



DFSMS™: Latest and Greatest

Barbara McDonald
IBM - bawhite@us.ibm.com

March 3, 2015
Session 17107



SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

Permission is granted to SHARE, Inc. to publish this presentation slide in the SHARE Inc. proceedings. IBM retains the right to disavow copies of this presentation if however it occurs.



© Copyright IBM Corp. 2015

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.

Connect with IBM z Systems on social media!

Subscribe to the new [IBM Mainframe Weekly](#) digital newsletter to get the latest updates on the IBM Mainframe!

LinkedIn
System z Advocates**
IBM Mainframe-Unofficial Group
IBM System z Events
Mainframe Experts Network
SHARE

Facebook
IBM System z**
IBM Master the Mainframe Contest
IBM Destination z
SHARE Inc.

Twitter
IBM System z**
IBM System z Events
Destination z
SHARE

Blogs
IBM Mainframe Insights**
Millennial Mainframer
#MainframeDebate blog
SHARE blog
IBM Destination z
DancingDinosaur
Bob's Mainframe Blog

Google+

IBM Evangelists

- Willie Favero- Twitter [@wfavero](#) - [Getting the Most out of DB2 for z/OS](#)
- Robert Catteral DB2 expert - [Robert's DB2 blog](#)
- Ed Woods Tivoli expert
- Steven Dickens Twitter [@StevenDickens3](#) - [MainframeDebate blog](#)
- Eberhand Pasch, Linux R&D [Linux and Mainframe blog](#)
- Matt Whitbourne Twitter: [@inventorMatt](#)
- Christopher Spaight Twitter: [@Cspaight](#)
- Mary Hall Twitter [@SystemsandTech](#), Blog, [Mainframe Insights](#)

YouTube
IBM System z**
Destination z

tumblr
IBM Mainframe50

Include the hashtag #mainframe and #mainframe50. Also see #z13 and #IBMz

** Indicates official IBM System brand owned accounts, managed by Carly Exum (STG) and Pratin Ashtekar (SWG)

✓ **SHARE**

Internal System z SMEs that are also social.

- Willie Favero- Twitter [@wfavero](#) - [Getting the Most out of DB2 for z/OS](#)
- Robert Catteral DB2 expert - [Robert's DB2 blog](#)
- Ed Woods Tivoli expert - [Tivoli with a z blog](#)
- Eberhand Pasch, Linux R&D- [Linux and Mainframe blog](#)
- [Matt Whitbourne Twitter: @inventorMatt](#)

IBM Social Evangelists

- Christopher Spaight [Twitter: @Cspaight](#)
- Mary Hall Twitter [@SystemsandTech](#), Blog, [Mainframe Insights](#)
- Steven Dickens Twitter [@StevenDickens3](#) - [MainframeDebate blog](#)

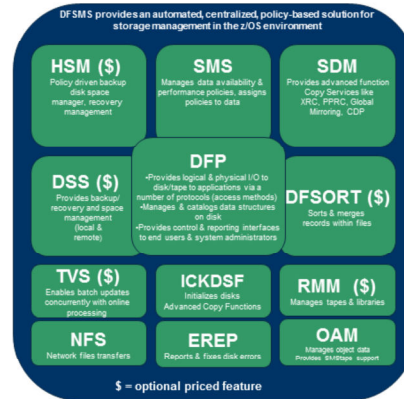
DFSMS™: Providing System Managed Storage on z/OS®

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



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

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

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



- How and where a customer will choose to store this data depends on many things: type of data type, quantity, cost, availability, etc.
- With "data" being key to the System z strategy and DFSMS playing a critical role in the storage and management of that data, the DFSMS strategy addresses the explosive growth and management of customer data.
- In addition to the continuous work in response to customer requirements around growth in the areas of availability, scalability and performance, we see additional growth opportunities...such as data protection through continued zCDP initiatives as opportunities in the area of data archiving.
- The DFSMS strategy also addresses the management of that data through storage simplification and synergy initiatives involving SWG (DB2, IMS, CICS, and Content Management) as well as Tivoli and STG.

Data Facility Product (DFP)

- Provides storage, data, program, and device management functions and DFSMS Copy Services capabilities

Hierarchical Storage Manager (HSM)

- An optional feature providing backup, recovery, migration, and space management functions

Data Set Services (DSS)

- An optional feature providing data movement, copy, backup, and space management functions

DFSMS Optimizer

- Optional feature of DFSMS that provides performance and management class analysis capabilities and real-time monitoring and automation of HSM. To be replaced by Tivoli Storage Optimizer (TSO).

Removable Media Manager

- An optional feature providing management functions for removable media such as tape cartridges and 3420 reels

Network File System (NFS)

- Provides a way to optimize performance and efficiency in a distributed network while still capitalizing on the capacity, security, and integrity of z/OS.

DFSORT

- Separate complementary product that provides high performance, sorting, merging, copying, reporting, and analysis of data.

Transactional VSAM Services

- Optional priced feature that enables batch jobs and CICS online transactions to update shared VSAM data sets concurrently.

System Data Mover

- DFSMSdftp through its System Data Mover Component provides DFSMS Copy Services which include Extended Remote Copy (XRC) and Concurrent Copy

Device Support Facilities (ICKDSF)

Error Recovery Executive Program (EREP)

z/OS® DFSMS™ V2.2 Preview Highlights (January 2015)



- **DFSMSshm**
 - DFSMS Storage Tiers Enhancements
 - Fast Replication Enhancements
 - DFSMSshm / DFSMSdsszEDC Exploitation
- **Catalog / IDCAMS**
 - Support for > 255 GDGs
 - MODIFY CATALOG Security Enhancement
 - Catalog RAS Enhancements
 - IDCAMS VERIFY RECOVER
 - IDCAMS RAS Enhancements
- **SMS**
 - Space Constraint Relief
 - RAS Enhancements
 - Storage Group Space Alert Messages
- **Access Methods**
 - RLS Index Level Locking
 - Secondary Space Reduction
- **Open/Close/EOV**
 - Dynamic Exits
 - RAS Enhancements
- **DADSM/CVAF**
 - RAS Enhancements
- **DFSORT**
 - zHPF Exploitation
 - Functional Enhancements
- **DFSMS / Storage Synergy**
 - Multi-Target PPRC
 - IBM zHyperWrite™
 - Easy Tier Application Hints
 - XRC Workload Write Pacing
 - Multiple Incremental FlashCopy

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Additional Sessions to Note



Session 16725: What's New with OMEGAMON V5 Family, **Monday 11:15AM**

Session 17011: OMEGAMON XE for Storage News and Tips, **Tuesday 11:15PM**

Session 16896: IBM z Systems z13 and DS8870 I/O Enhancements, **Tuesday 3:15PM**

Session 16957: PDSE Version 2 Member Generations: Practical User Applications, **Wednesday 10AM**

Session 16956: PDSE Nuts and Bolts, **Thursday 8:30AM**

Session 16955: PDSE Best Practices, **Thursday 10AM**

Session 16958: PDSE Advanced Diagnostics and Recovery, **Friday 10AM**

Session 16786: Employing External Retention Management Techniques in a DFSMSrmm Environment, **Tuesday 3:15PM**

Session 17135: DFSMS Free for All, **Wednesday 5:45PM**

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

7



Note: Take count for Project Dinner

z/OS DFSMS Highlights



- **DFSMSshsm**
 - **V2.2 Preview**
 - DFSMS Storage Tiers Enhancements
 - DFSMSshsm Fast Replication Enhancements
 - DFSMSshsm / DFSMSdss zEDC Exploitation

Session 17108: DFSMS Exploitation of z/OS zEnterprise Data Compression (Revised from Pittsburgh!), **Monday 3:15PM**

Session 17101: DFSMSshsm CDS Deep Dive, **Tuesday 11:15AM**

Session 17102: Transitioning to Transitions, **Wednesday 3:15PM****

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

Session 17144: Disk Tiering Solutions from IBM, EMC and HDS, **Tuesday 3:15PM**

Session 17091: IBM EasyTier for z/OS Deep Dive, **Tuesday 4:30PM**

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



DFSMSHsm



- **DFSMS Storage Tiers**
 - Today, DFSMS provides policy-based...
 - Data Creation
 - Backup / Recovery Management
 - Space Management
 - Expiration
 - **With V2.1**, DFSMS Storage Tiers introduced automated, policy-based space management that moves SMS-managed data from tier to tier *within* the Primary (Level 0) Hierarchy, allowing for policy-based management of **“active”** data.
 - Automated movement provided via the existing DFSMSHsm Space Management function; movement is referred to as a 'Class Transition'
 - Space management functions will determine if a data set is eligible to be transitioned, based on management class criteria. SMS Management Class provides the Class Transition policies.
 - Data remains in its original format and can be immediately accessed after the movement is complete.
 - Enhanced support for DB2, CICS and zFS data
 - Open data temporarily closed to enable movement

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

9



The Space Management Environment is evolving

- Typical configurations have changed to leave data on Level 0 longer and then migrate directly to ML2
- When ML2 is a VTS, the VTS disk cache replaces the ML1 tier
 - Eliminates MIPS required for software compression to ML1
 - Eliminates DFSMSHsm ML1->ML2 processing

•Invocation:

- When a volume is selected for space management processing due to being over threshold, in addition to existing expiration and migration checking, space management functions will determine if a data set is eligible to be transitioned, based on management class criteria
- Existing Space Management functions: Primary space management (expiration, release unused space), Migration: On Demand Migration (ODM) – performs space mgmt. as soon as the volume goes over threshold vs Interval Migration (IM) - hourly

There is a distinction between using Migrate/Recall and Class Transitions

- When a data set is recalled, it will be returned to the class of storage as directed by the ACS routines, which would typically be higher than where a data set would be transitioned to
- When a data set transitions to a lower class of storage, it will remain there until it is transitioned again or until it migrates

DFSMSHsm



- **DFSMS Storage Tiers Enhancements**

- Today, the Class Transition function is limited to automatic space management processing.
- **New enhancement:** Enable Class Transitions to be initiated with the various Migration commands.
 - The various Migrate commands are enhanced to support class transitions at the data set, volume and storage group level.
 - The default behavior is to perform **both** migration and transition processing for VOLUME and STORAGEGROUP operations.
 - » New keywords: **BOTH, MIGRATIONONLY, TRANSITIONONLY**
 - The default for MIGRATE DATASET is to perform a migration. The TRANSITION keyword indicates that a transition should be performed.

? Why it Matters: Storage Tiers in general allows users to better align storage costs with changing business value and minimize the TCO of System z data. The new enhancements extend the use of tiering and improve the usability of the function.



- Data sets eligible for **both** a transition and migration will be migrated

DFSMSHsm



- **DFSMS Additional Storage Tiers Enhancements**
 - Today, for “classic” migration, command initiated volume migration is single threaded. This severely limits the throughput.
 - **New enhancement:** New **STORAGEGROUP** keyword to initiate migration for all volumes within a storage group.
 - Primary Space Management processing will be performed for the storage groups, as opposed to the standard volume-by-volume processing performed by `MIGRATE VOLUME`.
 - **Use Case:** Clients create DB2 image copies to a specific storage group and then invoke DFSMSHsm `MIGRATE` to move the copies to tape. This enhancement enables a single command to process all of the volumes in the storage group in **parallel**.
- **? Why it Matters:** Significantly improve the throughput for command initiated space management and improve usability.



Invocation:

- `MIGRATE STORAGEGROUP(sgname, sgname, ...) ...`
- Up to 30 storage group may be specified

DFSMSHsm



- **DFSMS Additional Storage Tiers Enhancements**

- Today, moving data to newly defined disk volumes within a storage environment can be manually intensive and cumbersome.
- **New enhancement:** Enable the DFSMS Class Transition function to move data laterally to the same tier of storage, in addition to the ability of moving data to different tiers storage.
 - **MOVE** keyword added to the MIGRATE DSNAMES, VOLUME and STORAGEGROUP commands.
 - Every data set will be processed, regardless of management class policy or threshold, and ACS routines will be invoked to determine the new storage class and/or storage group.
 - **Use Case:** Move DB2 data from existing smaller volumes to the new larger, newly defined EAVs.

? Why it Matters: Simplify the task of migrating data to newly defined disk volumes.

DFSMSHsm



- **DFSMSHsm Fast Replication Enhancements**

- Dump (Command, Auto, Fast Replication) commands today must be processed on the same DFSMSHsm host that initiated the request.
 - The resources of a single DFSMSHsm host do not scale to the large copy pool sizes that need to be processed.
- **New enhancement:** Distribute dump work across all hosts in a sysplex-wide Common Dump Queue (CDQ).
 - Introduces an architecture to distribute workload across host members in an HSMplex, and the new architecture allows for flexible configurations.
 - Hosts that both receive and process requests from the group
 - Hosts that just process requests
 - Hosts that receive requests, do but not process them
 - Invoked and controlled via SETSYS commands.

? Why it Matters: Improves overall throughput by distributing the workload across the sysplex instead of concentrating it on a single LPAR.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

13



zCDP for DB2 Overview:

- Continuous Data Protection for DB2 is a joint solution between DFSMS and DB2 (BACKUP SYSTEM utility) in which Point-in-Time backup copies managed by DFSMS can be combined with the DB2 log to recover a system or object to any point in time
- Moving entire Storage Groups via FlashCopy to disk and potentially to Tape.

Invocation:

- Today, all volumes associated with a copy pool must be dumped with a single command processed on a single host, and Copy pools may be comprised of hundreds of volumes.
- Introduces an architecture to distribute workload across host members in an HSMplex; architecture allows for flexible configurations.
- Appears similar to the Common Recall Queue (CRQ) BUT 'Common Queue' is managed by sending all requests to a 'Master Scheduler' that maintains a single, common queue; unlike the Common Recall Queue that implements a persistent CF List Structure

DFSMSHsm



• DFSMSHsm Fast Replication Dump Enhancements

- For Dump Stacking today, the stacking value takes precedence over the number of available tasks.
- Volumes from different copy pools are not stacked onto the same tape.
- Updating the expiration date is manually intensive.
- **New enhancements:**
 - Provide a minimum stack option to indicate that fully utilizing the available dump tasks is more desirable than optimizing stacking.
 - Added a new **MINSTACK**, minimum stack, parameter and an **MAXSTACK** alias for the STACK parameter.
 - Allow multiple Copy pools to be stacked onto the same tape.
 - This will occur when the copy pools are dumped at the same time and use the same dump classes. Today this only occurs with Automatic Dump.
 - Provide a new command to update the expiration date for a copy pool dump copy.
 - New **UPDTCDS** command updates **ALL** of the CDS records that need to be updated to reflect the new expiration date. Avoid using hundreds of FIXCDS commands.

? Why It Matters: Optimize Dump throughput and improve usability.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



14

Invocation:

- 1) For example, with STACK(100) if there are 100 volumes and 5 available dump tasks, all 100 volumes will be dumped using a single task to the same tape.
 - MINISTACK represents the preferred minimum number of dump copies that the system should use to place on a tape volume.
 - Added MINISTACK(20) → First 100 volumes dumped, 20 volumes dumped onto 5 tapes, utilizing all 5 dump tasks
 - Depending on the value of MINISTACK, it may be possible that the **use of multi-tasking will use more tapes**
- 3) Today, it can take hundreds of FIXCDS commands to update the expiration date of a single copy pool dump version. This new command updates all of the CDS records that need to be updated to reflect the new expiration date. This command will be extended to other functions in the future.

DFSMSHsm



- **DFSMSHsm Fast Replication Messaging Enhancements**
 - Finding messages related to a Fast Replication command is very difficult because they are written to shared logs.
 - You have to find the correct log and then search through all of the messages for all HSM activity.
 - **New enhancement:** Provide an option to dynamically allocate a unique message data set and write all DFSMSHsm and DFSMSdss messages for a specific FRBACKUP or copy pool Automatic Dump to that data set.
 - Invoked via SETSYS command
 - **SETSYS FASTREPLICATION(MESSAGE DATASET(YES|NO HLQ(hlq)))**
 - Indicates that all Fast Replication messages associated with an FR command or copy pool automatic dump should go to a unique data set with a unique **HLQ**.
 - Each message is prefixed with the date and time that the message was written to the data set, along with the ID of the host the wrote the message.
 - Message Data Sets must be SMS managed.
 - **SHARE Requirement Addressed:** SSMVSS14005

? Why it Matters: Greatly simplify problem determination.



Invocation:

- Each message is prefixed with the date and time that the message was written to the data set, along with the ID of the host the wrote the message.
- Some messages will appear multiple times in a message data set. - each time a message is written to the operator or a log, the message is also logged in the message data set. **Duplicates are not removed.**

DFSMShsm



- **DFSMShsm/DFSMSdss Exploitation of zEDC**

- **New enhancements:** DSS and HSM now exploit zEDC for dumping and restoring data and when HSM uses DSS to move data.
 - DSS
 - Supports DUMP, RESTORE, COPY, DEFRAG, CONSOLIDATE, RELEASE, and PRINT of zEDC compressed format data sets
 - Supports new compression keyword for DUMP
 - HSM
 - Supports MIGRATION, RECALL, BACKUP, RECOVER, FULL-VOLUME DUMP, RECOVER FROMDUMP, FRBACKUP, FRRECOV, ABACKUP, ARECOVER of zEDC compressed format data sets
 - Uses the DFSMSdss zEDC support for certain functions.

? Why it Matters: Intended to provide efficient compression and lower CPU overhead than the processor-based and software-based compression methods already available.

- When migrating data to ML1 Disk using zEDC, use up to 58% less disk space and up to 80% less CPU compared to using DFSMSHsm with the COMPACT keyword.*
- When using DFSMSHsm with zEDC, recalling data from ML1 Disk uses up to 69% less CPU as compared to using DFSMSHsm with the COMPACT option.*

Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration, and software levels. *Measurements for comparisons were completed as part of a formal performance evaluation on a dedicated, isolated test system.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval





z/OS DFSMS Highlights

- **Catalog / IDCAMS**
 - **V2.2 Preview**
 - Support for > 255 GDGs
 - MODIFY CATALOG Security Enhancement
 - Catalog RAS Enhancements
 - IDCAMS VERIFY RECOVER
 - IDCAMS RAS Enhancements

Session 16920: Introduction to ICF Catalog Tuning Techniques,
Wednesday 11:15AM

Session 17104: What's New with DFSMS ICF Catalog and IDCAMS,
Wednesday 4:30PM

Session 17103: What's New With the Catalog Search Interface,
Friday 8:30AM

Catalog



- **GDG Extended (GDGE)**

- Today, only up to 255 GDSs are controlled by a GDG.
- **New enhancement:** New DEFINE **EXTENDED** keyword allows up to 999 GDS to be active within 1 GDG.
 - DEFINE SMS.GDG GDG EXTENDED LIMIT(999)
 - Default is NOEXTENDED LIMIT is 255.
 - Enabled via SYS1.PARMLIB member IGGCATxx
 - New GDGEXTENDED variable set to "YES" (default is "NO").
 - **Coexistence considerations**
 - **ALL** systems in an installation should be at z/OS V2.2 or higher before enabling GDGEs, and it's expected that there will be no attempt to fall back a prior release.
 - Coexistence APAR OA46009 prevents access to GDGEs on lower level systems.
- **SHARE Requirement Addressed:** GO6SMG91037

? Why it Matters: Support a larger number of GDSs to make it easier to manage data sets that are created frequently. For example, it will be possible to specify a full year of GDSs be kept.



Catalog



- **Additional GDG Enhancements**

- When users take the default NOSCRATCH option on DEFINE, it can leave unwanted data sets when GDSes roll off.
- In addition, users sometimes inadvertently put retention periods on their GDSs. System administrators would like for these data sets to be deleted regardless of retention periods.
- **New enhancements:**
 - New **GDGSCRATCH (YES|NO)** parmlib variable overrides IDCAMS keyword on GDG DEFINE.
 - If not set or parmlib variable is GDGSCRATCH(NO), the keyword specified on the DEFINE GDG takes precedence.
 - New **PURGE** keyword added to IDCAMS DEFINE and ALTER GDG which allows user to specify what happens to unexpired GDSs during roll-off.
 - Specifies that when the SCRATCH is set, unexpired generation data sets (GDSs) are to be deleted during roll off processing. Default is NOPURGE.
 - New variable, **GDGPURGE (YES|NO)**, for SYS1.PARMLIB member IGGCATxx. Default is "NO". Overrides DEFINE GDG setting when PURGE is set.

? Why it Matters: Allows system administrators more control over and flexibility with GDG processing.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



19

GDGSCRATCH:

- **YES** - Causes Catalog to set attribute to SCRATCH in GDG catalog record regardless of the setting in the DEFINE GDG statement. SCRATCH specifies a generation data set (GDS) is to be deleted from disk when uncataloged from the GDG for EMPTY|NOEMPTY processing
- **NO** - the keyword specified on the DEFINE GDG takes precedence.
- Parmlib processing runs at CAS init or Catalog RESTART - does not change original setting once attribute is set in DEFINE GDG
- If not set or parmlib variable is GDGSCRATCH(NO), the keyword specified on the DEFINE GDG takes precedence.
- Applies to GDGs and GDGEs.

GDGPURGE

- Specified as GDGPURGE(YES) - overrides DEFINE GDG setting when the PURGE attribute is set.
- Only used when SCRATCH is also set
- Default is GDGPURGE(NO)

Catalog



- **MODIFY CATALOG Security Enhancement**

- Existing requirements to provide access to the Catalog Address Space reporting commands without also allowing ability to issue commands which might alter the catalog configuration.
 - A user via SDSF might be able to issue an "F CATALOG,ALLOCATED" command, but not be able to issue an "F CATALOG,CLOSE(usercat)" command.
- **New enhancement:** New RACF resource profile checked by catalog command processing to insure that the issuer of a command is authorized to execute that command.
 - Allows users utilizing SDSF to issue operator reporting commands, the results of which they can view at their terminal, but not permit them to issue commands that change the catalog configuration.
 - New optional OPERCMDS profile, MVS.MODIFY.STC.CATALOG.CATALOG.SECURE

? **Why it Matters:** Provide more granular security and better operational flexibility.



The only requirement for activating this facility is to define a new OPERCMDS Resource profile:

- MVS.MODIFY.STC.CATALOG.CATALOG.SECURE
- and then PERMIT users and operators (as needed) to this resource via the RACF PERMIT command, where either READ or UPDATE authorization is granted to a user or user group.

Catalog



- **RAS Enhancements**

- **Catalog Attributes Healthcheck**

- A shared catalog must be defined with SHAREOPTIONS(3 4) and must reside on a shared volume.
 - Catalogs that reside on shared volumes will become damaged if referred to by another system and the shareoptions are inconsistently set.
 - **New enhancement:** A new healthcheck will be designed to identify catalogs residing on volumes whose devices are defined as shared in the active IODF that might have their SHAREOPTIONS set incorrectly.

? Why it Matters: Damage to incorrectly shared catalogs can now be avoided. Further, for catalogs on non-shared devices, the shareoptions can be set to (3,3), which will provide for more efficient access and better overall performance.



Catalog



- **RAS Enhancements**

- **DFSMSDss Restore of a User Catalog to Any Volume**

- Today, DFSMSDss only allows a user catalog to be restored to the same volume from which it was dumped (ie source dump volume must match restore target volume).
 - **New enhancement:** Allows a user catalog to be restored to any volume.
 - Device geometry must match and applies only to logical restore.
 - Restriction for physical data set restore remains.

? Why it Matters: Improves the users ability to recover a user catalog.



When a user catalog is backed up there could be a point where the user catalog is moved to a different volume (volser).

IDCAMS



• VERIFY RECOVER Enhancement

- Currently, customers run IDCAMS EXAMINE on their datasets to look for errors. Some of the errors can be corrected by running the existing IDCAMS VERIFY or VERIFY RECOVER function.
 - However, both forms of IDCAMS VERIFY do not correct most of the errors returned from EXAMINE, and customers would need to either reorganize their datasets or obtain backup copies.
- **New enhancement:** EXAMINE and IDCAMS VERIFY RECOVER are enhanced to correct the following errors returned by EXAMINE:
 - IDC11718I DATA COMPONENT HIGH-USED RBA IS NOT EQUAL TO CA SIZE
 - IDC11728I DATA FOUND IN EMPTY CI
 - IDC11724I DATA COMPONENT CA NOT KNOWN TO SEQUENCE SET.
- Now, there are 3 ways to run IDCAMS VERIFY:
 - IDCAMS VERIFY (without RECOVER)
 - IDCAMS VERIFY RECOVER
 - **New enhancement:** EXAMINE followed by IDCAMS VERIFY RECOVER.
 - Functional only when EXAMINE and IDCAMS VERIFY RECOVER are run in the same job step.

? **Why it Matters:** Improved error recovery.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

23



IDCAMS VERIFY (without RECOVER):

- IDCAMS opens the data set for output, issues the VERIFY macro with no options, and closes the data set. It will only fix the index level number, high used RBA, and high level index RBA in the Catalog VVR.

IDCAMS VERIFY RECOVER:

- Currently, the RECOVER option causes VSAM Record Management VERIFY to back out or complete any interrupted CA reclaim in addition to regular IDCAMS VERIFY functions. This line item does not change this way of running VERIFY.

EXAMINE followed by IDCAMS VERIFY RECOVER:

- This is a new enhancement functional only when EXAMINE and IDCAMS VERIFY RECOVER are run in the same job step.
- With the line item, **EXAMINE always saves the info (such as index CI#, data CI#, error type, etc) of the data set errors it found in the error message area pointed to by RPLERMSA.** Enhancements to the VSAM VERIFY macro is made to process the area for errors to repair.

IDCAMS



- **RAS Enhancements**

- **REPRO MERGECAT output is too long**

- REPRO MERGECAT currently produces up to 6 lines for each entry that is processed; one blank line, one or two for the IDC0639I message (Sphere conversion started), a blank line, and one or two for the IDC01402I (Sphere conversion ended). The one or two lines depend on the length of the dataset name.
 - **New enhancement:** Add MESSAGELEVEL keyword to REPRO.
 - New keyword **MESSAGELEVEL(ALL|SHORT)**
 - **ALL** is the default and is the same as the current output
 - **SHORT** reduces the number of lines for each entry

? Why it Matters: Satisfy customers requirements to make REPRO more usable and flexible.



IDCAMS



- **RAS Enhancements**

- **PRINT and REPRO functions cannot process damaged VSAM CIs or records.**
 - **New enhancement:** Add CIMODE keyword to PRINT and REPRO.
 - New keyword **CIMODE**, which PRINTs or REPROs by CI rather than by record.
 - Opens VSAM data set in CIMODE and is primarily for use with ESDS, but could be used to access a KSDS data component directly.

? Why it Matters: Improved error diagnosis / recovery by allowing processing if CI or VSAM record is damaged.



z/OS DFSMS Highlights



- **SMS**
 - **V2.2 Preview**
 - Space Constraint Relief
 - RAS Enhancements
 - Storage Group Space Alert Messages

Session 17045: NaviQuest: Testing ACS and Streamlining SMS Work, Thursday 11:15AM

Session 17048: NaviQuest: Testing ACS and Streamlining SMS Work (Lab), Thursday 12:30PM



SMS

- **Space Constraint Relief**

- Currently for **non-Guaranteed Space** allocation, SMS Space Constraint Relief (SCR) processing simply reduces the requested space by the specified percentage, 'Reduce Space Up To (%)', in the data class if the originally requested space can not be satisfied.
- **New enhancement:** SMS will extend the existing space reduction function to **Guaranteed Space** allocations to reduce allocation failures.
 - A new sub-parameter, **Guaranteed Space Reduction (Y | N)**, is added to the data class panel for the user to specify whether space reduction on guaranteed space allocation is permitted or not.
 - Instead of allocating the smallest space quantity, SCR space reduction function is enhanced to allocate the largest possible space that satisfies the specified reduction percentage for both guaranteed space and non-guaranteed space requests.
 - **Note:** Space reduction will remain unsupported for striping allocation.

? Why it Matters: Helps reduce allocation failures when the requested space is not available.

SMS



- **RAS Enhancements**

- **Read-only Variable**

- Currently the user cannot provide installation specific values to the ACS routines via existing ACS read-only variables.
 - **New enhancement:** SMS provides a new user defined ACS Read-only variable in SMS IGDSMSxx parmlib member:
 - **USER_ACSVAR(value1,value2,value3)**
 - The new parameter has 3 positional values.
 - Each value is a user-defined 8 character string.
 - The positional values of the new parameter can be altered by the following SETSMS command:

SETSMS USER_ACSVAR(value1,value2, value3)

? Why it Matters: Each installation can specify unique values, similar to how the system symbols are defined, to assign proper constructs without a need to modify the ACS routines constantly, making it easier to reuse ACS routines on different systems.

SMS



- **Storage Group Space Alert Messages**
 - Currently, SMS VTOC Data Set Services (VDSS) issues IGD17380I to notify a user that the cumulative space allocated on the selected storage group has exceeded the high allocation threshold.
 - *IGD17380I STORAGE GROUP(sgname) IS ESTIMATED AT xx% OF CAPACITY, WHICH EXCEEDS ITS HIGH ALLOCATION THRESHOLD OF zz%*
 - **New enhancement:** SMS will be designed to allow you to specify new storage group space warning thresholds separately from the high allocation thresholds.
 - Create new alert threshold attributes in the storage group panel:
TOTAL SPACE ALERT THRESHOLD %..... (0-99)
TRACK-MANAGED SPACE ALERT THRESHOLD %.....(0-99)
 - If the alert threshold value is not specified, it defaults to zero and no alert messages will be issued.
 - New alert messages, IGD400I and IGD401I, will be issued to the console when the alert thresholds have been reached.



Alert message may not be ideal:

- The user may need to change it to a higher or lower value to capture more meaningful alerts
- It may inadvertently affect SMS volume selection behavior.

The equivalent fields on **NaviQuest**:

TOTSPALERT()

TRKSPALERT()

The **DCOLLECT** record type 'SG' has 2 new one-byte fields:

DSGTOTAP at offset 916(X'394')

DSGTMSAP at offset 917(X'395')

SMS



- **Storage Group Space Alert Messages**
 - IGD400I TOTAL SPACE ALERT ON STORAGE GROUP (sgname)
CURRENT USAGE (xx%), ALERT THRESHOLD (yy%)
 - IGD401I TRACK-MANAGED SPACE ALERT ON STORAGE GROUP
(sgname)CURRENT USAGE (xx%), ALERT THRESHOLD (yy%)
 - The alert messages will be issued at an incremental interval to prevent them from being issued too frequently.
 - IGD401I is issued only for the pool storage group containing one or more EAV volumes.
- ? Why it Matters:** Allows users to set a lower threshold for warning messages, which can provide more time to react to storage group space shortage conditions.



Alert message may not be ideal:

- The user may need to change it to a higher or lower value to capture more meaningful alerts
- It may inadvertently affect SMS volume selection behavior.

The equivalent fields on **NaviQuest**:

TOTSPALERT()

TRKSPALERT()

The **DCOLLECT** record type 'SG' has 2 new one-byte fields:

DSGTOTAP at offset 916(X'394')

DSGTMSAP at offset 917(X'395')

SMS



- **Storage Group Space Alert Messages**
 - **New enhancement:** A new optional keyword, **ALERT**, is introduced for the DISPLAY SMS command.
 - **D SMS[,{STORGRP|SG}{(sgname|ALERT|ALL)}[,LISTVOL]**
 - The new optional keyword, **ALERT**, is used to display any pool storage groups which have already reached either the total space alert threshold or the track-managed alert threshold.
 - The existing message, IGD002I, will be enhanced to include the space information: total space, total track-managed space, usage percentages and alert thresholds on the system at which the command is issued if the information is available.
 - The space information is shown for pool storage groups only.
 - **? Why it Matters:** Makes it easier to see when it might be necessary to change a storage group's space management settings or add volumes to a storage group.



SMS



▪ Sample Output for D SMS,SG(ALERT)

IGD002I 11:50:04 DISPLAY SMS 021

```
STORGRP  TYPE      SYSTEM= 1 2 3 4 5 6 7 8
SXP02    POOL      + + + + + + + + +
SPACE INFORMATION:
TOTAL SPACE = 306458MB USAGE% = 99 ALERT% = 65
TRACK-MANAGED SPACE = 106220MB USAGE% = 100 ALERT% = 50
***** LEGEND *****
. THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
+ THE STORAGE GROUP OR VOLUME IS ENABLED
- THE STORAGE GROUP OR VOLUME IS DISABLED
* THE STORAGE GROUP OR VOLUME IS QUIESCED
D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
> THE VOLSER IN UCB IS DIFFERENT FROM THE VOLSER IN CONFIGURATION
SYSTEM 1 = SYSTEM1  SYSTEM 2 = SYSTEM2  SYSTEM 3 = SYSTEM3
SYSTEM 4 = SYSTEM4  SYSTEM 5 = SYSTEM5  SYSTEM 6 = SYSTEM6
SYSTEM 7 = SYSTEM7  SYSTEM 8 = SYSTEM8
```

- Note: Only one storage group SXP02 has its usages which have reached the alert thresholds.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



32

z/OS DFSMS Highlights



- **Access Methods**
 - **V2.2 Preview**
 - RLS Index Level Locking
 - Secondary Space Reduction

Session 17108: DFSMS Exploitation of z/OS zEnterprise Data Compression (Revised from Pittsburgh), **Monday 3:15PM**

Session 16731: Getting the Most out of your VSAM Data Sets in CICS by Using RLS, **Thursday 8:30AM**

Session 17137: VSAM New Functions in z/OS 2.1 and 2.2, **Friday 11:15AM**

Access Methods



• RLS Index Level Locking

- Today when processing VSAM data sets, there is a single point of contention (DIWA) preventing concurrent VSAM RLS data CI splits in different CAs, CI reclaims, and spanned-record insertions and updates.
- **New enhancement:** Introduce a new level of serialization at the sequence set index CI level.
 - This function is enabled by default on all systems on V1R13 and above.
 - New SMF fields were added to MVS System Management Facilities (SMF) for the new index record locks.
 - **Toleration Considerations:**
 - Toleration APAR OA42676 is required to be installed on all V1R13 and V2R1 systems across the sysplex prior to bringing a V2R2 system into the sysplex.
 - Additional feature will chain the I/O for the multiple segments of a spanned record so that only one I/O is needed to write out the spanned record. This will be available for NSR buffering only (not LSR or GSR).

? Why it Matters:

- Reduces contention and improves performance for many frequently updated VSAM key-sequenced and variable-length relative record (KSDS and VRRDS) data sets.
 - Side benefit: Reduces the window for inconsistent spanned-record segments.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

34



Toleration Required:

- This function is enabled by default on all systems on V1R13 and above. There is no disablement for this line item.
- Toleration is a **MUST** – without it broken data sets can occur.

New SMF42 records: new SMF 42 fields for component_1 class 4

Access Methods



• Secondary Space Reduction

- Currently to allocate a secondary extent on the current volume, VSAM, RLS, PDSE, BAM and SAM will try to satisfy the user's specified amount on the DEFINE by calling DADSM extend services.
 - When there is not enough free space available on the current volume, EOVS will call SMS to switch volumes and try to extend onto a new volume.
 - This algorithm can be inefficient in leaving large amounts of free space behind on the current volume which is less than what the user specified for secondary amount.
- **New enhancement:** Allow the system to extend data sets by less than their originally specified secondary space allocation amounts when doing so would avoid allocation of space on additional volumes.
 - DADSM will provide a new minimum allocation amount on the DADSM Extend interface. This new parameter along with the current secondary allocation amount on the IDCAM DEFINE, will allow DADSM the ability to determine the largest free space to allocate.
 - The allocated space must be equal to or greater than the minimum quantity derived from the attribute **Reduce Space Up To (%)** on the ISMF data class definition under "Space Constraint Relief", and be less than or equal to the original specified secondary allocation amount.
 - This support will be for SMS managed non-striped VSAM data sets and Non-VSAM data sets.

? Why it Matters: Help to improve disk space utilization.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



35

For example:

- **OLD:** secondary allocation amount of 100 cyl. specified, and the largest free spaces is 90 cyl. We would leave behind 90 cyl. and go to a new volume
- **NEW:** If 40 was specified on the "Reduce Space Up To (%)", for a secondary allocation of 100 cylinders, you may now see 80, 90 or even 60 cylinder extents, but you won't see an extent for less than 60 cylinders. With this support, the allocated extent sizes on the same volumes can vary depending on the availability of free storage on the volume.
- A new bit field (**SMF64SSR**) in the SMF record for VSAM/RLS will be added to indicate that the current extent was reduced during EOVS processing.
- A new field (**SMF64NTA**) in the SMF record for VSAM/RLS will be added to indicate the number of tracks acquired during this EOVS processing.

z/OS DFSMS Highlights

- **Open/Close/EOV**
 - **V2.2 Preview**
 - Dynamic Exits
 - RAS Enhancements





OCE

- **Dynamic Exits**

- Current design of DFSMS Open/Close/EOV Installation Tape exits requires IPL for changes to be put into effect.
 - In addition, there is no way to add multiple exit routines for the exit.
- **New enhancement:** Several DFSMS Open/Close/EOV Installation Tape exits will be converted to be dynamic.
 - Impacted tape exits include: **Label Anomaly, Volume Mount, File Validation, File Start, and File End.**
 - All exits are AMODE 31, reentrant and persistent until the next IPL.
 - New rules for parameter list and execution, as well as return code decisions when multiple exits are called. See z/OS DFSMS Installation Exits, SC23-6850, for specific details and considerations.

? **Why it Matters:** No need to IPL to manage changes with Tape Installation exits, and multiple exit routines may be associated with each exit.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



OCE



- **OCE RAS Enhancements**

- **Abend code Elimination**

- **New enhancement:** Eliminate abend code A13, reason code 18 by automatically selecting the correct volume during Open processing for multivolume tape data sets.

- **? Why it Matters:** Provide improved usability.

- **Add JOBID & SYSPLEXID to SMF 14/15**

- **New enhancement:** JOBID and SYSPLEXID are now added to SMF 14/15 (non-VSAM data set activity) records.
 - **SHARE Requirement Addressed:** SSMVSE11033 (SHARE Top 75)

- **? Why it Matters:** Make chargeback and capacity planning easier.



OCE



- **OCE RAS Enhancements**

- **Improved DEVSUPxx Processing**

- There is a subset of DEVSUPxx PARMLIB member keywords that are initially reset to their default values before processing the member. However, other keywords and their values are left intact.
 - This creates an inconsistency in how all keywords are processed.
 - **New enhancement:** All keywords will be treated in a consistent manner so that none will be initially reset to their default values both during IPL and with processing SET DEVSUP command.
 - The **SET DEVSUP=xx** target can specify only the keyword(s) that need to be changed.
 - Continued processing of the DEVSUP member stops at the point where a keyword error is detected. There is no indication which DEVSUP member is processed.
 - **New enhancement:** If syntax error is detected, the system continues processing and IEA252E message is issued:
 - IEA252E DEVSUPxx: INVALID SYNTAX ON LINE yyyy. MEMBER PROCESSING CONTINUES

? Why it Matters: Provide improved usability.



OCE



- **OCE RAS Enhancements**

- **Improved DEVSUPxx Processing**

- SET DEVSUP MVS command can only support single value
 - In addition, IEASYSxx PARMLIB member supports only single value for the DEVSUP keyword.
 - **New enhancement:** Allow you to specify more than one DEVSUPxx member in a single SET command.
 - SET DEVSUP=(xx,yy,zz...)
 - IEASYSxx PARMLIB member now supports multiple values for DEVSUP keyword: DEVSUP=(xx,yy,zz...)
 - **New enhancement:** Display additional information about tape-related DEVSUPxx parameters with new support in the DEVSERV command.
 - A new **DEVSERV QLIB** operand will display settings for TAPEAUTHDSN, TAPEAUTHF1, TAPEAUTHRC4, and TAPEAUTHRC8.

? Why it Matters: Provide improved usability.



z/OS DFSMS Highlights

- **DADSM / CVAF**
 - **V2.2 Preview**
 - RAS Enhancements



DADSM / CVAF



- **DADSM RAS Enhancements**

- As documented the DADSM pre-processing exit, IGGPRE00, does not support space changes of SMS managed datasets.
- **New enhancement:** Allow the IGGPRE00 exit to be able to alter SMS managed allocations (ie primary-space and secondary-space quantities).
 - IGGPRE00 parm list will include a pointer to the IGGDacre control block.

? Why it Matters: Improved usability and flexibility with the DADSM user exit.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



42

z/OS DFSMS Highlights



- **DFSORT**
 - **V2.2 Preview**
 - zHPF Exploitation
 - Functional Enhancements

**Session 17106: The ICETOOL Cometh,
Monday 4:30PM**

**Session 17105: What's New with
DFSORT? (V2R1 and V2R2 Features),
Wednesday 8:30AM**

DFSORT



- **zHPF Exploitation**

- Currently DFSORT does not exploit zHPF.
 - DFSORT normally uses EXCP for processing of basic and large format sequential input and output data sets (SORTIN, SORTOUT, OUTFIL).
 - DFSORT already uses BSAM for extended format sequential input and output data sets (SORTIN, SORTOUT and OUTFIL). BSAM already supports zHPF.
- **New enhancement:** Update DFSORT to prefer BSAM for SORTIN/SORTOUT/OUTFIL when zHPF is available.
 - DFSORT will automatically take advantage of zHPF if it is available on your system; no user actions are necessary.

? Why it Matters: Taking advantage of the higher start rates and bandwidth available with zHPF is expected to provide significant performance benefits on systems where zHPF is available.



DFSORT



- **Functional Enhancements**

- **Date conversion AGE function**

- Customers often want to calculate the date duration that specifies the number of years, months, and days between an input date and current date.
 - **New enhancement:** A date conversion function **AGE** for the BUILD and OVERLAY operands of DFSORT's INREC, OUTREC and OUTFIL statements can now be used to calculate the date duration (ie time between a given date and the current date):
 - **AGE=YMD** produces a 8 byte result which has duration in years (0-9999), months (00-12), and days (00-31).
 - **AGE=YM** produces a 6 byte result which has duration in years (0-9999), months (00-12).
 - **AGE=YD** produces a 7 byte result which has duration in years (0-9999), days (00-366).

- **Date conversion WEEKNUM function**

- Customers often want to calculate the week of the year for an input date.
 - **New enhancement:** A date conversion function WEEKNUM for the BUILD and OVERLAY operands of DFSORT's INREC, OUTREC and OUTFIL statements can now be used to calculate the week number that represents the week of the year.
 - **WEEKNUM=USA** returns an integer in the range of 1 to 54 that represents the week of the year. The week starts with Sunday, and January 1 is always in the first week.
 - **WEEKNUM=ISO** function returns an integer in the range of 1 to 53 that represents the week of the year. The week starts with Monday and includes 7 days.

? Why it Matters: Provide additional flexibility in creating reports and to help improve the usability of reports generated with these new functions.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

45



For both functions, previously users would have had to code a program to achieve the same results.

z/OS DFSMS Highlights



- **DFSMS / Storage Synergy**
 - **V2.2 Preview**
 - Multi-Target PPRC
 - IBM zHyperWrite™
 - Easy Tier Application Hints
 - XRC Workload Write Pacing
 - Multiple Incremental FlashCopy

Session 16723: Effective Backups: Selecting the Right Backup Mechanism to Match the Business Requirement, **Wednesday 1:45PM**

Session 16796: DB2 for z/OS and Advanced Copy Services, **Wednesday 3:15PM**

Session 17112: Multi-Target Replication with IBM DS8870, **Thursday 8:30AM**

Session 17113: IBM DS8000® Replication Performance Considerations, **Thursday 10AM**

Session 17017: Significant Advancements in System z Resiliency Provided by GDPS, **Thursday 10AM**

Session 17022: Extending z/OS Mainframe Workload Availability with GDPS/Active-Active, **Thursday 3:15PM**

Session 17092: IBM DS8870 Enterprise Storage News and Review, **Thursday 4:30PM**

Complete your session evaluations online at www.SHARE.org/Seattle-Eval




Storage Synergy



- **Multi-Target PPRC (MT-PPRC)**

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

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



Storage Synergy



- **IBM zHyperWrite™ - DS8870 Log Write Accelerator for DB2**
 - New technology that combines DS8000 and z/OS enhancements that deliver performance benefits for writing operations to DB2 logs in the Metro Mirror (PPRC) environment.
 - Reduces latency overhead compared to normal storage based synchronous mirroring.
 - Can help reduce up to 43% of the DB2 log write time.*

? **Why it Matters:** Provide improved DB2 transactional response time and log throughput improvements.

- **Easy Tier Application Hints**
 - System Storage Easy Tier is now designed to allow software-defined policy information to be communicated to Easy Tier control units, to help them efficiently deploy storage.
 - Easy Tier currently optimizes data placement and tiering based on workload history and this does not always reflect the future performance requirements of the data.
 - New planned exploitation by DB2 to provide hints to Easy Tier on the DS8000 when a database reorganize is done so that the target data set for the database will be reorganized and optimized based on the activity of the application.

? **Why it Matters:** Intended to help steer data placement within Easy Tier volumes to meet application performance objectives.

*Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels. †Measurements for comparisons were completed as part of a formal performance evaluation on a dedicated, isolated test system.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Storage Synergy



- **XRC Workload Write Pacing**

- XRC is designed to work with z/OS Workload Manager to control XRC Write Pacing.
 - Reduces administrative overhead on hand managing XRC write pacing.
 - Avoids the need to define XRC write pacing on a per volume level allowing greater flexibility in configurations.
 - Prevents low priority work from interfering with the Recovery Point Objective of critical applications.

? **Why it Matters:** Simplifies XRC usage by eliminating the need to constantly adjust the write pacing settings and/or allocation of data to volumes..

- **Multiple Incremental FlashCopy**

- New FlashCopy function that supports up to 12 targets for incremental FlashCopy.
 - Previously only a single incremental FlashCopy was allowed for any individual volume.
 - Incremental FlashCopy allows a full physical copy to be created much more quickly than a non-incremental FlashCopy.

? **Why it Matters:** Provide more flexibility and resilience for FlashCopy helps to better protect application availability and provide improved data protection across physical volume failure events.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

49



•An Incremental relationship is a persistent ~~FlashCopy with change recording active, so that when another copy is needed, the change recording that has been occurring will allow only the changed tracks since the last increment was taken to be copied to the target.~~

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>
 - R13
 - DFSMS Using the New Functions (SC26-7473-08) – (<http://publibz.boulder.ibm.com/cgi-bin/bookmgr/Sshelves/ez2zo111?filter=DFSMS+Using+the+New+Functions+&SUBMIT=Search+titles>)
 - z/OS V1.13 DFSMS Technical Update - (<http://www.redbooks.ibm.com/abstracts/sq247961.html?Open>)
 - R12
 - DFSMS Using the New Functions (SC26-7473-07) - <http://publibz.boulder.ibm.com/epubs/pdf/dqt2q570.pdf>
 - z/OS V1.12 DFSMS Technical Update - <http://www.redbooks.ibm.com/abstracts/sq247895.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.s.zdatamgmt/zsysprog_c_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>



5

2

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*	FICOM*	Lotus*	System Storage
DB2*	FlashCopy*	MQSeries*	Tivoli*
DFSMS	GDDM*	Multiprise*	TotaraStorage*
DFSMSdftp	GDS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMSshsm	HyperSockets	Parallel Sysplex*	VMEESA*
DFSMSrmm	HyperSwap	PR/SM	VSEESA
DFSORT	IBM*	QMF	VTAM*
DFSMS	IBM xggo*	RACF*	WebSphere*
DS4000	ImagePlus*	Rational*	zArchitecture*
DS6000	IMS	RMF	zOS*
DS8000	Intelligent Miner	System i	zVM*
Enterprise Storage Server*	Language Environment*	System z	zVSE
ESCON*		System z9	zSeries*
			zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies:

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

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

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

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

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

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

NOTES:

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

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

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

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



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.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

