#SHAREorg

SHARE

DFSMS™: Latest and Greatest

Barbara McDonald
IBM DFSMS Product Management
bawhite@us.ibm.com

August 12, 2013
Session 14141

SHARE in Boston

Disclaimer

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. 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 our sole discretion.

System z Social Media Channels

- Top Facebook pages related to System z:
  - IBM System z
  - IBM Academic Initiative System z
  - IBM Master the Mainframe Contest
  - IBM Destination z
  - Millennial Mainframer
  - IBM Smarter Computing
- Top LinkedIn groups related to System z:
  - System z Advocates
  - SAP on System z
  - IBM Mainframe- Unofficial Group
  - IBM System z Events
  - Mainframe Experts Network
  - System z Linux
  - Enterprise Systems
  - Mainframe Security Gurus
- Twitter profiles related to System z:
  - IBM System z
  - IBM System z Events
  - IBM DB2 on System z
  - Millennial Mainframer
  - Destination z
  - IBM Smarter Computing
- YouTube accounts related to System z:
  - IBM System z
  - Destination z
  - IBM Smarter Computing
- Top System z blogs to check out:
  - Mainframe Insights
  - Smarter Computing
  - Millennial Mainframer
  - Mainframe & Hybrid Computing
  - The Mainframe Blog
  - Mainframe Watch Belgium
  - Mainframe Update
  - Enterprise Systems Media Blog
  - Dancing Dinosaur
  - DB2 for z/OS
  - IBM Destination z
  - DB2utor

DFSMS™: Providing System Managed Storage on z/OS®

Storing, managing, protecting, and serving data on a zEnterprise System

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

DFSMS strategy addresses explosive growth and management of customer data

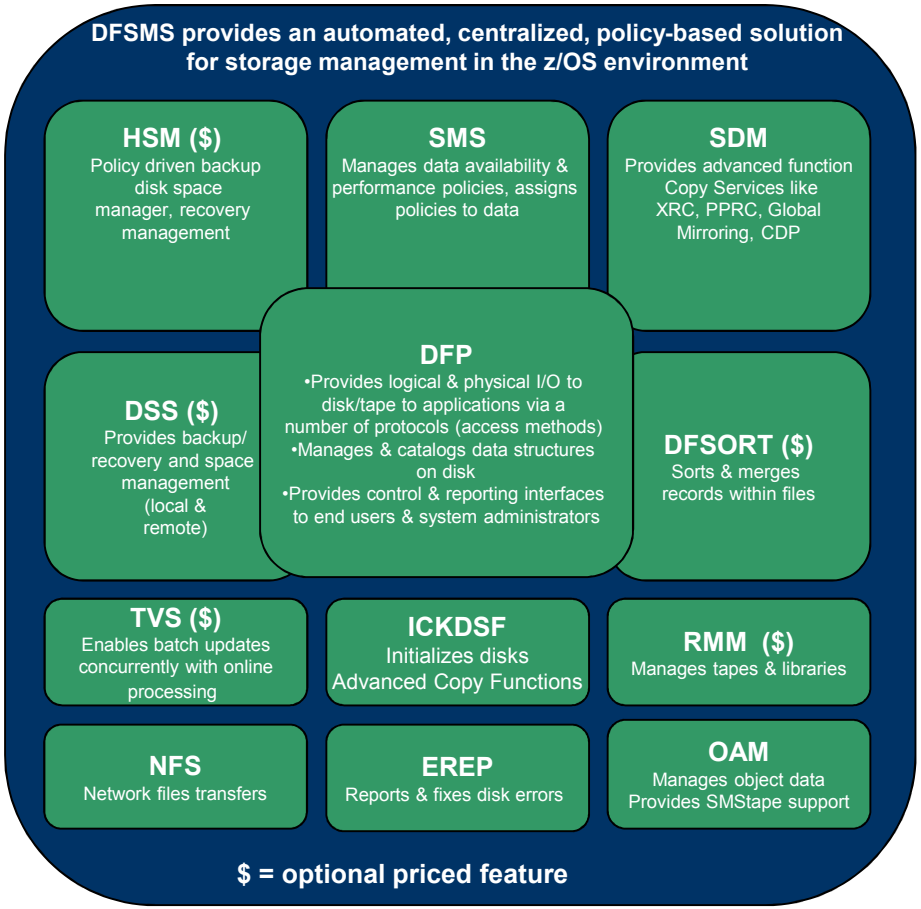

- **DFSMS drives value as the data hub for System z:**
  - **Creates integrated solutions by exploiting new hardware features**
    - Improved Security with exclusive media encryption capability, integration with z/OS Key Management
  - **Maintains leadership in policy based storage management**
    - 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

z/OS® DFSMS™ V2.1 Highlights
(September 2013)

- **DFSMShsm**
  - DFSMS Storage Tiers
  - DFSMShsm Tape Enhancements
  - DFSMShsm Fast Replication Enhancements
  - DFSMShsm RAS & Usability Enhancements
- **Catalog**
  - RLS Support for Catalog
  - Catalog Contention Detection Enhancements
  - Catalog Alias Enhancements
  - Catalog CSI Enhancements
  - Catalog Parmlib Member Enhancements
  - Catalog DFSMS GDG Enhancements
  - Catalog: RNLs HealthCheck
- **IDCAMS**
  - Larger BlockSize for LBI Support
  - IDCAMS Support for RLS
  - DELETE PDS/PDSE with Mask
  - ALTER NULLIFY Management Class
  - DIAGNOSE of GDGs
- **Access Methods**
  - VSAM RLS Directory Only Caching
  - VSAM RLS 64-bit Enhancements
  - VSAM RLS Dynamic Volume Count
  - VSAM SMB Enhancements
  - VSAM SHOWCB Enhancements
  - DFSMS Support for zHPF
  - SAM EF Support for FlashCopy
- **PDSE**
  - PDSE Version 2 Enhancements
  - PDSE Member Generations
  - PDSE Larger Member Size
  - GDG support for PDSEs
  - IEBCOPY Enhancements
- **SMS**
  - Provide Accurate Volume Space Statistics
  - Alter ACDS/COMMDS to SHAREOPTIONS(3 3)
  - New PARMLIB Option to Generate PDSE
  - SMS ACS Read-only Variable for EAVs

z/OS® DFSMS™ V2.1 Highlights
(September 2013)

- **DFSMSdfp**
  - OCE Partial Release Enhancements
  - OCE RAS Enhancements
  - XTIOT HealthCheck
- **DSS**
  - Reset with RESTORE
  - zFS Change Activity Support
- **SDM**
  - XRC Offline Volumes
- **OAM**
  - OAM Usability & Reliability Enhancements
- **DFSMSrmm**
  - DFSMSrmm SMS Mgmt Class for Tape
  - DFSMSrmm RAS Enhancements
- **DFSORT**
  - Dynamic Sort Enhancements
  - Functional Enhancements
  - 64-bit Enhancements
- **NFS**
  - NFS Server 64-bit Enhancements
  - NFS RPCSEC Performance Enhancements
  - NFS Server RPCBIND Enhancements

z/OS V2.1 Statement of Direction

- **DFSMS Exploitation of zEDC**
  - zEnterprise Data Compression (zEDC) for z/OS V2.1, running on zEC12 and zBC12 servers with the zEDC Express adapter, is designed to support a new data compression function designed for low-latency compression. Initially, z/OS is designed to allow you to specify that SMF data written to log streams be compressed.
    - *In addition, IBM intends to provide support for the BSAM and QSAM access methods. This function, planned to be made available by the end of the first quarter of 2014, is intended to help you save disk space, improve effective channel and network bandwidth without incurring significant CPU overhead, and improve the efficiency of cross-platform data exchange.*
    - *IBM also plans to provide support for DFSMSdss to exploit zEDC by the end of the third quarter 2014. This function is designed to be available for dumping and restoring data, and also when DFSMShsm uses DFSMSdss to move data. This is intended to provide efficient compression with lower CPU overheads than the processor- and software-based compression methods already available.*
  - There are several MVS Core Technologies sessions around the new zEC12 and zBC12. For more on zEDC and SW exploitation:
    - **Session 13705**: z/OS Software Support for IBM zEnterprise EC12 or zBC12 Server, **Wednesday 11AM**

z/OS DFSMS Highlights

- **DFSMShsm**
  - **V2.1**
    - DFSMS Storage Tiers
    - DFSMShsm Tape Enhancements
    - DFSMShsm Fast Replication Enhancements
    - DFSMShsm RAS & Usability Enhancements

Session 14101: A New Frontier in the Evolution of Space Management... Buckle In!, **Monday 1:30PM****

Session 14135: A New Frontier in the Evolution of Space Management...Blast Off!, **Wednesday 1:30PM****

Session 14092/14095: RMM and HSM Report Generator and Hands-on Lab, **Thursday 9:30AM and 11AM**

Session 14137: What's New in DFSMShsm, **Thursday 1:30PM**

Session 14108: What You Need to Know About the Way HSM Uses SMS, **Thursday 3PM**

Session 13772: The Life and Times of a Data Set, You Wouldn't Want Your Relatives Hanging Around, Why Your Data?, **Friday 8AM**

**** also multiple HW Storage Tiering sessions in the Storage project**

DFSMShsm

- **DFSMS Storage Tiers**
  - Today, DFSMS provides policy-based...
    - Data Creation
    - Backup / Recovery Management
    - Space Management
    - Expiration
  - No policy-based automation for moving data within the Primary Storage Hierarchy (Level 0)
  - No policy-based management of Active (open) data
  - ***New enhancement***: Automated, policy-based space management that moves SMS-managed data from tier to tier within the Primary (Level 0) Hierarchy.
    - Movement is referred to as a 'Class Transition'
    - Data remains in its original format and can be immediately accessed after the movement is complete
    - Policies implemented via the existing Class Transition policies and updated Management Class policies
    - Enhanced support for DB2, CICS and zFS data
      - *Open data temporarily closed to enable movement*

DFSMShsm

- **DFSMS Storage Tiers**
  - Data sets that would benefit from this solution:
    - Those currently not eligible for migration because they always need to be immediately accessible (ie recall delay is unacceptable)
      - *Data sets could be allocated on a particular class of storage and then later transitioned to a less expensive class of storage for permanent retention*
    - Data sets that are eligible for migration today, but there would be a benefit to keep them online for a longer period of time.
      - *Convert the migration of data sets to transition to a lower cost storage, then migrate directly to ML2*

? **Why it Matters:** Better align storage costs with changing business value; minimize the TCO for System z Data by actively managing data on the lowest cost storage that meets the business needs of the data.

DFSMShsm

DFSMShsm Tape Enhancements

- Today, although the migration function consists of multiple tasks, the actual migration of a data set is performed by a single task.
  - The migration of a data set may be broken down into three phases: setup (enqueue, verification), data movement, post processing (recatalog, scratch, dequeue).
  - In a single task, those three phases are performed sequentially.
- ***New enhancement:*** Allow a migration task to start multiple data set migration sub-tasks, where each sub-task manages the data set migration from beginning to the end as it does today.
  - By supporting multiple sub-tasks, the setup and post-processing phases of data set migration can be run concurrently across the migration sub-tasks.
  - The synchronization of data to tape can be deferred after a certain number of sub-tasks/data sets having completed data movement.

? **Why it Matters:** Potential marked reduction of elapsed time for any migration function managing multiple data sets such as interval migration, primary space management, ODM and volume migration. Greatest improved expected with moving large numbers of small data sets to tape.

DFSMShsm

- **DFSMShsm** Tape Enhancements
  - DFSMShsm migration and backup data sets can span up to a maximum of 40 tape volumes.
    - Since the size of a virtual tape is limited to 6GB, DFSMShsm cannot migrate or back up data sets larger than 600GB to virtual tape volumes (assuming 2.5:1 compaction).
  - ***New enhancement:*** Extend the maximum number of volumes that a migration or backup tape data set can span from 40 to 254 volumes.
    - Existing DFSMShsm architecture prevents the limit from being extended to the Allocation limit of 255 volumes.
    - Allow Recycle to process connected sets of up to 254 volumes.

? **Why it Matters:** Allows migration and backup of larger data sets, which is particularly useful when using the typically small tape volume sizes configured for virtual tape subsystems.

DFSMShsm

- **DFSMShsm** Fast Replication Enhancements
  - Fast replication is an HSM function that manages Point-in-Time copies
    - Combined with DB2 BACKUP SYSTEM, provides non-disruptive backup and recovery to any point in time for DB2 databases and subsystems (SAP) (ie continuous data protection).
    - Recovery at all levels from either disk or tape - entire copy pool, individual volumes and data sets.
  - **FlashCopy Consistency Group Support**
  - ***New enhancement:*** Add FlashCopy Consistency Group support
    - DFSMShsm will exploit the existing DFSMSdss capabilities.
    - Requirement from DB2 to use the FlashCopy consistency group feature to create a backup of the log copy pool with point-in-time data consistency so that a conditional restart of DB2 is not required.

? **Why it Matters:** By adding support for consistency groups, DFSMShsm will enable the DB2 BACKUP SYSTEM / RESTORE SYSTEM utilities to be used as a method to easily clone DB2 systems.

DFSMShsm

- **DFSMShsm** Fast Replication Enhancements
  - **Physical Data Set Recovery to Any Volume**
    - Previous enhancements in R8 and R11allowed data set recovery support for data sets that were currently cataloged on the same volumes and recovery for deleted and moved data by capturing catalog information.
    - Requirement to remove the restriction that data sets must be recovered back to the original volumes.
      - *This is a very undesirable restriction because there may not be space on the volumes at the time of the recovery.*
  - ***New enhancement:*** Add Fast Replication recovery to any volume
    - Enables DFSMS to select the volumes with the most highest preference weight, to which the data sets will be recovered.
    - Enables DB2 customers to discontinue the creation of object-level backup copies.
      - *These object-level backup copies must be created today in case the recovery cannot be performed back to the original volumes.*

? **Why it Matters:** Avoids of out-of-space conditions that can occur during recovery.

DFSMShsm

- **DFSMShsm** Fast Replication Enhancements
  - **Physical Data Set Recovery RENAME**
  - ***New enhancement:*** Enhance the FRRECOV DSNAME command to enable data sets to be renamed.
    - Renaming a data set during recovery enables users to be able to recover a broken data set to another name for analysis to determine what caused the breakage and to determine to which point in time the data set should be recovered, before replacing the production version of the data set.
    - ***Additional enhancement:*** Physical VSAM Data Set Restore RENAME support
      - *Currently, DFSMShsm does not support renaming of a VSAM data set using the RECOVER data set FROMDUMP command.*
      - *With the DFSMSdss support for physical data set RENAME described above, DFSMShsm will extend the RECOVER data set NEWNAME FROMDUMP support to VSAM data sets.*

? **Why it Matters:** Allows users to unload data from the renamed data set; also allow DB2 to create object-level backups from the system-level backup without impacting the availability of their objects to applications.

DFSMShsm

- **DFSMShsm RAS and Usability Enhancements**
  - **UCB Capture**
    - HSM's large customers face periodic 878-Abends in the DFSMShsm address space.
  - ***New enhancement:*** HSM will no longer capture UCBs into below the line storage.

  ? **Why it Matters:** Provides storage constraint relief by increasing the available storage below-the-line.

  - **SMSVSAM Server Termination Handling**
    - Today, when DFSMShsm accesses the CDSes in RLS mode and an SMSVSAM server error occurs, DFSMShsm simply takes a fatal abend and shuts down.
      - *If RESTART is not specified in the DFSMShsm startup procedure or if DFSMShsm is unable to restart within the allotted amount of time, the user is required to determine when the SMSVSAM server has initialized and manually restart all DFSMShsm hosts.*
      - *All DFSMShsm requests that were in progress at the time of the SMSVSAM server error are lost and must be reissued to complete.*
    - ***New enhancement:*** When an SMSVSAM server error occurs, DFSMShsm detects the error, and quiesces all CDS I/O activity.
      - *Once the SMSVSAM server initializes, DFSMShsm automatically closes and reopens the CDSes and resumes all requests waiting on CDS I/O operations.*

  ? **Why it Matters:** greatly improve the usability and robustness of DFSMShsm in regards to it's response to SMSVSAM server errors.

DFSMShsm

- **DFSMShsm RAS and Usability Enhancements**
  - **New Recycle Command**
    - Today, when a take-away occurs from a Recycle task, the Recycle task relinquishes the tape it is processing and terminates.
      - *This requires customers to potentially have to re-issue the Recycle request multiple times until all of the data on the original tape gets moved to new tapes..*
    - ***New enhancement:*** Automatically generate a new Recycle command for the same original tape when the original Recycle must terminate due to the take-away process.
    - ***Share Requirement MET:*** SSMVSS10004

? **Why it Matters:** Provides improved usability and recovery during Recycle processing.

DFSMShsm

- **DFSMShsm RAS and Usability Enhancements**
  - **Use Recycle instead of Tape Copy for failed alternate tape**
    - Today, when an alternate tape of a duplex copy pair fails to be created, DFSMShsm generates an automatic request to create a tape copy of the original.
      - *DFSMShsm tape copy function is single tasked and does not have a resume capability; almost impossible to create a copy of a tape (a recall request can take-away the tape from tape copy, causing tape copy to have to start from the beginning).*
    - ***New enhancement:*** Provide an option for customers to specify that DFSMShsm should automatically generate a Recycle request in response to an alternate tape failure, as opposed to generating a tape copy request.
    - ***Share Requirement MET:*** SSMVSS10003

? **Why it Matters:** Provides improved usability, recovery, and potentially performance for the Duplex Tape function.

z/OS DFSMS Highlights

- **Catalog**
  - **V2.1**
    - RLS Support for Catalog
    - Catalog Contention Detection Enhancements
    - Catalog Alias Enhancements
    - Catalog CSI Enhancements
    - Catalog Parmlib Member Enhancements
    - Catalog DFSMS GDG Enhancements
    - Catalog RNLs HealthCheck

Session 14145: What's New with DFSMS ICF Catalog and IDCAMS, **Monday 11AM**

Session 14142/14144: Unclog Your Systems with z/OS 2.1 - Something new and Exciting in Catalog Part I & II, **Tuesday 1:30 and 3PM**

Session 13771: Simplifying ICF Catalog Management with Tivoli Advanced Catalog Management, **Wednesday 9:30AM**

Catalog

- **RLS Support for Catalog**
  - Customers have raised a number of issues and requirements around Catalogs and Catalog processing:
    - **Performance**
      - *Contention on SYSIGGV2 bcsname when updating catalogs.*
      - *Limited catalog buffering and buffer invalidation.*
      - *Limited VSAM buffers/strings/storage.*
    - **Availability**
      - *Catalogs need to be split for contention issues.*
      - *Catalogs unavailable for splitting, recovering, and other maintenance activities.*
    - **Integrity**
      - *Catalogs damaged by utilities updating a catalog while catalog is opened by CAS.*
      - *Lack of sysplex control of closing and serializing catalogs.*
    - **Recovery**
      - *Long/error prone forward recovery procedures.*

Catalog

- **RLS Support for Catalog**
  - ***New enhancement:*** Exploit VSAM Record Level Sharing (RLS) for Catalogs.
    - Addresses customer concerns related to catalog performance, usability, availability, and recovery.
    - Replaces current BCS sharing and buffering protocols with more efficient functionality.
    - RLS will provide record level locking, and 64 bit local/global buffer pools.
      - *Performance is expected to improve by eliminating contention on current catalog serialization (SYSIGGV2 resource), reduced i/o activity (via larger local/global buffering), and improved cross system buffer invalidation (via XES/XCF Cross Invalidation).*
        - *SMSVSAM will hold SYSIGGV2 bcsname SHARE while a catalog is opened for RLS access (ensures catalog data integrity from programs relying on SYSIGGV2 to serialize the catalog)*
    - Enhance IDCAMS commands PRINT, REPRO, IMPORT and EXPORT to be able to open VSAM data sets using RLS.
    - New Catalog MODIFY commands to switch access between RLS and non-RLS.
    - New vendor interface to quiesce updates or suspend lock/catalogs to improve integrity.
    - New performance measurements at a catalog level.
    - Usability and availability are expected to improve since there is no longer a need to split catalogs in order to reduce contention and improve performance.
      - *Catalog size should no longer be a factor in maintaining and managing user catalogs.*
      - *Improve integrity and availability with new sysplex wide commands to control access to individual usercatalogs within a parallel sysplex.*

Catalog

- **RLS Support for Catalog**
  - ***Additional enhancements (applies to ALL catalogs RLS and non-RLS):***
    - *Preserve user catalog connector alias entries when you temporarily delete a user catalog so they need not be redefined when the catalog is reallocated, and prevent new catalog entries using those aliases from being defined until the new catalog is available.*
    - *Suspend / Resume catalog requests for a specified catalog across a sysplex to allow users to minimize application disruption during catalog maintenance.*
    - *DSS DUMP processing will invoke a Quiesce for Copy (QUICOPY/QUICEND) for catalogs opened for RLS.*
      - *The QUICOPY will suspend update requests only in order to obtain a sharp copy of the catalog. The QUICEND will resume update requests.*
      - *For nonRLS catalogs, DSS will use existing serialization (SYSIGGV2).*
    - *DSS RESTORE processing has new LOCK / SUSPEND options*
      - *LOCK will invoke a sysplex wide close of the catalog and lock the catalog (failing new unauthorized requests), if the catalog is not already locked or suspended.*
      - *SUSPEND will invoke a sysplex wide close of the catalog and suspend new unauthorized requests in the client space if the catalog is not already locked or suspended.*

? **Why it Matters:** Addresses several critical customer requirements around Batch Constraint, Mean Time to Recover (MTTR), and Catalogs as a Single Point of Failure (SPOF).

Catalog

- **Catalog Contention Detection Enhancements**
  - CAS (Catalog Address Space) Contention Management was introduced in Z/OS R1.12
    - CAS Contention Management monitors Catalog Address space for possible contention for resources among Catalog tasks.
    - It was designed initially to ONLY to detect SYSZTIOT contention with in Catalog and introduced a new and an altered catalog modify command.
      - *The only action taken when the wait threshold was crossed was notification to the console and a one time symrec to the logrec.*
  - ***New enhancement:*** Expand the resources being monitored to include:
    - SYSZVVDS – Serialization on the VVDS dataset of a volume.
    - SYSIGGV2 – provides an essential mechanism to facilitate cross system sharing of catalogs.
    - ALLOCLCK -- an internal CAS lock which protects allocations, de-allocations, opens and closes.
    - A new action, REDRIVE, can be associated with a resource and triggered when the wait threshold is breached.
    - Contention wait-time and actions per resource can be set in the Catalog parmlib.

? **Why it Matters:** More efficient use of storage resources; better diagnostics to determine the cause of serialization contention problems that impact CAS.

Catalog

- **Catalog Alias Enhancements**
  - ***New enhancement:*** Improvements in the processing catalog aliases:
    - Save creation date for Catalog aliases; listed by IDCAMS
      - *Clients are attempting to cleanup obsolete HLQs (High Level Qualifiers). If an alias has no associated datasets, there is no easy way to determine whether this is a new alias and no data sets have been created or this is an obsolete alias that should be deleted.*
      - ***SHARE Requirement Met:*** **SSMVSE10018**
    - Add a check when deleting a catalog entry that has an associated alias to verify that the alias is only related to the entry being deleted, before deleting the alias record.
    - Ensure requests are oriented to the correct catalog when the data set aliases in the master catalog specify a different high-level qualifier.

? **Why it Matters:** Ease of use and improved catalog alias processing.

Catalog

- **Catalog CSI Enhancements**
  - Today, some data set information can only be accessed by group field names.
  - ***New enhancement:*** New field names will be externalized for clients to retrieve the requested information:
    - STRNO: Number of concurrent requests
    - BUFND: Number of buffers requested for Data component
    - BUFNI: Number of buffers requested for Index component
    - INDXLVLS: Number of Index Levels SEQ-SET-RBA
    - HILVLRBA: RBA of High Level Index Record
    - ASSOC: A repeating list of catalog records associated with this entry.
    - ASSOCSYB: Indicates if the entry is a symbolic-relate
    - TRACKS: Total tracks per volume.

? **Why it Matters:** Improved usability for CSI requests.

Catalog

- **Catalog Parmlib Member Enhancements**
  - The Catalog PARMLIB member, IGGCATxx, was introduced in R13 and will allow:
    - The specification of most Catalog parameters that now can only be specified in SYSCATLG/LOADxx and/or by the Modify command
    - Parameters in the catalog PARMLIB member override equivalent parameters in SYSCATLG or LOADxx.
  - The options take effect on the next IPL or catalog address space restart
  - ***New enhancement:*** Include additional parameters for the remaining Modify Catalog command and for other specifications not required early during IPL processing.

? **Why it Matters:** These enhancements are intended to make it easier to specify options for catalog processing.

Catalog

- **Catalog DFSMS GDG Enhancements**
  - ***New enhancement:*** Allows users to specify that all the members of a generation data group (GDG) be returned in order from oldest to newest when the generation data set (GDS) name is specified without a generation number.
    - New GDGORDER JCL DD statement keyword to specify that you get the generation datasets oldest generation first to newest or the reverse.
    - System default remains newest-to-oldest

? **Why it Matters:** Allows all the members of a GDG to be processed in chronological order without being sorted or concatenated.

Catalog

- **Catalog RNLs HealthCheck**
  - IBM makes specific recommendations about what to specify in global resource serialization resource name lists (GRS RNLs) to prevent catalog-related deadlocks when using shared volumes and catalogs.
  - ***New enhancement:*** New GRS RNLs Health Check to help prevent lockouts due to shared volumes:
    - The check will perform the system check and indicate when GRS RNLs do not match IBM recommendations for SYSIGGV2, SYSZVVDS and SYSVTOC.
    - IBM now recommends that ALL reserves be converted; the SYSIGGV2, SYSZVVDS and SYSVTOC reserves should all be converted, unless the user shares DASD outside the sysplex.
    - Deadlocks can be prevented by always converting the SYSIGGV2, SYSZVVDS and SYSVTOC reserves to SYSTEMS ENQUEUEs using GRS or an equivalent product.

? **Why it Matters:** Automatic checking (via Healthchecker) of IBM best practice recommendations.

z/OS DFSMS Highlights

- **IDCAMS**
  - **V2.1**
    - Larger BlockSize for LBI Support
    - IDCAMS Support for RLS
    - DELETE PDS/PDSE with Mask
    - ALTER NULLIFY Management Class
    - DIAGNOSE of GDGs

Session 14145: What's New with DFSMS ICF Catalog and IDCAMS, **Monday 11AM**

IDCAMS

- **Larger BlockSize for LBI Support**
  - The large block interface (LBI) was introduced ten years ago.
    - It uses less tape storage and transfers data faster
  - IDCAMS has not supported LBI, and this restriction can stop the use of the LBI feature
    - IDCAMS REPRO and PRINT utilities fails with error message IDC3300I followed by IDC3321I and a return code of 12.
    - The max block size for IDCAMS REPRO and PRINT is 32760.
  - ***New enhancement:*** IDCAMS REPRO and PRINT will support a block size up to the access method limit, which currently is 256 KB.
    - The block size is still limited to 32 KB when the data set does not support LBI, such as with a unit record device or TSO terminal.
    - SYSIN and SYSPRINT do not support LBI.

? **Why it Matters:** Allow processing of data sets created using LBI.

IDCAMS

- **IDCAMS Support for RLS**
  - Customer requirement to copy, print and backup VSAM data sets while sharing the VSAM data set with other applications.
  - ***New enhancement:*** Enhance IDCAMS commands PRINT, REPRO, IMPORT and EXPORT to be able to open VSAM data sets using RLS.
    - A new optional keyword RLSSOURCE ({NO|YES| QUIESCE}) and/or RLSTARGET ({NO|YES|QUIESCE}) will be implemented for the PRINT, REPRO, IMPORT and EXPORT commands.
  - ***SHARE Requirement Met:*** **SSMVSS01007**

? **Why it Matters:** Provides applications the ability to read or print records from a data set being used in RLS mode.

IDCAMS

- **DELETE PDS/PDSE with Mask**
  - In R12, IDCAMS provided a DELETE option to delete all members of a partitioned data set in a single operation.
  - ***New enhancement:*** Enhance the IDCAMS DELETE command to be more flexible in performing the deletion of the members in a partitioned data set (PDS/PDSE).
    - Allows a mask for member names to be specified in a DELETE command for PDS/PDSEs.
    - A mask for a member name can contain an asterisk (*) or percent sign (%).
      - *Asterisk means 0 or more characters*
      - *% means 1 and only 1 character*
    - Double asterisk (**) still means delete all members in the PDS/PDSE.

? **Why it Matters:** Improved usability and flexibility of the DELETE command for PDS/PDSE processing.

IDCAMS

- **ALTER NULLIFY Management Class**
  - Currently, IDCAMS ALTER does not allow a user to nullify a management class, but ISMF provides support to remove the management class from a dataset by specifying '-'.
  - ***New enhancement:*** Allow a user to specify 'ALTER NULLIFY' to nullify a Management Class.
    - Users can specify NULLIFY(MANAGEMENTCLASS) to nullify the Management Class of a dataset.
    - The abbreviation is NULLIFY(MGMTCLAS).

? **Why it Matters:** Improved usability of the ALTER function.

IDCAMS

- **DIAGNOSE of GDGs**
  - Currently, IDCAMS DIAGNOSE command does not detect a mismatch in the actual number of extension cells versus how many GDG BCS records exist.
    - This can lead to GDG processing errors that are undetectable until batch processing fails.
  - ***New enhancement:*** Enhance AMS DIAGNOSE to crosscheck the extension cells and the GDG BCS record, so the mismatch can be detected.
    - DIAGNOSE will return a failing return code, if there is a mismatch of the actual number of extension cells and the number of BCS records of a GDG.

? **Why it Matters:** Intended to help users easily and quickly identify the cause of GDG processing errors.

z/OS DFSMS Highlights

- **Access Methods**
  - **V2.1**
    - VSAM RLS Directory Only Caching
    - VSAM RLS 64-bit Enhancements
    - VSAM RLS Dynamic Volume Count
    - VSAM SMB Enhancements
    - VSAM SHOWCB Enhancements
    - DFSMS Support for zHPF
    - SAM EF Support for FlashCopy

Session 14142/14144: Unclog Your Systems with z/OS 2.1 - Something new and Exciting in Catalog Part I & II, **Tuesday 1:30 and 3PM**

Session 14156: DFSMS Basics: Transactional VSAM (TVS) Basics/Implementation, **Wednesday 4:30PM**

Session 14153: DFSMS Advanced: RLS Diagnostics and Recovery, **Friday 11AM**

Access Methods

- **VSAM RLS Directory Only Caching**
  - Accessing data stored in the local buffers is the quickest way for a user to access shared data.
    - However, in a sysplex environment if a system has invalidated the local copy because another user has updated the data, users must gain access to the data in another way.
  - Accessing data from the cache structure in the coupling facility is the next fastest way for the user to access the shared data.
    - Data in the XCF cache structure is directly accessible to any system in the sysplex that has access to the structure.
  - ***New enhancement:*** Provides a new keyword, DIRONLY, for SMS DATACLAS RLSCFCACHE to bypass caching all RLS data for files, including the index component, when the cost of caching any data in the coupling facility outweighs the benefits.

Access Methods

- **VSAM RLS Directory Only Caching**
  - **? Why it Matters:**
    - Directory only cache will benefit RLS customers with limited coupling facility storage but who still need to share VSAM data sets across a parallel sysplex.
      - They will be able to define small cache structures and exploit them only to maintain data consistency.
    - Customers with single system plexes configuration still need to define XCF cache structures and have their data sets connected to them to use RLS; however these cache structures can be very small with directory only caching.
    - Depending on an applications design and workload, some RLS customers, including single system sysplex users, will also experience performance improvements when using Directory Only Cache.
      - RLS will skip writing the data to XCF cache structure every time the data is updated.

Access Methods

- **VSAM RLS 64-bit Enhancements**
  - As customers have taken advantage of 64 bit buffers, the number of control blocks needed to support additional buffers has increased dramatically.
  - ***New enhancement:*** Move of a number of RLS buffer-related control blocks from the SMSVSAM data space into 64-bit storage.

? **Why it Matters:** Increase the amount of available SMSVSAM data space storage and is expected to help improve performance when processing a large amount of VSAM RLS data.

Access Methods

- **VSAM RLS Dynamic Volume Count**
  - Currently setting Dynamic Volume Count (DVC) along with the Space Constraint Relief attribute in the SMS data class used for a data set can be used to determine the maximum number of volumes it will be allowed to span, to increase the original volume count specified for data sets in JCL or when using Dynamic Allocation.
    - Enables the data set to be extended later should it run out of space on the volumes on which it was originally allocated, and is intended to help prevent space-related abends, and is only supported by base VSAM.
  - ***New enhancement:*** Remove the restriction and extend DVC to support VSAM RLS data sets.
    - Note: To be activated, all sharing systems must be running z/OS V2.1

? **Why it Matters:** Prevent space-related abends when data sets grow during VSAM RLS processing.

Access Methods

- **VSAM SMB Enhancements**
  - VSAM supports the use of system-managed buffering (SMB) for VSAM data sets. In prior releases, SMB access bias (ACCBIAS) specifications could be made in JCL, but not specified at the system level.
    - Currently, to change the SMB and RMODE31 options for the data sets in a data class, potentially hundreds of JCL DD AMP statements would need to be changed.
  - ***New enhancement:*** Specify SMB Record Access Bias values for VSAM data sets in the SMS data classes and override the ACB RMODE31 parameter with SMS data class specifications.

? **Why it Matters:** Users, who need to update SMB in hundreds or thousands of their jobs, can save time by simply modifying the data class to satisfy their SMB needs instead of individually changing the JCL “AMP=” SMB attributes in each job.

Access Methods

- **VSAM SHOWCB Enhancements**
  - The VSAM SHOWCB macro provides information about open VSAM data sets.
  - ***New enhancement:*** Enhances two sub-parameters for the SHOWCB macro to display fields of an Access Method Control Block (ACB):
    - BUFNOL – to return the number of buffers that was obtained during BLDVRP or SMB for a particular data set component and its LSR buffer pool.
    - BUFUSE – to return the number of buffers in the LSR or NSR buffer pool that are currently being used.

? **Why it Matters:** Help application programs to tune their VSAM buffering; for example, determine whether to change their LSR buffer pool sizes.

Access Methods

- **DFSMS Support for zHPF**
  - z/OS R11 provided the initial support for System z High-Performance FICON (zHPF) and exploited data sets accessed using the media manager component of DFSMS, including VSAM data sets.
  - z/OS R13 added support for QSAM, BSAM, and BPAM and allowed EXCPVR callers to use zHPF channel programs.
  - ***New enhancement:*** Add zHPF support for EXCP
    - The function is also available for z/OS V1.12 and V1.13 with the PTF for OA38185 and OA40697* (DFSMS).

? **Why it Matters:** provide function that programmers can use to achieve significant I/O performance improvements for programs using EXCP.

Access Methods

- **SAM EF Support for FlashCopy**
  - Multi-volume, single-striped extended format sequential data set cannot use FlashCopy with DFSMSdss because the volume boundary cannot be changed.
    - The suffix in each block contains the block number on that volume, and the FlashCopy process cannot adjust those numbers.
  - Similarly a single-volume, single-striped extended format data set cannot use FlashCopy if the destination requires multiple volumes.
  - ***New enhancement:*** Support the use of FlashCopy by removing volume boundary awareness from sequential extended format data sets.
    - The system will create a new format, Format 2, of an extended format data set which contains no volume boundary awareness within the suffix associated with each physical block.
    - The data set's catalog entry will indicate the data set's version., and LISTCAT and DCOLLECT will display the version.
    - **Do not set version 2 until all sharing and backup systems are at V2.1 or compatibility PTFs have been installed.**

? **Why it Matters:** Removes a long standing restriction and allows the exploitation of FlashCopy with single striped extended format data sets.

z/OS DFSMS Highlights

- **PDSE**
  - **V2.1**
    - PDSE Version 2 Enhancements
    - PDSE Member Generations
    - PDSE Larger Member Size
    - GDG support for PDSEs
    - IEBCOPY Enhancements

Session 14147: The Future of PDSE: New Features in z/OS 2.1, **Tuesday 9:30AM**

Session 14146: Using PDSEs in your SYSPLEX: Best Practices and Troubleshooting, **Wednesday 11AM**

PDSE

- **PDSE Version 2 Enhancements**
  - ***New enhancement:*** Simplify PDSEs to be able to make partial release more effective and to improve general PDSE performance.
    - The system will create a new PDSE format, Format 2, which will allow all unused space to be released, consolidate directory pages when possible, improve read performance, and reduce virtual storage utilization for PDSE processing.
      - *Toleration of the new PDSE format is planned for z/OS V1.12 and z/OS V1.13*
    - Overall PDSE performance will be improved.
      - *The path length of almost all PDSE related operations will be reduced, and index searches will be improved.*
      - *Unnecessary structures from the directory will be removed allowing space to be used more efficiently.*

? **Why it Matters:** These enhancements are intended to provide additional scalability and usability benefits of using PDSEs in place of PDSs, make it feasible to use PDSEs instead of multiple large sequential data sets, and help reduce the space required for PDSEs.

PDSE

- **PDSE Member Generations**
  - ***New enhancement:*** When you create a PDSE, you will be able to specify the maximum number of generations for the system to retain for replaced members.
    - MAXGENS keyword on the DD statement or dynamic allocation or Data class
    - PARMLIB will have a system limit on the generations limit.
    - Users will need to allow for more space for the data set
    - Applies to both SMS-managed and non-SMS-managed PDSEs and is supported for program objects and data members
    - Each time that a member is replaced (not updated in place) the replaced generation will be retained.
      - *If the generations limit has been reached, the oldest generation will be deleted permanently.*
    - Old generations will not be visible to current APIs
      - *Reading the directory will see no change*
      - *New options on the DESERV macro will reveal old generations and their aliases.*
    - The aliases for each generation will be retained with it
    - ISPF and IPT will provide ways to display information about old generations and to recover them
      - *TSO commands or JCL cannot be used to see old generations*
  - **Support for recovering prior levels of a PDSE member is planned to be made available with a PTF for APAR OA42358 in the first quarter of 2014.****
    - Additional PTFs will tolerate and ignore old generations.

? **Why it Matters:** Allows deleted or replaced PDSE members to be recovered.

** IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

PDSE

- **PDSE Larger Member Size**
  - Currently PDSE members are limited to 15,728,639 records.
    - Message IEC036I 002-A8 is issued if a PDSE member exceeds 15,728,639 lines. This limit does not exist for PDS datasets.
  - ***New enhancement:*** Increase the limit on PDSE member size.
    - PDSE member size is planned to be over 125 times larger (approximately 2,146,435,071 records) than the current limit in many circumstances, and substantially larger than the maximum supported size of a PDS member.
  - ***SHARE Requirement met:*** **SSMVSS11010**

? **Why it Matters:** Provide additional scalability and usability benefits of using PDSEs in place of PDSs and make it feasible to use PDSEs instead of multiple large sequential data sets.

PDSE

- **GDG support for PDSEs**
  - Currently PDSs are supported as GDSs (generation data sets). However, PDSEs are not supported as GDSs.
  - ***New enhancement:*** Provide support for using generation data groups (GDGs) comprising PDSE generation data sets.
    - This support, planned to be similar to existing GDG support for PDS data sets.
    - PDSE can be defined as a GDS by specifying DSNTYPE=LIBRARY or a DATACLAS that specifies DSNTYPE=LIBRARY on the GDS define command.

? **Why it Matters:** Extend usage of PDSEs.

PDSE

IEBCOPY Enhancements

- Currently, IEBCOPY can be used to copy to or from a PDSE data set a member and its aliases, together as a group. In addition, copy functions require that member names be fully qualified.
  - COPYGRP supports all combinations of group copy requests for PDSE to PDSE, PDSE to PDS and PDS to PDSE, but PDS to PDS group copies are currently treated as a COPY operation.
- ***New enhancement:*** New COPYGROUP function designed to copy members and all aliases for any combination of PDS and PDSE data sets.
  - A superset of the existing COPYGRP function.
  - All aliases in a group will be copied with the member or neither the aliases nor the member in a group will be copied.
  - The EXCLUDE statement is not supported.
- ***New enhancement:*** Allow the user to pass a filter pattern mask on COPYGROUP functions
  - The MEMBER sub parameter of the SELECT and EXCLUDE statements will be enhanced to accept the "*" and '%' filter control characters.

? **Why it Matters:** Improved usability for IEBCOPY.

z/OS DFSMS Highlights

- **SMS**
  - **V2.1**
    - Provide Accurate Volume Space Statistics
    - Alter ACDS/COMMDS to SHROPT(3 3)
    - New PARMLIB Option to Generate PDSE
    - SMS ACS Read-only Variable for EAVs

Session 14157: DFSMS Intermediate: Naviquest – Streamlining SMS work**, Wednesday 8AM**

Session 14158: DFSMS Intermediate: Understanding Somebody Else's ACS routine**, Wednesday 1:30PM**

Session 13773: What You Need to Know About the Way HSM Uses SMS, **Thursday 3PM**

SMS

- **Provide Accurate Volume Space Statistics**
  - Currently, SMS updates space statistics of SMS-managed DASD volumes in the Active Configuration when
    - CVAF informs SMS that volume space usage has changed
    - A volume is varied online at the first time
    - User issues IGDCNS call to retrieve volume definition that does not contain updated space statistics
    - SMS issues LSPACE and updates the volume definition with most current space statistics
  - SMS does **not** refresh volume space statistics when resizing is done from a different SMS-plex.
  - ***New enhancement:*** SMS is enhanced to detect the size change and also provides a new command for the user to refresh volume space statistics when needed.
    - DCE contains the newly resized value
    - SMS issues LSPACE to obtain space statistics after resize and updates volume definition in active configuration
    - V SMS,VOL(volser)|SG(sgname),SPACE

? **Why it Matters:** Ensures volume space statistics are current and accurate.

SMS

- **Alter ACDS/COMMDS to SHROPT(3 3)**
  - ***New enhancement:*** SMS will be enhanced to check whether or not the SHAREOPTIONS(3,3) or higher is specified for ACDS or COMMDS.
    - If a lower share option is detected by SMS during CDS activation, SMS will attempt to alter the share option to (3,3) or higher and issue a message to inform the user of the result of the change.
    - If the change is successful, SMS will issue new message IGD098I. Otherwise new message IGD099I will be issued. In both cases the activation process continues as before.

? **Why it Matters:** "Soft" enforcement of SMS best practice to help users avoid potential problems with SMS due to incorrect sharing options.

SMS

- **New PARMLIB Option to Generate PDSE**
  - Currently, the DSNTYPE={LIB|HFS} does not guarantee that a partitioned data set will be created
    - A partitioned data set is created if DSORG=PO or directory blocks are specified
    - Otherwise, a sequential data set is created.
  - ***New enhancement:*** SMS provides a new parameter in IGDSMSxx that directs SMS to create a PDSE:
    - HONOR_DSNTYPE_PDSE(YES|NO); NO is the default and the processing remains as before
    - HONOR_DSNTYPE_PDSE(YES) applies only when DSNTYPE={LIB|HFS}
      - *A PDSE is created regardless of the specification of DSORG and directory blocks when HONOR_DSNTYPE_PDSE(YES)*
    - SMS provides a new SETSMS command to modify the setting

? **Why it Matters:** Allows the user to specify that partitioned data sets be unconditionally allocated as PDSE when DSNTYPE=LIBRARY is specified, whether or not directory space is also specified in JCL.

SMS

- **SMS ACS Read-only Variable for EAVs**
  - Currently, the EATTR keyword can be specified on
    - JCL
    - Dynamic allocation
    - AMS DEFINE
    - DATACLAS
  - However, it's NOT available to the installation's ACS routines
  - ***New enhancement:*** Provide a new ACS Read-only Variable for EAV:
    - &EATTR - contains the extended attributes for EAV
    - Expected Values:
      - *OPT: extended attributes are optional*
      - *NO: no extended attributes*
      - *Blank : not specified (This is a default value)*

? **Why it Matters:** ACS routines can be more intelligent to select proper SMS constructs for EAV, so a data set can be allocated to an EAV storage group.

z/OS DFSMS Highlights

- **DFSMSdfp**
  - **V2.1**
    - OCE Partial Release Enhancements
    - OCE RAS Enhancements
    - XTIOT HealthCheck

OCE

- **OCE Partial Release Enhancements**
  - Currently unused space at the end of a data set is released under these conditions:
    - Sequential or partitioned data set
    - RLSE was coded on the DD statement or the management class specifies it
    - It happens during HSM space management or when a program closes the data set that is open for writing.
  - If the data set has multiple volumes, the space is released only on one volume, not on subsequent volumes that the data set might have been extended to previously.
    - An exception is striped data sets, where space is released on all stripes if possible.
  - ***New enhancement:*** If SMS-managed, then all the space in the data set on subsequent volumes will be released.
    - The format 1 or 8 DSCB will remain with no extents.
    - The catalog entry will still show the volume serials.
    - Space is released even if the storage class says “guaranteed space”.
- ***SHARE Requirement Fully Addressed:*** SSMVSS08002
  - Partially addressed in R12 which addressed SMS, Extended Format (EF) datasets.
  - Now fully addressed with support for sequential data sets

? **Why it Matters: More efficient use of storage resources**

OCE

- **OCE RAS Enhancements**
  - Eliminate ABEND 837 RC08
    - Currently during EOV tape output processing when another volume needs to be added to the JFCB volume list and a JFCB extension is required but does not exist, an abend 837 RC08 is issued
    - ***New enhancement:*** Eliminate the abend by dynamically calling Allocation to create the JFCB extension required to add the current volume to the volume list.

    ? **Why it Matters:** Automatic error recovery and avoidance of abends.

  - Parmlib member IEAAPP00 should allow comments
    - Parmlib member IEAAPP00 can be used to define authorized I/O appendage routines. Currently IEAAPP00 processing in IEAVNP16 fails if comments are included.
    - ***New enhancement:*** Allow comments in parmlib member IEAAPP00
      - Comments will be allowed both at the start of a member or interspersed throughout the member.

    ? **Why it Matters:** Improved communication regarding changes to parmlib members.

OCE

- **OCE RAS Enhancements**
  - Externalize reason for DCBE invalidation
    - Function was introduced in z/OS V1.13 that allows users to specify that explanatory text, for a number of DFSMS abends, be included in job output.
    - ***New enhancement:*** Externalize in a new version of message IEC190I issued to joblog and syslog the reason for OPEN processing invalidating the DCBE.
      - *When OPEN processing finds a problem with the DCBE, OPEN processing continues but without processing the DCBE.*
      - *The additional diagnostics is intended to help avoid problems related to unexpected processing results due to not processing the DCBE.*

? **Why it Matters:** Improved problem diagnostics.

OCE

- **XTIOT HealthCheck**
  - IBM recommends that all users set NON_VSAM_XTIOT=YES in the DEVSUPxx member of PARMLIB.
    - Enables application programs to exploit options on dynamic allocation with BSAM, BPAM and QSAM and affects EXCP.
  - ***New enhancement:*** To remind system programmers to enable it, a new OPEN/CLOSE/EOV health check, OCE_SMSOCE1 will be supplied.
    - It needs to be run only once per IPL.
    - It has no options.
    - It issues a warning message if NON_VSAM_XTIOT is not set to YES.

? **Why it Matters:** Using XTIOTs is recommended because it provides virtual storage constraint relief (VSCR) below the 16 MB line.

z/OS DFSMS Highlights

- **DSS**
  - **V2.1**
    - Reset with RESTORE
    - zFS Change Activity Support

Session 14138: What's New in DFSMSdss, **Thursday 8AM**

DFSMSdss

- **Reset with RESTORE**
  - The DFSMSdss Full Volume RESTORE command allows users to restore a volume from a DFSMSdss dump data set.
    - The DS1DSCHA (data-set-changed) bit that is found in the Format-1/8 DSCB in the VTOC indicates whether or not the data set has changed since its last backup.
    - During DUMP processing, the user may specify the RESET keyword which tells DFSMSdss to turn off the DS1DSCHA bit once the data set is successfully dumped.
    - During full volume RESTORE, DFSMSdss unconditionally resets (turns off) the data-set-changed indicator for each data set restored to the target volume.
  - ***New enhancement:*** New RESET keyword, RESET(DUMP|YES|NO), for RESTORE Full and Tracks.
    - Used to turn off the "data set changed since last backup" indicator (DS1DSCHA) for all data sets restored from the dump.
    - DUMP – If RESET was specified on the DUMP command at z/OS R2.1 or higher, then the indicators will be turned off. If RESET was not specified, then the indicators will not be touched.
    - YES – Turn the DS1DSCHA indicators off
    - NO – Do not change the DS1DSCHA indicators from what they are in the backup.
  - ***SHARE Requirement met:*** **SSMVSS07002**

? **Why it Matters:** Intended to make policy-based storage management more effective for recently restored volumes.

DFSMSdss

- **zFS Change Activity Support**
  - DFSMSdss will provide support to reset the data set changed (DS1DSCHA) flag in the Format-1/8 during logical data set dump processing.
  - ***New enhancement:*** Set an indication that a file system has changed, allow its use in DFSMSdss dump command filtering to back up changed file systems, and reset it after a successful dump.

? **Why it Matters:** Helps reduce unnecessary backups for mounted file system data sets.

z/OS DFSMS Highlights

- **SDM**
  - **V2.1**
    - XRC Offline Volumes

Session 13636: What's New in GDPS 3.10**, Wednesday 3PM**

Session 14139: System z and Storage Synergy**, Wednesday 11AM**

Session 14097: Enhanced Availability and IT Resilience: An Integrated TS7700 Grid**, Thursday 4PM**

Session 12259: IBM Data Replication Solutions for System z and Beyond**, Thursday 4:30PM**

SDM

- **XRC (ie z/OS Global Mirror) Offline Volumes**
  - The current implementation of XRC requires that primary volumes (the volumes being written to by application programs) be online to the XRC system when the volumes are added to an XRC session, and when a XRC session is restarted.
    - As customer installations have grown from hundreds, to thousands, to tens of thousands of primary volumes, the time to vary these volumes online has become quite large.
      - *"Vary online processing before XSTART takes excessive amount of time, and if vary 'missed', XADD processing fails."*
      - *"Volumes offline, XDELPAIR fails after XSTART, unable to locate UCBs"*
  - ***New enhancement:*** Allow XRC primary volumes to be offline when the XSTART and XADDPAIR commands are issued to start or restart mirroring for existing volumes.

? **Why it Matters:** Improve availability by eliminating the need to wait for all devices to be varied online.

z/OS DFSMS Highlights

- **DFSMSrmm**
  - **V2.1**
    - DFSMSrmm SMS Mgmt Class for Tape
    - DFSMSrmm RAS Enhancements

Session 14092/14095: RMM and HSM Report Generator and Hands-on Lab, **Thursday 9:30AM and 11AM**

DFSMSrmm

- **SMS Management Class for Tape**
  - Continuation of the strategy of moving DFSMSrmm expiration and retention decisions outside VRSEL inventory management, and enabling the use of DFSMS Management Class attributes for tape data sets.
  - Currently, users who want the benefits of the EXPDT retention method must assign the expiration date and Last reference days to each data set explicitly by the JCL, DataClass, TSO RMM command or use the DFSMSrmm parmlib default.
    - New requirement to automate these settings.
  - ***New enhancement:*** To automate the setting of expiration date and Last reference days DFSMSrmm now gives the user the possibility to obtain management class attributes relevant for tape data set management.
    - New RMM parmlib option MCATTR
      - *When Management Class attributes are enabled, the expiration information is retrieved by DFSMSrmm during OPEN for output and used to set the expiration date and the LASTREF extra days for the tape data set.*
      - *Provides an option to use the Management Class attributes for all volumes **except** the Management Class expiration attributes for volumes managed by the VRSEL retention method.*
      - *Regardless of whether the Management Class attributes are used or not for a tape data set, the Management Class name is recorded, just as is done today, in the data set record, and if the volume is managed by VRSEL retention method will be used in the normal VRS matching (same as done today).*
      - *APAR OA35808 provides toleration support.*
  - ? **Why it Matters:** Provides simplification by allowing the user to automate setting the expiration date or LASTREF extra days, without using the JCL expiration date or Data Class or TSO command.

DFSMSrmm

- **RAS Enhancements**
  - **Expire after days non-usage/Last reference date**
    - Currently, for volumes managed by the VRSEL retention method, users can define vital record specifications to retain all copies of the data set based on the number of days since the data set was last read or written.
      - *New requirement to manage data based on the number of days since the data set was last read or written on volumes managed by the **EXPDT retention method**.*
    - ***New enhancement:*** A new attribute, LASTREF extra days is added to the data set record for data sets on volumes managed by the EXPDT retention method.
      - *The LASTREF extra days can only be set for data sets on volumes managed by the EXPDT retention method. It can be set explicitly, via PARMLIB, or by Management Class.*
      - *APAR OA35808 provides toleration support.*

? **Why it Matters:** The user can manage his data based on the number of days since the data set was last read or written without VRSEL processing.

DFSMSrmm

- **RAS Enhancements**
  - **'Retain by' options for volumes managed by the EXPDT retention method**
    - Currently, for volumes managed by the VRSEL retention method, the user can specify whether a volume set is to be retained as a whole set expiring all at the same time or as individual volumes.
      - *New requirement to manage volume sets that are managed by the EXPDT retention method with the choice of retaining them as individual volumes or as whole sets or based on the expiration date of the first file.*
    - ***New enhancement:*** A new attribute, RETAINBY is added to the volume record for volumes managed by the EXPDT retention method.
      - *Possible values for RETAINBY attribute are VOLUME, FIRSTFILE and SET.*
      - *A new sub-parameter RETAINBY for the EXPDT retention method is added to the DFSMSrmm PARMLIB member to specify a system wide default.*
      - *RETAINBY attribute is the same for all volumes in a set, it cannot explicitly be set for volumes other than the first. DFSMSrmm ensures that volumes managed by the EXPDT retention method inherit the RETAINBY attribute from their previous volumes if existing.*
      - *APAR OA35808 provides toleration support.*

? **Why it Matters:** The user can manage his EXPDT retention method volume sets as individual volumes, as a whole set or based on the expiration date of the first file.

DFSMSrmm

- **RAS Enhancements**
  - **Conversion Support Changes**
    - RMM implemented a new data set attribute LASTREF and a new volume attribute RETAINBY, for data sets and volumes managed by the EXPDT retention method.
      - *If a customer wants to use these new attributes he has to use TSO RMM CHANGEDATASET and CHANGEVOLUME commands.*
    - ***New enhancement:*** set the LASTREF and RETAINBY attributes during conversion.
      - *A default value can be set via new EDGCNVT SYSIN statements :*

      OPTION EXPDT_RETAINBY

      OPTION EXPDT_LASTREF.
      - *APAR OA35808 provides toleration support.*

? **Why it Matters:** At conversion time, for the EXPDT retention method, similar attributes of the input data can be translated to the RMM data set and volume attributes. No extra updates of these attributes are needed after the conversion.

z/OS DFSMS Highlights

- **OAM**
  - **V2.1**
    - OAM Usability & Reliability Enhancements

OAM

- **OAM RAS Enhancements**
  - OAM Support for Tape Block Sizes > 32 KB
    - ***New enhancement:*** Improved tape performance by supporting larger block sizes for tape.
      - *Optionally enabled via new keyword on SETOAM statement in CBROAMxx member of PARMLIB*
      - *Volume supports larger block sizes if support enabled when first object written to that volume*
      - *Volumes that support block sizes > 32 KB not accessible from prior releases*

    ? **Why it Matters:** Improved performance and scalability for OAM tape processing.

OAM

- **OAM RAS Enhancements**
  - OAM Support for Automatic Backup Deletions
    - ***New enhancement:*** Improved tape performance by automatically removing unneeded backup copies.
      - Optionally enabled via new keyword on SETOSMC statement in CBROAMxx member of PARMLIB
      - When enabled, OSMC processing will delete existing backup copies that exceed the number required by object's current management class (e.g. when object transitions to TS7700 Virtualization Engine with multi-cluster grid support, OAM backups may no longer be needed).
      - SMF – Existing OAM SMF Record Type 85 (x'55') Subtype 32 changes:
        - New flag will indicate whether automatic backup deletion was enabled at the time of OSMC processing
        - New fields will show number of backup copies and total bytes deleted as a result of automatic backup deletion

    ? **Why it Matters:** More automated space management and exploitation of IBM Tape HW functionality.

OAM

- **OAM RAS Enhancements**
  - Reduced Size Limit for OAM Store Sequence Processing
    - ***New enhancement:*** Improved performance by enabling the OSREQ Store Sequence support on small object sizes.
      - *Minimum size required for OSREQ Store Sequence processing (STOREBEG/STOREPRT/STOREEND) reduced from 256 MB+1 to 50 MB+1*
      - *Provides option for storing objects in 50 MB+1 to 256 MB size range*
      - *Not supported when writing to optical*

    ? **Why it Matters:** Improved OAM usability and reduced application virtual storage requirements.

OAM

- **OAM RAS Enhancements**
  - OAM ATAM Coexistence Enhancement
    - ***New enhancement:*** Allow users to specify how long OAM's object support should wait before issuing a message when no tape devices are available.
      - New keyword on SETOAM statement in CBROAMxx member of PARMLIB to specify retry time period.
        - Can be set to 0 to indicate no retries.

    ? **Why it Matters:** Improved OAM interoperation with products such as IBM Tivoli Automated Tape Allocation Manager for z/OS (ATAM, 5698-B15),.

z/OS DFSMS Highlights

- **DFSORT**
  - **V2.1**
    - Dynamic Sort Enhancements
    - Functional Enhancements
    - 64-bit Enhancements

See San Francisco Conference Proceedings:

Session 12983: DFSMS DFSORT: The ICETOOL Cometh

Session 12985: DFSMS DFSORT: Resource Usage and Understanding

DFSORT

- **Dynamic Sort Enhancements**
  - While DFSORT has installation defaults designed to control the use of central storage resources by all SORT applications running concurrently on a system, customers still often experience degraded performance due to over commitment of resources and high paging.
    - *This is usually due to changes in available resources after a sort has already committed to using large amounts of central storage.*
  - Often the first submitted DFSORT job can possibly use most of the central storage for creating dataspaces, hiperspaces or memory objects.
    - *Any DFSORT job(s), submitted later, use very little or no dataspace, hiperspace or memory objects because the central storage is not available at that time.*
  - ***New enhancement:*** A new TUNE option is designed to allow users to specify that DFSORT obtain storage incrementally and check on storage availability before allocating additional storage.
    - *Helps to better balance utilization for sort operations and other workloads initiated within a short time.*
  - ***New enhancement:*** Increase the memory object work space maximum from 64 GB to 1 TB, to allow users to sort larger amounts of data in memory object work files.

? **Why it Matters:** Improved memory resource management to better balance the memory requirements of multiple large concurrent sort operations and other workloads.

DFSORT

- **Functional Enhancements**
  - **Alphanumeric Tests**
    - ***New enhancement:*** Support new alphanumeric tests, using binary format, for both compare fields and parse fields, including combinations of alphanumeric character sets (uppercase and lowercase, and numeric).

    ? **Why it Matters:** Allows users to specify various sets of characters using a single compare condition or PARSE keyword rather than using compare conditions or PARSE keywords.

  - **Symbol and PARSE Enhancements**
    - ***New enhancement:*** Allow symbols to be used for more DFSORT operands, and the number of parse fields supported is planned to be increased from the prior limit of 100 fields to 1,000 fields.

    ? **Why it Matters:** Improved DFSORT usability

  - **Add String at End of Variable-Length Records**
    - ***New enhancement:*** Allow users to specify that a string up to 50 characters in length be appended to variable-length output records.

    ? **Why it Matters:** Improved DFSORT usability.

DFSORT

- **64-bit Enhancements**
  - Programs can call (invoke) DFSORT and use an invocation parameter list to pass information from the program to DFSORT.
  - During DFSORT processing, exits can be optionally called. Information is passed between DFSORT and exits in 2 different ways:
    - User exit parameter lists
    - COBOL interface used only for E15, E32, and E35 exits written by customers in COBOL
  - ***New enhancement:*** Provide Blockset sorting support for programs running in 64-bit addressing mode.
    - This new function is designed to be available to programs, using new parameter lists for DFSORT applications that use E15, E35, or E32 exits to process 64-bit addressed records.

? **Why it Matters:** Helps relieve storage constraints for programs calling DFSORT to perform certain sort operations.

z/OS DFSMS Highlights

- **NFS**
  - **V2.1**
    - NFS Server 64-bit Enhancements
    - NFS RPCSEC Performance Enhancements
    - NFS Server RPCBIND Enhancements

NFS

- **Server 64-bit Enhancements**
  - The z/OS NFS Server provides two data paths, one for the z/OS UNIX file systems (such as zFS), and the other for the traditional z/OS MVS data sets.
    - When writing to z/OS MVS data sets, the z/OS NFS Server has to buffer the RPC WRITE data so the buffered data logically appears sequential before the z/OS NFS Server call DFSMSdfp to write the blocks of data.
    - With the large data sets and the random write from the NFS Clients, the z/OS NFS Server 2GB Address Space is now required to handle concurrent and random writes to large data sets.
  - ***New enhancement:*** NFS Server exploits 64-bit addressing to support larger sequential data sets, PDS members, and PDSE members.
    - Designed to support processing for files as large as 4 TB, up from the prior limit of 800 MB.

? **Why it Matters:** Potentially helps to improve application performance for random access.

NFS

- **RPCSEC Performance Enhancements**
  - RPCSEC_GSS is an authentication flavor of the Remote Procedure Call (RPC) protocol that is supported by the z/OS NFS server for its NFS V4 workloads.
    - Kerberos V5 is the underlying security mechanism supported for this authentication flavor.
    - Currently the z/OS NFS server invokes the GSS and Kerberos API's from a single thread (Main task). This affects the performance of the secure mount workloads, as those requests have to be routed to a single main task for the invocation of the relevant Kerberos and GSS APIs.
  - ***New enhancement:*** NFS Server is designed to use multi-tasking for the RPCSEC_GSS authentication type of the Remote Procedure Call (RPC) protocol, which is supported by z/OS NFS server for NFS V4 workloads.
    - The Kerberos and GSS API's will now be invoked from the Transport and Worker tasks of the NFS server and will no longer need to be routed to the Main task.

? **Why it Matters:** Potentially helps to improve the performance of secure mount workloads in the NFS server.

NFS

- **NFS Server RPCBIND Enhancements**
  - Currently on the z/OS system, when the RPCBIND fails, the z/OS NFS Server has no capability to re-register with RPCBIND when it is restarted.
    - All existing connections to the NFS Server remain operative. However, no new NFS V2 or NFS V3 mounts can be established to the NFS Server.
    - In order for the NFS Server to reconnect to RPCBIND, it must be restarted, which impacts the existing mounts to the NFS Server.
  - ***New enhancement:*** The RPCBIND and NFS Servers are designed to allow the NFS Server to re-register with RPCBIND when RPCBIND is restarted, without an NFS Server restart.
    - This is designed to help preserve existing connections to the NFS Server and to allow new mounts when RPCBIND is restarted.

? **Why it Matters:** Potentially helps to improve the NFS Server resiliency and availability by eliminating a reason for NFS Server restarts.

z/OS® DFSMS™ Highlights

- Where to find additional information:
  - **V2.1**
    - **DFSMS Using the New Functions (SC23-6857-00) -** http://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss?CTY=US&FNC=SRX&PBL=sc23-6857
    - **z/OS V1.12 DFSMS Technical Update** (available 9/30/2013)
  - **R13**
    - **DFSMS Using the New Functions (SC26-7473-08) –** (http://publibz.boulder.ibm.com/cgi-bin/bookmgr/Shelves/ez2zo111?filter=DFSMS+Using+the+New+Functions+&SUBMIT=Search+titles)
    - **z/OS V1.13 DFSMS Technical Update -**(http://www.redbooks.ibm.com/redpieces/pdfs/sg247961.pdf)
  - **R12**
    - **DFSMS Using the New Functions (SC26-7473-07) -** http://publibz.boulder.ibm.com/epubs/pdf/dgt2g570.pdf
    - **z/OS V1.12 DFSMS Technical Update -** http://www.redbooks.ibm.com/abstracts/sg247895.html?Open

Sources for more information

- Information about DFSMS and components
  - http://publib.boulder.ibm.com/infocenter/zos/basics/index.jsp?topic=/com.ibm.zos.zdatamgmt/zsysprogc_dfsmselements.htm
- Information about DFSORT
  - http://www-01.ibm.com/support/docview.wss?rs=0&uid=isg3T7000077
- Information about z/OS Storage Management Tools
  - http://www-03.ibm.com/systems/storage/software/zos/index.html
- Information about IBM Tivoli Storage Productivity Center
  - http://www-03.ibm.com/systems/storage/software/center/index.html
- Information about IBM System Storage Disk systems
  - http://www-03.ibm.com/systems/storage/disk/ds8000/index.html
- Information about IBM System Storage Tape systems
  - http://www-03.ibm.com/systems/storage/tape/?lnk=mprST-tsys

Simplifying access to z/OS® skills!

- **IBM Academic Initiative -** https://www.ibm.com/developerworks/university/academicinitiative/
  - Global program that facilitates the collaboration between IBM and educators to teach students the information technology skills they need to be competitive and keep pace with changes in the workplace.
  - Links to many publicly available tutorials, IBM Redbooks, technical articles, and white papers to help you learn about System z and Power systems.

- **System z Academic Initiative -** https://www.ibm.com/developerworks/university/systemz/
  - The System z program provides significant benefits to everyone who participates.
    - Educators: can build relationships with industry, create enhanced employment opportunities for their students, and gain no-charge access to mainframe systems, courseware, and training.
    - Students: can compete in contests and get scholarships and IBM professional certification. Find out more.
    - Customers: can establish relationships with schools worldwide that are teaching Enterprise Computing, and post jobs at Systemzjobs.com to locate qualified students for internships and jobs.
- **Education Opportunities**
  - IBM is designing education packages for System z and would like your input and feedback.
  - Please talk to Neal Bohling for more detailed information.

Thank you!

- IBM's commitment to the mainframe helps deliver:
  - Extreme scalability, and availability
  - Reduced costs and simplified IT infrastructure
  - High performance and energy efficient technologies
  - Resilient and security rich system

Trademarks and Disclaimers

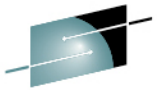

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

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