSE
ESCON*		System z9	zSeries*
			zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies:

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

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

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

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

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

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

## NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

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

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

# Trademarks and Disclaimers (continued)

## NOTES:

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 are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.