

What's New in z/OS® DFSMS™ Overview

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

August 2, 2010
Session 8051



SHARE in Boston

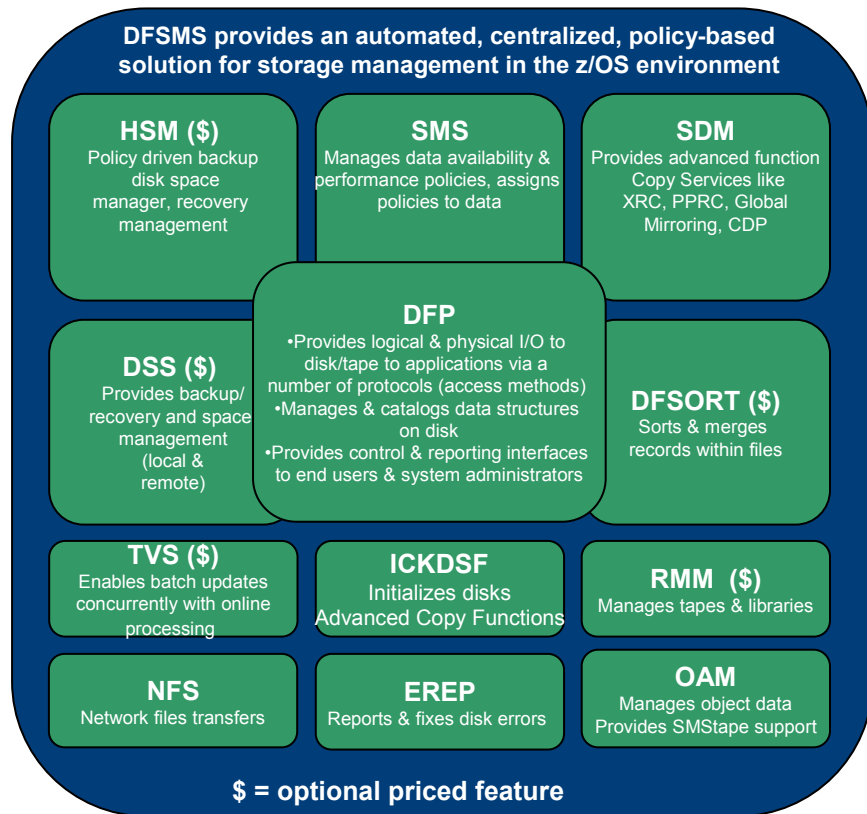
Disclaimer

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.

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



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



DFSMS Supports Information Infrastructure Strategy

Innovative, Integrated, Available Today



Information Compliance

- Ability to cope with increased security and compliance requirements



Information Availability

- Improved storage administrator productivity and simplified management of the z/OS environment
- Increased data storage capacity and scalability to cope with explosive growth of data volumes and database sizes
- Seamless, reliable, performance-sensitive data sharing
- Support for deployment of new processors/systems while exploiting their capabilities efficiently
- Point-in-time copy, fast replication, and continuous data mirroring functions while preserving consistency
- High Availability with simpler, faster, and more reliable recovery operations



Information Retention

- Data availability at all levels of the storage hierarchy
- Long standing provider of critical data management functions



Information Security

- Improved Security with Exclusive media encryption capability, integration with z/OS Key Management³

z/OS® DFSMS™ V1.11 Highlights

(September 2009)



- **Ease of Use**
 - ★ DFSMSrmm Enhancements
 - Catalog HealthCheck Enhancement
 - **IDCAMS Delete Masking**
 - **DFSMS DADSM Enhancement**
 - NFS Client Message Globalization
 - **Application Integration**
 - **DFSMSdfp Enhancements**
 - NFS Server Symbolic Links
 - **Scalability and Performance**
 - **Extended Address Volumes (EAV) – Support non-VSAM data types**
 - DS8000® Dynamic Volume Expansion Enhancements
 - **OAM Scalability – 2GB Tape objects**
 - ★ DFSMSrmm Enhancements
 - VSAM SMB Enhancements
 - NFS Server File Delegation
 - **Optimization & Management Capabilities**
 - ★ DFSMSShsm Fast Replication Enhancements
 - ★ DFSMSShsm ML1 Enhancements
 - ★ DFSMSShsm Backup Copy Retention Period
 - **OAM Archive Retention Enhancements**
 - **SMS Dataset Separation by Volume**
 - SMS Volume Selection Enhancements
 - DFSMS Open/Close/End-of Volume Enhancements
 - ★ DFSMSrmm Enhancements
 - **Security**
 - NFS Elimination of mvslogin and mvslogout for RPCSEC_GSS
 - NFS RPCSEC_GSS V4 Client Security
 - **Availability**
 - NFS Delay Detection Interval
 - **Serviceability**
 - VSAM Serviceability Enhancements
 - NFS Serviceability Enhancements
- ★ **For more details on R11 enhancements see sessions:**
- **8045 What's New in DFSMSShsm**
 - **8041 What's New in DFSMSrmm**

For more details: DFSMS Using the New Functions (SC26-7473-06)
<http://www-03.ibm.com/systems/z/os/zos/bkserv/r11pdf/#dfsms>

z/OS® DFSMS™ V1.12 Highlights

(September 2010)

- **Ease of Use**
 - SMS Healthchecker Enhancements
 - DFSORT Enhancements
 - ★ DFSMSrmm Simplified Monitoring & Management
 - ISMF COPY Storage Group Enhancements
 - Catalog Partial Release Enhancements
 - Catalog DEFINE RECATALOG Enhancements
 - PDSE EMPTY Command
 - PDSE Message Enhancements
 - OCE Enhancements
- **Application Integration**
 - SDM Support of ATTREXX Interface
- **Optimization & Management Capabilities**
 - DFSMS IMBED/REPLICATE Removal
 - ★ DFSMS Fast Reverse Restore Enhancements
 - ★ DFSMSrmm TS7700 Reporting Enhancements
 - SMS Storage Group Mgmt & Volume Selection Enhancements
 - IDCAMS DCOLLECT Enhancements
 - IDCAMS GDG Enhancements
 - NFS Enhancements
- **Availability**
 - ★ DFSMSrmm Active and Queued Task Management
- **Scalability, Performance, & Constraint Relief**
 - VSAM KSDS CA Reclaim
 - EAV Additional Data Set Support
 - ★ DFSMSrmm EAV Exploitation
 - DFSORT Memory Object Sorting
 - DFSMS Support for XTIO and Uncaptured UCBs
 - OAM Enhancements
 - Extended Addressable Catalogs
 - Catalog CAS Contention Detection
 - VSAM RLS Striping
 - ★ DFSMSdss Larger Block Support for DUMP, COPYDUMP, and RESTORE
 - ★ DFSMSShsm Dump Tape Recovery Enhancements
 - ★ DFSMSShsm Space Management Performance
 - ★ DFSMSShsm Dump Stacking
- **Security**
 - NFS Enhancements
- ★ **For more details on R12 enhancements see sessions:**
 - 8045 What's New in DFSMSShsm
 - 8041 What's New in DFSMSrmm
 - 8048 DFSMS Basics: What's New in DFSMSdss and SDM

z/OS® DFSMS™ V1.11 Highlights (September 2009)

- **Ease of Use**
 - IDCAMS Delete Masking
 - DFSMS DADSM Enhancement
- **Application Integration**
 - DFSMSdfp Enhancements
- **Scalability and Performance**
 - OAM Scalability – 2GB Tape objects
- **Optimization & Management Capabilities**
 - OAM Archive Retention Enhancements
 - SMS Dataset Separation by Volume

Ease of Use: Catalog and IDCAMS

■ IDCAMS Delete Masking

- IDCAMS DELETE command is enhanced to include a new function called DELETE MASK. It allows users to specify the data set name selection criteria desired with a mask-entry-name and a keyword “MASK”.
- A mask-entry-name (also called a filter key) can have two consecutive asterisks (**) or one or more percentage signs (%).
 - **Enhancement:** OA29880 - Delete Masking will accept multiple qualifiers
- Delete Masking not deleting entries from the master catalog unless the catalog name is specified.
 - **Enhancement:** OA30208 – Correction to have DELETE MASK follow normal catalog search order
- Delete up to 100 data sets without specifying multiple entry-names
 - **Enhancement:** OA30916/OA31658/OA32220 – Removes the limitation of allowing deletion of up to only 100 data sets.
- Additional fixes to be provided with OA31658:
 - Delete Masking not deleting entries from the master catalog unless the catalog name is specified.
 - Delete Masking does not follow normal catalog search order..
- TSO Support for Delete Masking
 - **Enhancement:** OA31526 – TSO support for Delete Masking
 - **Enhancement:** OA32763 – Help text for TSO DELETE updated to include Delete Masking

? Why it Matters: Single command deletion of related DSN's without qualifier considerations and 100 DSN Limit

Ease of Use: DADSM Enhancement

- **Save Jobname when data sets are created**
 - New fields in the Format 9 DSCB added for the job name, step name and time of data set creation in mapping macro IECSDSL1 and documented in the DFSMSdfp Advanced Services.
 - Prior to this item being implemented, the only way to determine what JOB and STEP created the data set was to track down the SMF records associated with that data set's creation, which took a considerable amount of work.
 - IEHLIST LISTVTOC FORMAT and DUMP options will externalize the value.
- ***SHARE Requirement Partially MET:*** SSMVSS064955
 - DFSMS: Add Jobname to Catalog Record When Data Sets Are Created.
 - Remaining support for DCOLLECT provided in z/OS V1R12.

? Why it Matters: For new DSN's, the creating job name will stay with the DSN instead of possibly getting lost if SMF records "wrap" or get lost.

Application Integration: DFSMSdfp

■ DFSMSdfp Enhancements

- Processing will be changed to indicate end-of-file (EOF) during the allocation of data sets on DASD that are not SMS-managed and have either sequential or an undefined data set organization
- This is intended to make this processing for both SMS-managed and non-SMS-managed data sets consistent, to make it unnecessary to open data sets solely to indicate EOF, and to help prevent programs from accidentally* reading old data when a data set is read after being allocated.
- ***NOTE:** This writing of a file mark at the beginning of new data sets clears only the first track of the new data set. All subsequent tracks are unchanged, and there is still a need to use the erase-on-scratch function to clear any residual data.
- **SHARE Requirement MET:** SSMVSS07016
 - Writing EOF (end-of-file) for non-SMS data sets at create time

? Why it Matters: Avoids accidentally reading old data – this is a data integrity improvement, not a security improvement.

Scalability and Performance: OAM



■ OAM Scalability – 2GB Tape Objects

- In z/OS® V1.10, OAM implemented 2 GB Object Support to enable applications to store objects up to 2000 MB (2097152000 bytes) in size using DB2 on direct access storage
- In z/OS V1.11, OAM will extend that support to the tape tier of the OAM storage hierarchy
 - Intended to provide full support for objects up to 2000 MB in size on both DASD and tape and is expected to reduce the need to separate large binary strings into multiple objects
 - Simplifies the application interface by eliminating the need for applications to materialize entire objects before storing them

? Why it Matters: Performance – avoids having to break large objects up in order to store them

Optimization & Management Capabilities : OAM



■ OAM Archive Retention Enhancements

- Added new retention-related enhancements
- Enhance protection against inadvertent object deletion and provide a mechanism for complying with government regulations
 - Deletion-hold: Prevent object deletion while object is in deletion-hold mode.
 - Deletion-protection: Prevent object deletion prior to object's expiration date.
 - Retention-protection: Prevent object deletion prior to object's expiration date, and don't allow expiration date to be changed (explicitly or implicitly) to an earlier date.
 - *Note: RP in effect for life of object. If expiration date is ever set to 'forever' the object can never be deleted.*
 - Event-based-retention: Object expiration date dependent on external event notification.

? Why it Matters: Adding archival features to z/OS to enable it to meet industry and Legal requirements

SHARE in Boston

Optimization & Management Capabilities: SMS



■ SMS Data Set Separation by Volume

- SMS will support the allocation of critical data sets (such as DB2partitions) on different volumes
 - Reduces I/O contention, provides improved performance and helps to avoid single points of failure
 - Designed to expand the existing data set separation function, to allow you to specify that critical SMS-managed data sets be separated across extent pools and volumes that are not used by other data sets specified in the separation group.

? Why it Matters: The system will ensure DSN's that need to be on separate Disk "spindles" will actually be physically allocated that way.

z/OS® DFSMS™ V1.12 Highlights

(September 2010)



■ Ease of Use

- SMS Healthchecker Enhancements
- DFSORT Enhancements
- DFSMSrmm Simplified Monitoring & Management
 - *For more details see Session 8041 What's New in DFSMSrmm*
- ISMF COPY Storage Group Enhancements
- Catalog Partial Release Enhancements
- Catalog DEFINE RECATALOG Enhancements
- PDSE EMPTY Command
- PDSE Message Enhancements
- OCE Enhancements

Ease of Use: SMS Healthchecker Enhancements

- **New SMS health checks for the communications and active configuration data sets (COMMDS and ACDS)**
 - One new check is designed to alert you that the COMMDS and ACDS are on the same volume
 - As a best practice, an ACDS/COMMDS must reside on a volume, accessible from all systems in the SMS complex.
 - However, to ease recovery in case of failure, the ACDS should reside on a different volume than the COMMDS.
 - The other is intended to identify COMMDS and ACDS data sets that were defined without the REUSE attribute, which is recommended
 - It's a best practice to specify the REUSE option when you define an ACDS or a COMMDS to avoid running into space problems (SMS reason code 6068) as result of subsequent ACDS updates, or IMPORT/EXPORT functions.

? Why it Matters: Helps ensure that your SMS DSN's are allocated correctly according to best practices.

Ease of Use: DFSORT Enhancements

- **Provide VSAM work space estimates**
 - DFSORT needs a good file size estimate to adequately allocate work space to successfully sort input records
 - When VSAM data sets are not closed after updating them (ie when a job ends abnormally), the information stored about their sizes is often incorrect
 - DFSORT can detect when one or more statistics from the Catalog may be incorrect, and in this case, NOT use the statistics to determine file size or number of records in the data set
 - DFSORT will use an alternate method to determine the file size and calculate the workspace needed
 - ICE255I message is changed to be issued only for a SORT application
 - Use of automation software to monitor this message may need to be updated
 - New message ICE264I is issued for a COPY operation if the VSAM dataset is “broken”.
 - The COPY operation will complete successfully. This message is issued only so the user is aware that they may want to fix the statistics.
 - ICE264I ddname STATISTICS MAY BE INCORRECT

? Why it Matters: Users can now run jobs to sort “broken” VSAM data sets and are less likely to receive abends

Ease of Use: DFSORT Enhancements

- **Provide additional diagnostic information**
 - DFSORT will issue new messages in addition to the ICE083A message
 - ICE254I (DFSORT FAILED TO DYNAMICALLY ALLOCATE THE REQUIRED WORK DATA SET SPACE)
 - ICE258I (DFSORT COULD NOT DYNAMICALLY ALLOCATE THE OPTIMAL WORK DATA SET SPACE)
 - Intended to help determine why DFSORT was unable to dynamically allocate all of the requested disk work space
 - Additional messages provide information about how much space DFSORT was attempting to allocate and how much of that it was able to allocate
 - Message also includes the number of volumes to help the user determine if the space amount is too large or if it needs to be spread across more volumes
 - Enables users to
 - Determine possible causes of ICE083A
 - Calculate required work space and number of work volumes for failed sort
 - Take corrective action to resolve or prevent failures
- ? Why it Matters: Easier correction or prevention of errors, improved first failure data capture, and reduced PMRs**

Ease of Use: DFSORT Enhancements

- **Issue diagnostic messages automatically**
 - In error situations, issue DFSORT messages without the need to specify a SORTDIAG DD statement or the DIAGSIM=YES installation option
 - Diagnostic messages (ICE75x, ICE8xx and ICE9xx) are issued for all error situations that result in ICExxxA messages
 - Additional messages are only for diagnosis and are intended only for use by L2, L3 and development and will **not** be documented
 - Users still need to code the SORTDIAG DD or DIAGSIM=YES to cause DFSORT to issue diagnostic messages when DFSORT completes successfully

? Why it Matters: Improved first failure data capture and reduced frequency of requests to recreate customer problems

Ease of Use: DFSORT Enhancements



- **Provide capability for dynamic allocation of additional work data sets that are only used if needed**
 - DFSORT's dynamic allocation of work data sets provided limited capability to react to unexpected increases in disk work space requirements
 - Incorrect file size information is provided to DFSORT
 - Resource contention limits the amount of central storage that can be used as intermediate work space
 - New DYNAPCT installation and run-time option allows you to specify additional work data sets to be used if more workspace is required
 - New DYNAPCT=10 installation default increases number of work data sets allocated for all sorts using dynamic allocation
 - *Does not increase total space allocated*
 - *Additional work data sets are allocated with zero space*
 - New DFSORT message ICE236I shows DYNAPCT value in effect
 - ICE236I OPTIONS: DYNAPCT=a
 - New DFSORT message ICE278I indicates if additional work data sets were required to complete a sort
 - ICE278I x WORK DATA SETS WERE INSUFFICIENT TO COMPLETE THIS SORT SO y
ADDITIONAL WORK DATA SETS WERE USED
 - If additional work data set allocations are a concern, DYNAPCT=OLD will force dynamic allocation to function the same as it did prior to this line item

? Why it Matters: Expected to reduce sort failures that occur due to insufficient work space

Ease of Use: DFSMSrmm Enhancements

- Multiple enhancements provided to support simplified monitoring and management of DFSMSrmm:
 - **Reason why a DFSMSrmm retention limit was reached added to the ACTIVITY file**
 - Available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30881
 - New reports created from the ACTIVITY and extract files are planned to help you see why retention limits were triggered
 - User will have to tailor the provided sample job EDGJACTP in order to fit their environment
 - **Set a volume hold attribute to prevent expiration and to search and report on volumes which have the hold attribute**
 - Difficult to determine how data sets and volumes are retained and changing policies or EXPDT might cause loss of actual retention policy
 - New command prevents a volume being set to pending release and can be changed per TSO command or RMM dialog
 - New sample report EDGGAHLD is provided
 - **OPENRULE ignore processing available for duplicate tape volumes**
 - Used to determine which volser should DFSMSrmm take in case of a „duplicate volume“
 - *For duplicate volumes only the external / requested volser is used to select the OPENRULE entry*

? Why it Matters: Adding/Updating archival features to z/OS to enable it to meet industry and Legal requirements

Ease of Use: DFSMSrmm Enhancements

- Multiple enhancements provided to support simplified monitoring and management of DFSMSrmm:
 - **DFSMSrmm ISPF dialog search results can be bypassed when using the CLIST option**
 - RMM added the CLIST option to the dialog, but without NOLIST, the SEARCH command often terminates for lack of TSO storage
 - **SHARE Requirement MET:** SSMVSS09006
 - *Implement NOLIST option for RMM CLIST commands in dialog*
 - **? Why it Matters:** Bypasses system memory size limitations and reduces the run time
 - **DFSMS plans to provide a system option to control how the system handles multivolume tape label anomalies.**
 - Possible to prevent applications processing tape volumes out of sequence without coding an installation exit.
 - **RMM exploitation of a new Timed Auto Reply Function**
 - Provides an additional way for the system to respond automatically to write to operator with reply (WTOR) messages.
 - **SHARE Requirement MET:** SSMVSS09007
 - *Prevent RMM impact when issuing WTOR in parallel*

Ease of Use: ISMF Enhancements

- **ISMF COPY Storage Group Enhancements**

- Currently installations can copy storage group definitions from one control data set (CDS) to another.
- However, the volumes defined in the storage groups cannot be copied. The Storage administrator will have to manually add volumes to the storage groups.
- ***New enhancement:*** Specify that the volume list for pool-type storage groups be copied at the same time.
 - This allows you to copy entire storage groups from one configuration to another without having to add their volumes to the destination CDS afterward.

? Why it Matters: Saves (potentially much) time and effort

Ease of Use: Catalog Enhancements

- **Catalog Partial Release Enhancements**
 - Currently VSAM partial release, only releases space on volumes where the high-used RBA (HURBA) and the high-allocated RBA (HARBA) are on the same volume, **NOT** space which spans multiple volumes.
 - **New enhancement:** VSAM partial release will release unused volumes in addition to releasing space on the last volume of a multi-volume VSAM data set that contains data.
 - **SHARE Requirement Partially Addressed:** SSMVSS08002
 - Space Release Will not Release Over-Allocated Space for Multi-Volume Files
 - **Note:** This line item only addresses SMS, Extended Format (EF) datasets.
- ? Why it Matters:** More efficient use of storage resources

Ease of Use: Catalog Enhancements

- **Catalog DEFINE RECATALOG Enhancements**
 - IDCAMS DEFINE RECATALOG command will be enhanced for multivolume and striped data sets
 - Automatically create catalog entries with correctly ordered volume lists, while eliminating any duplicate volumes that might have been specified.
 - Makes it easier to recatalog multivolume and striped VSAM data sets.
 - Addresses FIN APAR closure for OA24010

? Why it Matters: Saves administration time, helps avoid errors

Ease of Use: PDS/PDSE EMPTY Command

- Currently, IDCAMS DELETE can only delete one specific member in a PDS or PDSE by specifying the member name in parenthesis.
 - **Example:** DELETE pdsname(member name)
 - Wildcards are not allowed for member names.
 - ***New enhancement:* DELETE all members of a partitioned data set in a single operation**
 - Specifying a wildcard character (*) as the member name for a data set when using the DELETE command
 - Allows you to remove all members of a PDS or PDSE data set in a single command.
 - ***SHARE Requirement MET:* SSMVSS063069/70**
 - DFSMS – Method to Empty a PDS/PDSE
- ? Why it Matters: Allows you to remove all members of a PDS or PDSE data set in a single command**

Ease of Use: PDSE Message Enhancements

- **Identify Corrupt PDSE in LNKLST**

- Currently, when a corrupt PDSE is detected in the link list during IPL, the system enters a wait state
- **New enhancement:** The system will issue a message identifying the corrupt PDSE prior to entering the wait state
 - *Success message (IGW700I): If the PDSE validation tool is not able to find anything wrong with a data set, and the exploiting application chooses to, a message indicating that no problems were found is printed on the console.*
 - *Global Errors (IGW701I): This message is issued when the PDSE validation finds an error with the overall structure of the data (i.e. a missing or doubly allocated page).*
- Allows the user to attempt to restore the corrupt PDSE, re-IPL the system and avoid taking a standalone dump to debug the problem

? **Why it Matters:** Better problem determination and error recovery

Ease of Use: OCE Enhancements

- **New DCBabend exit option when exit sets IGNORE option**
 - When an OCE determinant ABEND is detected and DCB ABEND exit selects the ignore option, the associated ABEND message is issued even though there is no abnormal termination of the task
 - Provide a new DCB ABEND exit ignore option to additionally bypass the associated determinant ABEND message
 - Eliminates the contradiction of externalizing an ABEND when the task does not actually abnormally terminate
- **Recovery for SMS DASD input when missing last volume externalized via IEC710I**
 - When reading an SMS DASD data set and message “IEC710I another volume expected” is issued, the job completes successfully even though not all the data is read
 - Recover by locating the next volume via the catalog to read. If recovery is not possible, then ABEND the task with ABEND 637-BC
 - Maintain data integrity by ensuring all data is read
- **New reason codes when FREE=CLOSE is not honored in IEC988I message**
 - FREE=CLOSE is requested (unallocate when the data set is closed) but is then bypassed for expected reasons which are not externalized
 - Provide a new message IEC988I to externalize the reason for not honoring the FREE=CLOSE request
 - Simplifies diagnosis of FREE=CLOSE failures
- **New reason codes for ABEND 413 and 637 when detected for missing or out of order tape volumes of a multivolume tape data set**
 - When the label anomaly installation exit requests an ABEND due to a missing or out sequence volume condition when reading a multivolume tape data set, there is no specific reason code that describes the anomaly
 - Provide a new ABEND reason code for each anomaly and provide a recovery option via a DCB ABEND exit
 - Provides better control at both the installation and application levels when missing or out of order tape volume conditions are detected when reading

? Why it Matters: Improved first failure data capture and diagnostics

z/OS® DFSMS™ V1.12 Highlights

(September 2010)



■ Application Integration

- SDM Support of ATTREXX Interface

- *For more details see Session 8048 What's New in DFSMSdss and System Data Mover*

Application Integration: SDM Enhancements



- **SDM Support of ATTREXX Interface**

- Currently, SDM has an unsupported REXX interface that has been used as a test tool for the ANTRQST API and has been used by a small set of customers.
- **New enhancement:** The System Data Mover (SDM) component will provide a REXX interface for many of the functions of the SDM programming interface (ANTRQST)
 - Provides interfaces to FlashCopy®, Global Mirror, z/OS Global Mirror (XRC), Metro Mirror (PPRC), and other misc. SDM services.
 - CLISTs will be provided in SYS1.SAMPLIB for the invocation of ANTTREXX.
 - Results from ANTTREXX will be in the form of messages indicating the success or failure of the request.
 - Complete documentation provided in z/OS V1R11.0 DFSMS Advanced Copy Services (SC35-0428-16).
- **SHARE Requirement MET:** SSMVSS07014
 - Provide a supported REXX interface for the System Data Mover

? **Why it Matters:** Provides advanced users with a better programming interface

z/OS® DFSMS™ V1.12 Highlights

(September 2010)



■ Optimization & Management Capabilities

- DFSMS IMBED/REPLICATE Removal
 - *For more details see Session 8048 What's New in DFSMSdss and System Data Mover*
- DFSMS Fast Reverse Restore Enhancements
 - *For more details see Session 8045 What's New in DFSMShsm*
- DFSMSrmm TS7700 Reporting Enhancements
 - *For more details see Session 8041 What's New in DFSMSrmm*
- SMS Storage Group Mgmt & Volume Selection Enhancements
- IDCAMS DCOLLECT Enhancements
- IDCAMS GDG Enhancements
- NFS Enhancements

Optimization & Mgmt Capabilities: IMBED/REPLICATE Removal

- **DFSMS IMBED/REPLICATE Removal**

- The creation of new VSAM data sets with IMBED and REPLICATE attributes has been unsupported since z/OS V1.3.
- **New enhancement:** DFSMSdss will convert any indexed VSAM data set with the IMBED or REPLICATE attribute to a data set without these attributes when they are restored using logical data set restore.
 - New messages
 - *ADR507I will be issued identifying the data set being converted*
 - *ADR508I, will be issued during logical data set dump, physical data set dump and physical data set restore when an indexed VSAM data set with the IMBED, REPLICATE or KEYRANGE attribute is processed*

? Why it Matters: For DFSMSHsm logical dumps, these old attributes are automatically removed. Reduces customers effort to remove these obsolete attributes.

Optimization & Mgmt Capabilities: DFSMShsm Enhancements



• DFSMS Fast Reverse Restore Enhancements

- Currently Disk recovery cannot be performed until the physical background copy is complete (hours after the logical backup was created) and disk recovery from a NOCOPY version is not supported.
- **New enhancement:** DFSMSdss and DFSMShsm will exploit the Fast Reverse Restore feature.
 - Enable DFSMShsm to flash back for disk recovery even though the background copy has not completed and enable fast replication disk recovery from a NOCOPY version
 - Fast Reverse Restore function will support the recovery of a
 - *Copy pool without waiting for physical background copy to complete*
 - *Copy pool from a disk version in either COPY or NOCOPY environment*
 - *Copy pool from space efficient volumes in NOCOPY environment*
 - Only supported for recovery of entire copy pool and source volumes cannot be in any other FlashCopy relationships
 - *Since the background copy does not complete, the backup version is invalidated and backup volumes are initialized as part of the recovery*
 - New SMS copy pool setting indicates whether Fast Reverse Restore is enabled
 - *YES indicates it is acceptable to recover a version before the background copy is complete and for the backup version to become invalidated*
 - *The copy pool setting is stored for each backup version at the time of backup*
 - *Recovery of the copy pool uses Fast Reverse Restore if the backup version is FCFRR-enabled (FlashCopy fast reverse restore) and flash back is possible*
 - A new DFSMShsm SETSYS parameter indicates whether extent or full-volume FlashCopy relationships are to be established between volume pairs when DFSMShsm invokes DFSMSdss to perform fast replication backup and recovery.

? Why it Matters: Enables DFSMShsm to FlashBack for recovery even though the background copy has not completed – could save hours!

Optimization & Mgmt Capabilities: DFSMSrmm Enhancements

- **DFSMSrmm TS7700 Reporting Enhancements**

- DFSMSrmm will help with reporting of data sets and logical volumes which are copy exported from a TS7700 virtualization engine.
 - Reports can be created either from the export list file of up to three copy exports, or from the information created from the TS7700 Bulk Volume Information Retrieval Function (BVIR).
 - The information about stacked volumes, and logical volume copies is retrieved from this input and merged with the information that the DFSMSrmm extract file X records has for the stacked volumes and the logical volumes.
- Sample report job located in SYS1.SAMPLIB(EDGJCEXP)
- Reports are provided sorted
 - by data set name
 - by logical volume serial number
 - by stacked volume serial number

? Why it Matters: Quick and easy way to create a list of data sets on the volumes copy exported.

Optimization & Mgmt Capabilities: SMS Enhancements



- **SMS Storage Group Mgmt & Volume Selection Enhancements**
 - Currently the HIGH Allocation/Migration Threshold value is 1-99%
 - No specific recommendation, since the optimal value will depend on the user's specific requirements; however, the default value SMS will use, if no specific value is provided by the user, is 85%.
 - Best practices suggest a value for HIGH Threshold below 90%.
 - As volume sizes increase, one percent of a volume represents an increasingly large amount of storage. For example, on a 223 GB volume, 1% is over 2 GB of storage.
 - **New enhancement:** Increase the HIGH Allocation/Migration Threshold maximum to 100%
 - In most cases, IBM recommends a high threshold value less than 100% for storage groups. This allows data sets to expand without an increased risk of encountering out-of-space abends.
 - The 100% specification is intended to be used to make more storage capacity available for storage groups that hold static data.
 - Large installations frequently have thousands of volumes and the candidate Storage Groups for a single allocation may also have hundreds if not thousands of volumes.
 - **New enhancement:** SMS processing of volume lists will be changed in a way intended to improve allocation performance for large volume lists.
- **? Why it Matters: With the new 223GB EAV size volumes, getting that last 1% is worthwhile!**
Also large volume list processing is improved

Optimization & Mgmt Capabilities: IDCAMS DCOLLECT Enhancements

- **DCOLLECT data class (DC) records will be updated to include information about all data class attributes:**
 - VSAM SPEED and REUSE
 - Tailored Compression
 - CICSVR Forward Log
 - RLS Greater Than 4K Cache
 - Block Size Limit SPE
 - Dynamic Volume Count
 - RLS 64 Bits Virtual
 - Scaling Constants for Tape Support use
 - Tape Control Unit Performance Segmentation
 - SMB VSP
 - Tape Encryption
 - CA Reclaim
- **SHARE Requirement MET:** SSMVSS053125
 - DFSMS: DCOLLECT Needs to Include More Fields

Optimization & Mgmt Capabilities: IDCAMS DCOLLECT Enhancements (con't)

- **Data set (D) records will be updated to include job names:**
 - In z/OS V1R11, new fields in the Format 9 DSCB added for the job name, step name and time of data set creation in mapping macro IECSDSL1.
 - **SHARE Requirement MET:** SSMVSS064955 - DFSMS: Add Jobname to Catalog Record When Data Sets Are Created.
- **Storage group (SG) records will be updated**
 - Include information about OAM Protect Retention and Protect Deletion settings provided in z/OS V1R11
- **SHARE Requirement MET:** SSMVSS053125 - DFSMS: DCOLLECT Needs to Include More Fields

? Why it Matters: Ensure Data Class information is captured and available

Optimization & Mgmt Capabilities: IDCAMS GDG Enhancements

- Currently, IDCAMS, when deleting entire generation data groups (GDGs), invokes DFSMSHsm recalls for any generation data sets that are migrated.
- ***New enhancement:* IDCAMS will call DFSMSHsm to delete such data sets without recalling them**
 - Expected to reduce processing time, particularly when one or more generation data sets have been migrated to tape.
- ***SHARE Requirement MET:* SSMVSS064933**
 - Catalog Support of HDELETE for DELETE GDG FORCE

? Why it Matters: Potentially saves lots of time / avoids wasted recalls

Optimization & Mgmt Capabilities : NFS Enhancements

- **NFS Server SMF Records for File Operations**
 - NFS Server needs to provide an additional SMF record for auditing and security reasons
 - An audit trail is required specifying the user who created, removed, or renamed the file on the NFS mounted file system
 - An audit trail is required for the creation, deletion, and renaming of MVS datasets and members
 - **New enhancement:** New SMF record: type 42 subtype 26 used to provide additional SMF recording for every CREATE/REMOVE/RENAME NFS operation in z/OS NFS server
 - SMF record will include Client information, the type of operation (create, remove, rename), and object descriptive information
 - For z/OS Unix objects, the filesystem name, device number, object name, inode, and parent inode information will be saved.
 - For MVS objects, the volume name, full dataset name, and member name (if appropriate) will be saved.

? Why it Matters: Provides customer the audit capability and enhances NFS serviceability

Optimization & Mgmt Capabilities : NFS Enhancements

- **NFS Server Cache Monitoring and Reporting**
 - NFS Server needs to provide an additional SMF record for auditing and security reasons
 - Customer requirement for the z/OS NFS Server to issue messages to the console to warn of impending legacy data buffer shortage problems, before the server cache is exhausted and an error is generated
 - **New enhancement:** New console messages to warn about z/OS Network File System Server legacy data buffers shortage
 - Provide the ability to track current legacy data buffers utilization and to change maximum legacy data buffers pool size and warning thresholds from the console
 - New console commands:
 - MODIFY nfserver, bufhigh=(xx,yy)
 - MODIFY nfserver, bufhigh
 - MODIFY nfserver, bufferusage
 - Use BUFHIGH parameter to specify a percentage of the storage limit for allowed data buffer space, use the system default, or disable monitoring, and change the specification using the MODIFY operator command
 - From the console the maximum legacy data buffers pool size can be only increased
 - If monitoring is enabled, messages are issued:
 - When buffer utilization exceeds the alert threshold, as it approaches the maximum allowed, and when it falls below the threshold

? Why it Matters: Easier to tune z/OS Network File System Server attributes for current workload

Optimization & Mgmt Capabilities : NFS Enhancements

- **Display the Accounting Statistics in the NFS Server**

- The nfsstat (/usr/lpp/NFS/nfsstat) shell program displays the accounting statistics of ALL the z/OS NFS Server(s) within a LPAR along with the accounting statistics of the z/OS NFS Client (previously supported)
- The accounting statistics are
 - The number of received Remote Procedure Calls (RPC),
 - The number of the received Network File System (NFS) Version2, or Version3 procedures, or Version4 operations
- The nfsstat program still interfaces with the started z/OS NFS Client, and it additionally interfaces with all the started z/OS NFS Server(s) to report the Server(s) statistics through the new and enhanced option:
 - -s: report Server(s) RPC and NFS statistics (new option)
 - -r: report Server(s) and Client RPC statistics (enhanced option)
 - -n: report Server(s) and Client NFS statistics (enhanced option)
 - -z: report Server(s) and Client RPC and NFS statistics and reset to zero
 - -4: report Server(s) and Client NFSv4 statistics (new option)
 - -3: report Server(s) and Client NFSv3 statistics (new option)
 - -2: report Server(s) and Client NFSv2 statistics (new option)

? Why it Matters: easier to tune and debug server/client interactions

z/OS® DFSMS™ V1.12 Highlights

(September 2010)



■ Availability

- DFSMSrmm Active and Queued Task Management
 - *For more details see Session 8041 What's New in DFSMSrmm*

Availability: DFSMSrmm Enhancements

- DFSMSrmm Active and Queued Task Management
 - There was no way to retrieve information about the DFSMSrmm subsystem requests and task status except the MODIFY operator command
 - ***New Enhancement:* New TSO subcommand, RMM LC STATUS**
 - Provides information about the DFSMSrmm subsystem, subsystem requests, and task status
 - *The information returned is very similar to the results of the operator QUERY ACTIVE command*
 - Issued using TSO, the RMM API, HLL API, and the Web Service
 - RMM dialog panel EDGPCC00 can be used to view RMM status and to use Hold, Release and Cancel commands against tasks

? Why it Matters: Enables storage applications to monitor and act on the available information.

z/OS® DFSMS™ V1.12 Highlights

(September 2010)

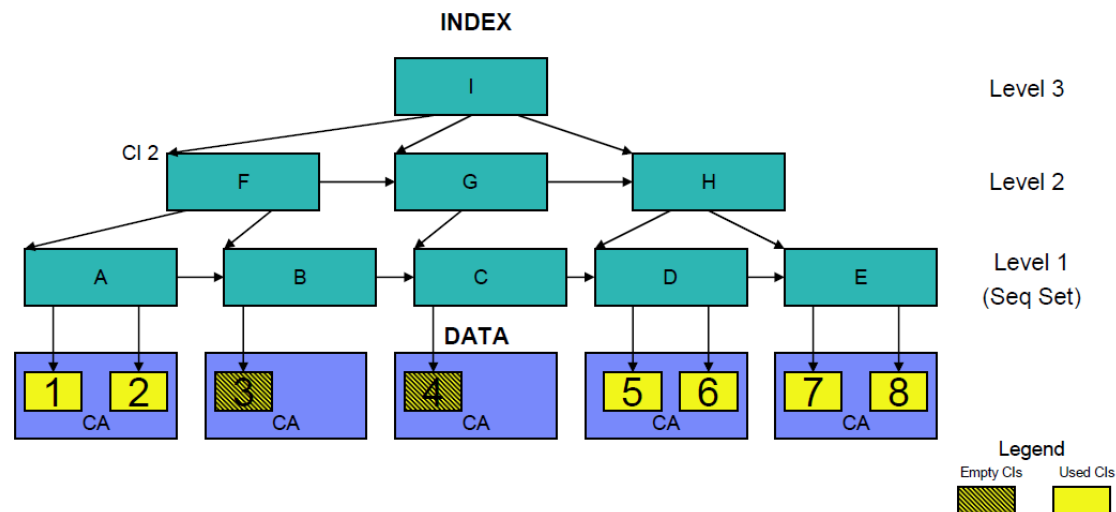


■ Scalability, Performance, & Constraint Relief

- CA Reclaim
 - *For more details see Session 8054 z/OS 1.12 New Feature - VSAM CA Reclaim*
- EAV Additional Data Set Support
- DFSMSrmm EAV Exploitation
 - *For more details see Session 8041 What's New in DFSMSrmm*
- DFSORT Memory Object Sorting
- DFSMS Support for XTIO and Uncaptured UCBs
- OAM Enhancements
- Extended Addressable Catalogs
- Catalog CAS Contention Detection
- VSAM RLS Striping
- DFSMSdss Larger Block Support for DUMP, COPYDUMP, and RESTORE
 - *For more details see Session 8048 What's New in DFSMSdss and System Data Mover*
- DFSMSHsm Dump Tape Recovery Enhancements
 - *For more details see Session 8045 What's New in DFSMSHsm*
- DFSMSHsm Space Management Performance
- DFSMSHsm Dump Stacking

Performance: VSAM KSDS CA Reclaim

- 24/7 Availability
 - This function circumvents VSAM KSDS files from being taken offline to reclaim unused space. When records are erased from a VSAM KSDS, the empty CAs are not often reused, resulting in a “fragmented” data set
 - Applications that use VSAM key-sequenced data sets (KSDS) can benefit from improved performance, minimized space utilization, and improved application availability.
 - For example, IBM system tests indicate performance of dedicated VSAM workload could improve by up to 44%, in addition to helping avoid outages that used to be required to defragment and reorganize this data.
 - Performance improvements are anticipated for many applications using CICS®, VSAM, VSAM RLS, IMS™ VSAM, and Catalog processing.



? Why it Matters: Avoids planned down time for VSAM KSDS reorgs

CA Reclaim Usage and Invocation



- CA Reclaim is disabled by default on a system level, but is enabled by default for all KSDSs without having to redefine the data set
- New system level parameter in SYS1.PARMLIB(IGDSMSxx):
CA_RECLAIM(NONE | {DATACLAS|DATACLASS})
 - When system level NONE is specified (default), CA Reclaim is disabled for all KSDSs on this system
 - When system level “DATACLAS|DATACLASS” is specified, the setting of new data class parameter: CA_Reclaim(Y/N) is used when the data set is defined to determine whether the data set is eligible for CA reclaim
- CA Reclaim can be enabled/disabled on a system level
 - **SETSMS CA_RECLAIM(NONE | {DATACLAS|DATACLASS})**
 - Takes effect immediately without an IPL
 - SYSPLEX – must route using RO *ALL, SETSMS
 - PARMLIB option CA_RECLAIM() or SET SMS=xx or RLS VSAM address space recycle
 - Takes effect after the IPL or recycle
- *For more details see Session 8054 z/OS 1.12 New Feature - VSAM CA Reclaim*

CA Reclaim Usage and Invocation

- When system level “DATACLAS|DATACLASS” is specified, the setting of a new data class parameter is used to determine whether the data set is eligible for CA reclaim
- CA Reclaim is enabled by default for individual KSDSs
 - ISMF data class keyword CA_Reclaim(Y/N)
 - DATACLAS attribute processed only when a catalog entry is defined (DEFINE, redefine, MIGRATE, RECALL)
 - Value is saved in the catalog regardless of the SYS1.PARMLIB setting
 - If DATACLAS is changed afterward, the change will not take effect unless the KSDS is re-DEFINED (or ALTER)
- CA Reclaim can be enabled/disabled on a data set level
 - **ALTER RECLAIMCA|NORECLAIMCA**
 - The ALTER specification for a KSDS changes the resultant cataloged CA reclaim attribute obtained from the data class
 - It will take effect at the first OPEN following the CLOSE of all open ACBs against the data-set control block structure

Scalability: What is an EAV?


- What is an Extended Address Volume (EAV)?
 - *A volume with more than 65,520 cylinders*
 - *Size limited to 223 GB (262,668 Max cylinders)*
 - *Supported in z/OS V1R10 and higher*



Breaking the 65,520 cylinder limit

EAS Eligible data set sets in z/OS

- **EAS Eligible:** A data set on an EAV that is eligible to have extents in the extended addressing space and described by extended attribute DSCBs (format 8/9)
- Can reside in track or cylinder-managed space
- SMS-managed or non-SMS managed
- Any data set type can reside in track-managed space
- Data set types supported
 - **VSAM data types (KSDS, RRDS, ESDS and linear)**
 - *This covers DB2, IMS, CICS, zFS and NFS*
 - *CA sizes: 1, 3, 5, 7, 9 and 15 tracks*
 - **Sequential (Extended Format)**
 - **Sequential (Basic and Large Format)**
 - **Direct (BDAM)**
 - **Partitioned (PDS, PDSE)**
 - **Catalog (VVDS and BCS)**
 - **All data sets used by DFSMSrmm (journal and dynamically allocated temporary files)**
 - *One exception is the RMM CLIST data set when created automatically by SEARCH subcommand processing.*



? Why it Matters: Scalability for customers constrained by 4 character UCB's

Non-EAS eligible data set list

- A data set that may exist on an EAV but is not eligible to have extents (through create or extend) in the extended addressing space or have extended attribute DSCBs (format 8/9)
- The following data sets are not EAS eligible:
 - VSAM data sets with incompatible CA sizes
 - VTOC (continues to be restricted to within first 64K-1 tracks)
 - VTOC index
 - Page data sets
 - VSAM data sets with IMBED or KEYRANGE attributes
 - HFS file system
 - XRC Control, Master or Cluster non-VSAM data sets
 - *State data set EAS eligible in z/OS V1R12*
 - *Journal data set EAS eligible in z/OS V1R11 and V1R12*

Scalability & Constraint Relief: DFSORT Memory Object Sorting



- **Memory object sorting**

- Preferred path is memory object (storage above the bar) sorting which sorts all records in this storage
 - Only limited by central storage (uses more CPU)
- Second in preference is to use hiperspace sorting which writes sets of sorted records to hiperspace storage (intermediate work space) then merges the sets of records
 - Capacity is limited to 32G (uses less CPU)
- The hiperspace path will now be preferred
 - This path has also been modified to allow use of memory objects or hiperspace as workspace
 - Central storage limit with memory objects is 64GB
- New MOWRK installation parameter created to control use of memory objects as work space.
 - Default is MOWRK=YES
- New DFSORT messages indicate whether memory object storage was used as main storage or work storage

? Why it Matters: Reduced DFSORT CPU time and contention and better balancing of central storage usage across concurrent sorts

Scalability & Constraint Relief: XTIOOT and Uncaptured UCBs

- **DFSMS Support for XTIOOT and Uncaptured UCBs**
 - In many z/OS environment, some workloads require an increasing number of open data sets.
 - BSAM, QSAM, and BPAM (basic and queued sequential, and basic partitioned access methods), OPEN/CLOSE/EOV, CVAF, DADSM, EXCP, and RACF processing support the use of an extended task I/O table (XTIOOT) with uncaptured UCBs, and support data set association blocks (DSABs) above the 16 MB line.
 - This is expected to allow more data sets to be allocated by an address space and to provide virtual storage constraint relief for programs using large numbers of DASD and tape data sets.
 - Exploiters include:
 - SNAP/SNAPX services and dump processing (including that for SVC, SYSABEND, SYMDUMP, and SYSUDUMP)
 - AMASPZAP
 - The Program Management Binder
 - TSO/E
 - DFSORT
 - *Programs that invoke DFSORT, ICETOOL, or ICEGENER can dynamically allocate input, output, and work data sets using the options for XTIOOTs, uncaptured UCBs, and DSABs above 16 MB*

? Why it Matters: VSCR (virtual storage constraint relief) and to support more than about 3200 allocated data sets in an address space.

Performance: OAM Enhancements

- **OAM Object Storage & Retrieval (OSR) Enhancements**
 - OAM is planned to provide API support for the Object Storage and Retrieval function (OSR) to run in a CICS threadsafe environment.
 - This is intended to allow exploiters to take advantage of the improved multitasking and throughput capabilities provided by threadsafe programming.
- **Volume Recovery Improvements**
 - Additionally, the Volume Recovery utility will be designed to improve performance in certain situations when recovering object data stored on optical and tape media.
 - Improvements are expected to be most noticeable when recovering a backup volume containing objects with primary copies in a large number of different collections on a large number of different volumes.

? Why it Matters: Exploiting CICS performance enhancements

Scalability: Extended Addressable Catalogs

- **DFSMS will allow catalogs to be defined with extended addressability (EA)**
 - This will make it possible to define and use Integrated Catalog Facility (ICF) Basic Catalog Structures (BCS) with EA, allowing catalogs larger than 4 GB.
- **SHARE Requirement MET: SSMVSS08006**
 - DFSMS – ICF Catalog Size Relief
- **Catalog CAS Contention Detection**
 - The Catalog address space (CAS) will be designed to check for SYSZTIOT enqueue contention periodically.
 - Based on an interval you specify and the reason for contention, CAS will be designed to write a logrec record and a notification message when tasks have waited longer than the specified interval and contention checking is active.
 - A new MODIFY CATALOG,CONTENTION command will allow you to specify a different interval than the 10-minute default or to disable CAS contention detection.
 - Intended to warn about tasks that take excessive time to complete, or never complete, from affecting Catalog performance.

? Why it Matters: Big (>4GB) Catalogs (more efficient use of storage resources; Better diagnostics)

Performance: VSAM RLS Striping

- **VSAM record level sharing (RLS) will support striped data sets.**
 - This will be designed to bring the benefits of VSAM striping, such as allowing single application requests for records in multiple tracks or control intervals (CIs) to be satisfied by concurrent I/O requests to multiple volumes.
 - Using striped data sets can result in improved performance by transferring data at rates greater than can be achieved using single I/O paths.

? Why it Matters: Performance improvement

Performance: DFSMSdss Enhancements

- **DFSMSdss Larger Block Support for DUMP, COPYDUMP, and RESTORE**
 - DFSMSdss will use larger blocks when possible for DUMP, COPYDUMP, and RESTORE operations, and to support Extended Format Sequential dump data sets on DASD for DUMP, RESTORE, and COPYDUMP.
 - The maximum supported block size for DASD is 32 760 bytes.
 - *DFSMSdss DUMP and COPYDUMP already support the entire range of valid block sizes when the output is DASD.*
 - The DUMP and COPYDUMP command will increase the default block size from 65 520 bytes to a maximum of 262 144 bytes when the output is to a tape device.
 - *DFSMSdss will choose the optimum block size for the particular tape device if BLKSZLIM was not specified.*
 - The use of larger block sizes is intended to improve performance for these operations and using Extended Format dump data sets is intended to support striping and compression.
 - DFSMSShsm functions that will exploit this enhancement
 - BACKVOL/FRBACKUP DUMP
 - AUTOMATIC DUMP
 - RECOVER/FRRECOV FROMDUMP

Performance: DFSMSHsm Enhancements



- **DFSMSHsm Dump Tape Recovery Enhancements**

- Currently, an entire copy pool can only be recovered from disk. DB2 or native DSS has to be used to recover from tape.
- **New enhancement:** DFSMSHsm will support parallel processing for recovery from dump tape volumes when the dumps reside on multiple tape volumes and multiple tape drives are available.
 - Allows recovery of entire copy pool from tape.
 - Multiple concurrently queued volume recovery requests on the same tape will be processed with a single tape mount.
 - Specify up to 64 concurrent tasks be used to help speed recovery processing.

? Why it Matters: Potentially could save a lot of time in large recoveries

- **DFSMSHsm Dump Stacking**

- The DFSMSHsm DUMP function used to copy source disk volumes to a target tape volume will allow up to 255 source volumes to be dumped to a single tape volume, an increase from the prior limit of 99.

? Why it Matters: Allows you take better advantage of large capacity tape cartridges.

Performance: DFSMShsm Enhancements

- **DFSMShsm Space Management Performance**
 - Intended to address performance issues due to vertical growth (increasing the number of data sets on individual disk volumes - EAVs) and horizontal growth (increasing number of data sets across more disk volumes).
 - Primary Space Management (PSM) or Interval Migration (IM) may not finish within expected timeframes.
 - ***New enhancement:*** New option to allow Primary Space Management, Interval Migration, and the MIGRATE PRIMARY command to overlap volume preprocessing with volume data movement in order to reduce elapsed time.

? Why it Matters: Potentially could save a lot of time in large DFSMShsm environments

z/OS® DFSMS™ V1.12 Highlights

(September 2010)



- **Security**
 - NFS Enhancements

Security: NFS Enhancements


- **NFS Password Phrase Support**

- In SAF or SAFEXP mode, client users are required to issue an mvslogin with the RACF User Id and the password
- Currently, NFS Server has 8 characters limit on password for the RACF users
- ***New enhancement:*** The NFS Server supports password phrases up to 100 characters in length for mvslogin, in addition to existing support for passwords up to 8 characters long
 - Requires password phrase support from RACF, or another external security manager
 - User can authenticate with either password or password phrase
 - New message is posted on the user console when user tries to change from current password to new password phrase and vice versa, which is not supported by RACF

? Why it Matters: lintended to allow you to migrate to password phrases, which offer a much larger name space than passwords

IBM System Storage™ & System z®

A winning combination



SHARE
Technology • Connections • Results

Collaborating to support the growth and protection of mission critical information

Heterogeneous Enterprise Disk Systems

- IBM System Storage™ DS8700
- IBM System Storage™ Easy Tiers
- IBM System Storage™ Advisor Tool
- DS8000 z/OS Distributed Data Backup
- System z Discovery and Auto-Configuration (zDAC)
- Softek Transparent Data Migration Facility (TDMF) for z/OS, v5.2
- Disk Encryption
- Solid State Drives
- Remote Pair FlashCopy
- High Performance FICON + Multi Track
- z/OS Global Mirror Incremental Resync
- Extended Address Volumes
- Basic HyperSwap™
- Extended Distance FICON®
- z/OS Global Mirror enabled for zIIP
- PAV and HyperPAV
- PPRC Manager and FlashCopy Manager
- Geographically Dispersed Parallel Sysplex™ (GDPS®)



SAN Volume Controller (z/Linux)

- SAN Volume Controller virtualization appliance

SHARE in Boston

Enterprise Tape Systems

- IBM System Storage™ TS7680 ProtecTIER® Deduplication Gateway for System z®
- IBM Virtualization Engine TS7700 4 site support
- IBM Virtualization Engine™ TS7700 R1.7
- IBM System Storage TS3500 Tape Library High Density Frames
- IBM System Storage™ Tape Encryption



Sources for more information

- Information about [DFSMS](#) and components
 - <http://www-03.ibm.com/systems/storage/software/sms/index.html>
- Information about [DFSORT](#)
 - <http://www-01.ibm.com/support/docview.wss?rs=0&uid=isg3T7000077>
- Information about [IBM TotalStorage Productivity Center for Replication](#)
 - <http://www-03.ibm.com/systems/z/advantages/resiliency/eventdriven/tpc.html>
- Information about [IBM System Storage Disk](#) systems
 - <http://www-03.ibm.com/systems/storage/disk/enterprise/index.html>
- Information about [IBM System Storage Tape](#) systems
 - http://www-03.ibm.com/systems/storage/tape/?cm_re=masthead-_-products-_-stg-tape
- Additional Information
 - [Redbooks](#)
 - <http://www.redbooks.ibm.com/>
 - z/OS V1R8 DFSMS Technical Update (SG24-7435-00)
 - [Techdocs](#)
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs>



Thank You!

Additional Material

OAM D-APAR Enhancements

- **OA29425 - HASH Table Performance (March 2010)**
 - Converts internal OAM volume search routines to hash algorithms
 - Potentially provides performance improvements for installations that have large number of tape and/or optical volumes managed by OAM object support
 - Volumes defined in OAM's TAPEVOL and VOLUME tables in DB2
- **OA25408 - OAM OSREQ QUERY Wildcard Enhancement (April 2009)**
 - REQ# MR0602084143
 - New wildcard characters '%' and '_' are added to OAM's OSREQ Query function
 - Support enhances the generic search (wildcard) support for OSREQ QUERY requests
 - new
 - DB2 cursor is added to take advantage of the following DB2 wildcard characters using the "LIKE" predicate as described in the DB2 SQL reference SC18-9854-xx
 - NOTE: new OSREQ wildcards will not work with the OSREQ TSO/E command processor (OAM IVP TSO interface)
- **OA25463 - OSREQ Update Collection Enhancement (Feb 2009)**
 - REQ#MR0602083050
 - Previous to this support, in order to change the default storage class and management class names associated with an OAM collection, installations need to run a manual procedure
 - With this support, OAM provides a utility that will change the management class
 - and/or storage class defaults associated with a given collection.

SDM RAS Enhancements

- The following RAS line items are included in the R12 changes to SDM:
 - **Service Virtual Machine (SVM) enhancements**
 - From a serviceability and new function perspective, need to document how to build SVM, to create new executables that match the source, to correct several known defects and deficiencies (one of which prevents ASID reuse), and to address FIN APAR KFI0535/OA26336, which causes a rare ABEND66D.
 - **Allow XRC address space IDs (ASIDs) to be reusable**
 - This enhancement, which is dependant on the SVM change noted above, will enable ASID reuse for the ANTAS000, ANTAS0xx and ANTCL0xx (but not ANTMMAIN) address spaces.
 - **Add TERTIARY parameter to the XRECOVER command**
 - This capability is being provided to enhance usability for customers and to prevents the customer from mistakenly recovering from the original secondaries and destroying the mirror that they belong to.
 - **KFI0490/OA04021 High CPU usage in sessions with large number of PND/CPY volumes**
 - This enhancement will reduce CPU utilization for environments with large numbers (e.g. 1000s) of volumes in CPY and PND state
- *For more details see Session 8048 What's New in DFSMSdss and System Data Mover*

DS8000 Business continuity Leadership for zEnterprise™ System

Availability to support mission-critical applications on System z



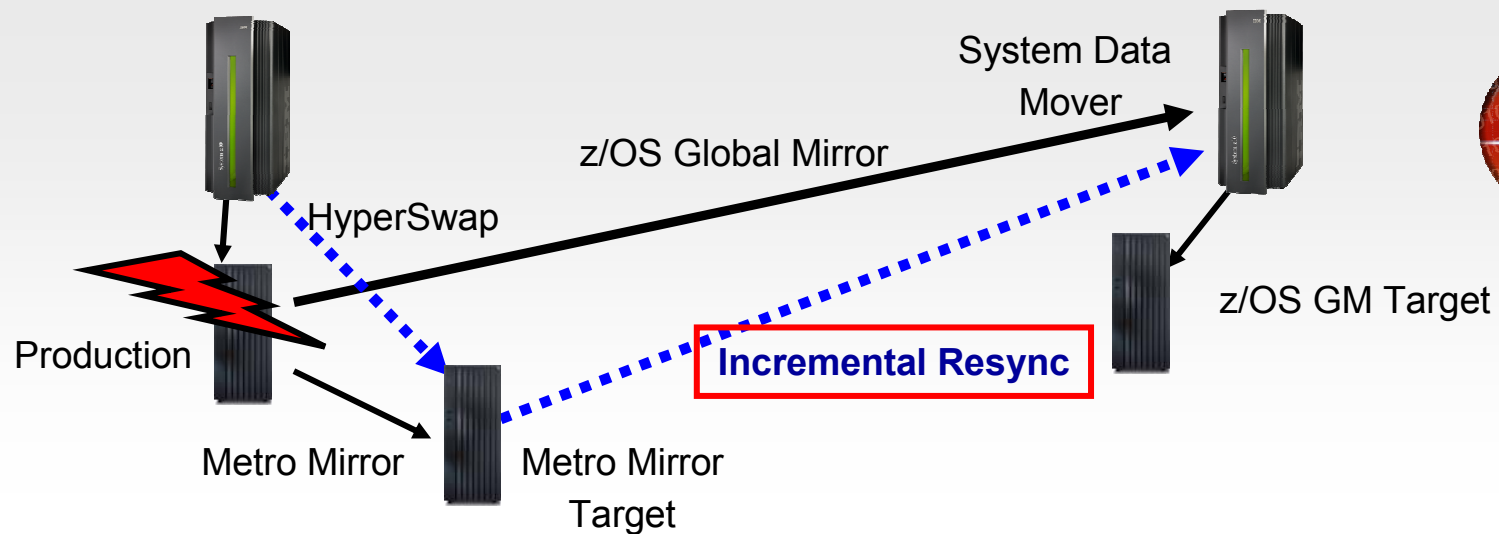
Advanced capabilities

- Designed for continuous operations with no single point of failure and hot-swappable components
- Data is RAID protected to guard against drive failure
- Outstanding copy and remote mirroring enhancements for mainframe environments
 - Data Set FlashCopy, Metro Mirror (PPRC), z/OS Global Mirror, z/OS Metro/Global Mirror
 - z/OS Metro/Global Mirror Incremental Resync - for more efficient resiliency
 - z/OS Global Mirror Multiple Reader – parallel processing for higher performance
 - z/OS Basic Hyperswap support – for economical, single-site resiliency
 - z/OS Global Mirror multiple sessions support – provides granular support
 - z/OS Global Mirror enabled for System z zIIP engine – cost effective replication on System z
- Fully-integrated HA/DR solutions for quick and reliable disaster recovery
 - HyperSwap, Globally Dispersed Parallel Sysplex™ (GDPS®)
- **Did you know?**
 - Synchronous data replication has high bandwidth utilization and supports up to 300km+ distance
 - Asynchronous data replication (long distance) is designed for data currency within 3-5 seconds and has self-managed consistency groups



z/OS® Metro/Global Mirror (XRC) Incremental Resync

Accelerates resynchronization for three-site remote mirroring for GDPS



- **Incremental Resync swaps zGM session volumes & reconnects to remote volumes on a HyperSwap**
- **Enables High Availability GDPS HyperSwap and Disaster Recovery... swap is automated, seamless and FAST**
- **Reduces amount of data transmitted – only sends changes from Metro Mirror Target to z/OS Global Mirror Target after HyperSwap**
 - Uses IBM System Storage Metro Mirror, z/OS Global Mirror (XRC) and GDPS
 - Production and Metro Mirror copy can be at different sites or in the same site for local/nearby HA/ DR and out of Region DR
 - Note: Three site mirroring with incremental resync using IBM Metro/Global Mirror (PPRC) for modular (UNIX®, Windows®, etc.) and System z is already available and is supported by TPC for Replication

Remote Pair FlashCopy®




Customer Value



More effective two-site business continuity with FlashCopy and Metro Mirror

- Helps improve data synchronization when a FlashCopy target is also a Metro Mirror source
- Prior to introduction of Remote Pair FlashCopy
 - The Metro Mirror pair could become temporarily “unsynchronized” and hinder a HyperSwap operation
 - This became problematic if client would perform many FlashCopies throughout the day
- Remote Pair FlashCopy now keeps Metro Mirror primary and secondary volumes synchronized during all FlashCopies
- Remote Pair FlashCopy can be combined with:
 - Incremental FlashCopy
 - Background copy or no background copy
 - Nocopy to copy
 - FlashCopy consistency groups
 - FlashCopy of open devices using a CKD access device
 - Dataset Level Flashcopy
- Remote Pair FlashCopy will GA on April 24 (Plant/Field)
- No charge for Remote Pair FlashCopy

Trademarks and Disclaimers



SHARE
Technology • Connections • Results

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*	FICON*	Lotus*	System Storage
DB2*	FlashCopy*	MQSeries*	Tivoli*
DFSMS	GDDM*	Multiprise*	TotalStorage*
DFSMSdftp	GDPS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMSHsm	HiperSockets	Parallel Sysplex*	VM/ESA*
DFSMSmm	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	zEnterprise
			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.