



z/OS DFSMS

Charlie Burger (cwburger@us.ibm.com)
San Jose Advanced Technical Support Center
408-256-0331 T/L 8/276-0331



Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.

- | | | |
|----------------|------------|------------|
| ▪ DFSMS | ▪ z/Series | ▪ CICS |
| ▪ DFSMS/MVS | ▪ z/OS | ▪ CICS/ESA |
| ▪ DFSMSdfp | ▪ OS/390 | ▪ CICSVR |
| ▪ DFSMSdss | ▪ S/390 | ▪ CICS/TS |
| ▪ DFSMShsm | ▪ ES/3090 | ▪ IMS |
| ▪ DFSMSrmm | ▪ FICON | ▪ RACF |
| ▪ DFSORT | ▪ ESCON | ▪ OS/400 |
| ▪ TotalStorage | ▪ MVS | |
| ▪ IBM | ▪ MVS/ESA | |

Other company, product, and service names mentioned may be trademarks or registered trademarks of their respective companies.



z/OS V1R5 DFSMS



Table of Contents

- Session Objectives
- z/OS V1R5 DFSMS (DFSMSdfp)
 - ƒ Enhanced data integrity for physical sequential data sets
 - ƒ File Sequence number > 9999
 - ƒ Multilevel security
 - ƒ Tape: Remove 8 Media Type Limit
 - ƒ BPAM Unix files support
- z/OS V1R5 DFSMS (DFSMSdss)
 - ƒ Exploitation of 64-bit Real Storage
 - ƒ Fast Replication Support
 - ƒ Physical DUMP and RESTORE of VSAM extended format data sets
 - ƒ UIM enhancements
- z/OS V1R5 DFSMS (DFSMShsm)
 - ƒ ABARS Catalog Search Interface (CSI)
 - ƒ CDS/EA Under CDSQ
 - ƒ Fast Replication Backup
 - ƒ RACF Facility Class

Table of Contents *(continued)*

- z/OS V1R5 DFSMS (DFSMSrmm)
 - ƒ Backup at Any Time
 - ƒ Duplicate VOLSER Support
 - ƒ OAM: Improved Volume Management
 - ƒ Multilevel security
- Miscellaneous z/OS V1R5 DFSMS DFSMSdftp Projects
 - ƒ Catalog Customer Satisfaction
 - ƒ SMS Availability/Usability Enhancements
 - ƒ VSAM Customer Sat
 - ƒ SAM 59 Stripes
 - ƒ Miscellaneous Items
- Summary
- References

Session Objectives

- To familiarize the audience with the enhancements introduced for DFSMS z/OS V1R5
- To provide the audience with basic implementation considerations and provide a basis for migration planning
- Only new functions introduced with z/OS V1R5 will be discussed
 - ƒ New functions introduced via PTF on earlier releases and brought forward into z/OS V1R5 will not be covered



IBM TotalStorage™

DFSMSdfp

© 2003 IBM Corporation

IBM TotalStorage™



z/OS V1R5 DFSMSdfp

- Enhanced Data Integrity for Physical Sequential Data Sets
- File Sequence Number > 9999
- Multilevel security
- Tape: Remove 8 Media Type Limit
- BPAM Unix files support

8

© 2003 IBM Corporation

Enhanced Data Integrity for Physical Sequential Data Sets

- Currently possible to have a physical sequential data set concurrently opened by multiple users for either output or update with DISP=SHR
 - f* Could cause destruction or loss of data
- Enhanced data integrity (EDI) if activated will allow user to prevent concurrent output access by multiple users
 - f* Can either fail the OPEN with an ABEND or issue a warning message

EDI Activation

- EDI will be shipped inactive
- New IFGPSEDI parmlib member in SYS1.PARMLIB
 - f* If member doesn't exist EDI will be inactive
 - f* MODE(WARN|ENFORCE|DISABLE)
 - WARN will result in warning message
 - IEC984I indicating whether the data set open was for input or output when already opened for output
 - IEC985I when data set is bypassed (in EDI table) and is already open for either input or output
 - ENFORCE will result in ABEND
 - DISABLE will deactivate EDI
- EDI activated automatically at IPL if IFGPSEDI specifies activation
- After IPL, IFGPSEDI can be modified and then start task IFGEDI
- To deactivate, remove IFGPSEDI and IPL or start task IFGEDI with MODE(DISABLE)

Enforcing EDI

- EDI enforced by placing an ENQ on major resource SYSZEDI minor volser+mode+dsn at OPEN and DEQ at close or forced closed during task termination
 - ƒ EDI not enforced for data sets shared across multiple GRSplexes
 - ƒ Only mode currently allowed is 'O' for open for output
- If EDI is active and if applications can maintain data integrity on their own:
 - ƒ Their data sets can be specified in IFGPSEDI or
 - ƒ Request OPEN not fail the request by either
 - Setting new DCBE flag DCBEEFPS
 - Program APF authorized or
 - Executing in system key(0-7) or
 - Executing in supervisor state
 - Set SCTDNDSI in the SCT
 - NODSI in Program Properties Table (PPT) for that application
 - Set DSABNODI in DSAB with dynamic allocation
 - S99NORES

11

© 2003 IBM Corporation

EDI Considerations

- Warning mode
 - ƒ No guarantee that all possible EDI violations will be found and messages issued
 - If an application opens a data set for output causing an ENQ on SYSZEDI and another application attempts to open the same data set for output, an EDI violation message will be issued but the second open will be allowed to proceed
 - If the first application closes the data set the DEQ on SYSZEDI will be released but the second application still has the data set opened for output
 - If opened again for output, no EDI violation message will be issued.
- Enforce mode
 - ƒ Possible to miss violations if an application requests EDI be bypassed but other applications don't
 - If selective EDI bypassing is allowed, the ENQ on SYSZEDI will not be done and another application could successfully open that data set for output
- If there is more than one data set of the same name, bypassing EDI will be done for all data sets of that name

12

© 2003 IBM Corporation

IFGPSEDI Syntax

- MODE(WARN | ENFORCE | DISABLE)
- DSN(=datasetname),DSN(=datasetname)....

MODE(ENFORCE),DSN(A.B),DSN(C.*.B),DSN(D.**),DSN(E%.**)

SMF Recording

- TYPE 14/15 records will be updated with new flags
 - f* SMF14/15EXT - Data set being processed is in the exclude list
 - f* SMF15OUT - Data set being opened for output is already opened for output and EDI warning mode is active
 - f* SMF14INO - Data set being opened for input is already opened for output
 - f* SMF14/15EPS - DCBE/SCT/DSAB indicated this data set was not to fail as a result of EDI processing

Migration/Coexistence

- No coexistence PTFs
- EDI will not be enforced on lower level systems
- All systems wanting EDI active must code IFGPSEDI

15

© 2003 IBM Corporation

File Sequence Number > 9999

- Prior to z/OS 1.5 the file sequence number maximum was 9999
- For IBM Standard Label tapes Maximum will now be 65535 :
 - f* SL
 - f* SUL
 - f* LTM
- For unlabeled tapes (NL) and using Bypass Label Processing (BLP) will now be 65535
 - f* The limit of 65535 files actually means that only 21845 logical data sets can be accessed since with BLP the labels are considered as files
- JCL LABEL parameter still limited to 9999
 - f* Use RDJFCB macro to obtain JFCB and update the file sequence number in JFCB and use OPEN, TYPE=J if data set isn't cataloged or...
 - f* If data set is cataloged and LABEL=(,label type) is coded not specifying file sequence number use OPEN and catalog will provide file sequence number
- Allow stacking of files on tape volumes to fully utilize large capacity tape cartridges

16

© 2003 IBM Corporation

Label Format

- The format of the data set sequence number field in the HDR1/EOV1/EOF1 label for SL tapes will be changed as follows:
 - f* If data set sequence number is between 1 and 9999 the format will stay the same-4 byte numeric value in EBCDIC format
 - f* If data set sequence number is between 10000 and 65535 then the first byte of the field will be the EBCDIC character '?'
 - The next 3 bytes will be the data set sequence number in binary format
 - f* OS/400 currently supports this label format

17

© 2003 IBM Corporation

DFSMSHsm ABARS

- The ABARS include data set can include cataloged tape data sets that reside on tape. ABARS has been modified so that it can back up and recover a cataloged tape data set with a file sequence number greater than 9999.
- The ABARS ACCOMPANY data set includes cataloged tape data sets. These data sets are not backed up, but the customer is instead telling ABARS that they will supply the tape that has the data set on it. ABARS needs only to catalog the data set at the recovery site. ABARS has been modified to ensure that an ACCOMPANY tape data set whose file sequence number is greater than 9999 can be successfully recataloged and opened at the recovery site
- DFSMSHsm does not write multi-file tapes for its backup and migration tapes
 - f* Therefore, this line item does not apply to the backup and migration functions
- DFSMSHsm DUMP tapes contain dumps of DASD volumes
 - f* The likelihood of ever needing a file sequence number greater than 9999 is remote but the DUMP function has been updated to handle a file sequence number greater than 9999

18

© 2003 IBM Corporation

Miscellaneous FSN>9999 Support

- It will be possible to catalog and define a nonvsam tape data set with a file sequence number greater than 9999 using the JCL DISP parameter or IDCAMS DEFINE NONVSAM
- Checkpoint/Restart will support processing a tape data set opened with a file sequence number greater than 9999 when the checkpoint is taken
 - f* Change DALDSSEQ dynamic allocation text unit to increase the maximum parameter value to x'FFFF' (65,535) allowing users of dynamic allocation to specify >9999 file sequence numbers
 - Data Set Sequence Number Specification - Key ='001F'

To specify data set sequence number of 2

| KEY | # | LEN | PARM |
|------|-----|------|------|
| 001F | 001 | 0002 | 0002 |

Migration/Coexistence

- In order to use this new function either:
 - f* An existing application must be modified or a new application written which would issue the RDJFCB macro to obtain the JFCB, modify the JFCBFLSQ field to be the file sequence number required and then pass the updated JFCB to OPEN via OPEN,TYPE=J or...
 - f* Ensure the data set is cataloged allowing existing applications to process a data set with a data set sequence number greater than 9999
- Tape data sets created with a file sequence number > 9999 will not be able to be accessed on any lower level system
 - f* All files on the tape with a file sequence number not greater than 9999 will still be accessible on all system levels
- IBM Utilities, such as IEHMOVE, which allows specification of a seqno parameter are not being changed
 - f* The only way for IBM Utilities, to access a data set with a file sequence number greater than 9999 is if that dataset is cataloged

Migration/Coexistence (continued)

- Since the data set sequence number in the SL label may now not be in EBCDIC format, any application that prints the tape labels or accesses the tape labels directly may need to be modified to take that into account
- Coexistence PTFs are planned for this support for all lower level systems supported at the time of general availability
 - f* These PTFs will issue a new ABEND 613-76 and message IEC147I if an attempt is made to open a tape data set with a file sequence number greater than 9999
 - f* Before using DFSMSShsm ABARS to ACCOMPANY a data set on a user tape whose file sequence number is greater than 9999, the customer must make sure their recovery site will be at z/OS V1R5
 - Even though the data set will be properly cataloged at the recovery site by ARECOVER, a subsequent OPEN of that data set will fail with a new 613-76 ABEND

21

Multilevel Security

- Multilevel security is an optional feature which supplies a method by which the user may specify a new security level to be active on the operating system
- The changes to be provided in DFSMS when multilevel security name hiding is active are:
 - f* Access to the VTOC will not be allowed unless the user has sufficient access authority
 - f* LISTCAT processing will include a RACF call for every data set to be listed when a generic name is provided to determine if the user has read access to the data set
 - The data set will not be listed if this access authority is not available. A specific, fully qualified, name may be used to obtain information about a data set even if the user does not have sufficient access authority
 - f* DADSM Obtain uses CVAFDIR to obtain DSCB information for a data set
 - CVAFDIR, CVAFSEQ, and CVAFFILT macros will not return information about a data set specified using a generic name unless the user has sufficient access authority
- Changes to other components covered later
 - f* Changes have been made to DFSMSShsm which will allow the use of specific commands to not be generally available
 - f* DFSMSrmm is changed to use authorization via data set profiles
 - f* DFSMSdss functions may be limited by using RACF program control which is currently available

22

Migration/Coexistence

- When multilevel security name hiding is active, the data set names which are returned to the caller are restricted by the authority which the user has to access the data when the data set name is obtained from the VTOC or catalog
 - f* On a system without multilevel security name hiding active, all data set names may be retrieved using IEHLIST with LISTVTOC, or using IDCAMS with LISTCAT, will also be able to provide the data set names for data sets which were processed with that function
- With multilevel security name hiding active, it may be more difficult to determine the cause of a program failure
- Different users on different systems may be able to get different output based on the access authority each user has on each system
- The SETROPTS command is used to activate name hiding as well as to list the current RACF facilities which are in affect
- Performance will be impacted negatively when multilevel security is in use
 - f* The degradation is a result of the requirement to check every name which is obtained to determine if the user has read access to the file

23

Tape: Remove 8 Media Type Limit

- When the SMStape library support was introduced, it was architected to support a maximum of 8 media types and 15 recording technologies
 - f* Today in use are:
 - 4 of the 8 media types
 - MEDIA1/CST
 - MEDIA2/ECCST
 - MEDIA3/HPCT
 - MEDIA4/EHPCT
 - 4 of the 15 recording technologies
 - 18TRACK
 - 36TRACK
 - 128TRACK
 - 256TRACK
- New support for 255 media types and recording technologies

24

Migration/Coexistence

- No coexistence PTFs
 - f* Lower release levels will be able to coexist with the expanded library records
- Examine size of VOLCAT as library records are now larger
 - f* Typically the number of library records is small compared to the number of volume records so there will probably be enough room to accommodate the larger library record
- Users of the Catalog Search Interface (CSI) will need an additional 1976 bytes of storage to store the library record
- DFSMSHsm has a table of reuse capacities that will be changed with this support
 - f* Anyone patching the recycle reuse capacities will have to use the new offsets

25

BPAM Unix Files Support

- Read-only access USS files in addition to PDS and PDSEs
 - f* HFS, zFS or NFS
- Increase in BPAM concatenation limit
 - f* From 120 PDSE data sets and/or PDS extents to 255
 - f* Each USS directory will be counted as 1 to this limit
- Member level security
 - f* Checking for PDS and PDSE at OPEN
 - f* With USS files OPEN will verify that the user has UNIX search authority to each specified directory and FIND will verify user has UNIX read authority to the file
- Before z/OS DFSMS 1.5, High Level Assembler should work with UNIX files for any data set except SYSADATA and SYSLIB
 - f* With z/OS DFSMS 1.5 SYSLIB also will work

26

Programming Interfaces

- No change to existing read-only BPAM interfaces
 - f* Files in file directory will appear as the directory of a PDS or PDSE
- Concatenation of PDS, PDSE, and USS will be allowed in any order
 - f* DD statements within a concatenation can have either value of FILEDATA: BINARY or TEXT
- RDJFCB will enable application program to learn the name of the USS directory and path
- FIND macro will have new reason code 20 for return code 8 meaning that the file open failed for a reason other than "file not found"

27

© 2003 IBM Corporation

Migration/Coexistence

- Programs that processes names of data sets will see a dummy name of "...PATH=.SPECIFIED..." for each DD statement that is for a USS directory
 - f* Could change program to use RDJFCB to learn file directory name
- Programs that use the UCB address from the TIOT entry or the DEB will not find one for a USS directory
 - f* The address will be 0
 - f* Use RDJFCB
- Existing restrictions for BSAM access to USS files apply
 - f* Cannot depend on user data in directory or on particular block boundaries within a member
 - f* Binary variable length records do not retain record boundaries unless all records except for possibly the last record are the maximum length for the DCB
- BPAM reads ahead for PDSE and USS to provide better performance and to guarantee data integrity
 - f* The member cannot be deleted between the time a program starts reading it until the programs connection and BPAM closes the file
 - f* For long running programs, consider issuing STOW DISC after no further need for the member is expected
 - TTR from BLDL is no longer valid so reissue BLDL or FIND

28

© 2003 IBM Corporation



DFSMSdss



z/OS V1R5 DFSMSdss

- Exploitation of 64-bit Real Storage
- Fast Replication Support
- Physical DUMP and RESTORE of VSAM extended format data sets
- UIM changes

DFSMSdss Exploitation of 64-bit Real Storage

- Storage for I/O buffers backed anywhere in 64-bit real storage when z/OS is running in z/Architecture mode
- Can be disabled by EXEC PARM='ZBUFF64R=NO'
- In ESA/390 mode storage for buffers will continue to be obtained from below the 2GB line

31

© 2003 IBM Corporation

Fast Replication Support

- APAR OW57347 introduces new fast replication support
- Enhanced support for:
 - f* DATASET COPY
 - f* FULL/TRACKS COPY
 - f* DEFRAG
- Replication techniques supported:
 - f* FlashCopy specified invokes fast replication technique for IBM ESS
 - f* SnapShot specified invokes fast replication technique for IBM RVA and STK SVA
 - f* Fast replication technique used by any other vendor that supports the fast replication API
- New keywords on COPY and DEFRAG commands
 - f* FASTREPLICATION(REQ | PRE | NO)
 - f* DEBUG(FRMSG(MIN | SUM | DTL))
 - Improved informational messages (ADR918I) and reason codes for cases when FlashCopy cannot be used
- Support for future fast replication techniques supported by DFSMSdss

32

© 2003 IBM Corporation

Combinations

| | Fast Replication Required | Fast Replication Preferred | Fast Replication None |
|--------------------------|--|---|---|
| CONCURRENT | <p>Not applicable</p> <p>FASTREPLICATION(REQUIRED) and CONCURRENT keywords are mutually exclusive</p> | <p>ESS: DFSMSDss COPY will attempt to use FlashCopy.</p> <p>If FlashCopy fails, DFSMSDss COPY will retry with CC.</p> <p>If CC fails, DFSMSDss will retry using traditional I/O.</p> <p>RAMAC RVA: DFSMS dss COPY will attempt to use native SnapShot.</p> <p>If native SnapShot fails, DFSMSDss COPY will retry with virtual CC. If virtual CC cannot be used, DFSMSDss COPY will retry using traditional I/O.</p> | <p>ESS: DFSMSDss COPY will not attempt to use FlashCopy. DFSMSDss will attempt to use CC.</p> <p>If CC fails, DFSMSDss COPY will retry using traditional I/O.</p> <p>RAMAC RVA: DFSMSDss COPY will not attempt to use native SnapShot. DFSMSDss COPY will attempt to use virtual CC.</p> <p>If virtual CC cannot be used, DFSMSDss COPY will retry using traditional I/O.</p> |
| CONCURRENT not specified | <p>DFSMSDss will use SnapShot or FlashCopy (depending on the hardware)</p> <p>If SnapShot or FlashCopy fails, DFSMSDss will not retry using traditional I/O. Processing of the current data set or volume fails.</p> | <p>DFSMSDss will attempt to use SnapShot or FlashCopy (depending on the hardware).</p> <p>If SnapShot or FlashCopy fails,, DFSMSDss will retry using traditional I/O.</p> | <p>DFSMSDss will not attempt to use SnapShot FlashCopy. Traditional I/O is used.</p> |

33

VSAM EF

- Physical Data Set DUMP
- Physical data set RESTORE from Physical Data Set DUMP
 - f* Preallocated target data set will be checked to insure same number of stripes
- Physical Data Set RESTORE from FULL volume DUMP
 - f* Preallocated target data set will be checked to insure same number of stripes

34

UIM Changes

- Provide source VOLSER to UIM during function startup
- Provide "no volume serialization" option to UIM during function startup
- Provide "bypass security verification" option to UIM during function startup

Developed to support DFSMSHsm Fast Replication

Migration/Coexistence

- If DFSMSDss is invoked via API or XMAPI and the application cannot tolerate DFSMSDss buffers backed by 64-bit real storage PARM='ZBUFF64R=NO' must be specified in OPTPTR
 - f* Applications may require modification if:
 - Perform I/O directly out of DFSMSDss buffers
 - Page fix DFSMSDss buffers
 - Care about the real addresses of the DFSMSDss buffers
- Recompile and linkedit using new macro libraries:
 - f* Options Installation Exit Routine ADRUIXIT can change options, defaults, or value. UFOPARM contains bits as flags and a new flag UFPZB64R is used for 64-bit real storage specification
 - f* Exit point EIOPTION 13 can also modify UFOPARM
- Coexistence PTFs for DFSMSDss functions to support DFSMSHsm Fast Replication are not required or provided



IBM TotalStorage™

DFSMShsm

© 2003 IBM Corporation



IBM TotalStorage™

z/OS V1R5 DFSMSHsm

- ABARS CSI
- CDS/EA Under CDSQ
- Fast Replication Backup
- RACF Facility Class

38

© 2003 IBM Corporation

ABARS CSI

- Prior to this support, ABARS used SVC 26 to perform locates on fully and partially qualified data set names to determine:
 - f* If the data set was migrated
 - f* If the data set was on disk or tape
 - f* Volume serial locations
- Catalog Search Interface (CSI) is a faster alternative to SVC 26
 - f* Multiple SVC 26 requests may be combined into a single CSI call since CSI can return multiple fields
- No external options are required for DFSMSHsm to use CSI

39

CDS/EA Under CDSQ

- Today the following types of CDS serialization methods:
 - f* No systems serialization - a single DFSMSHsm host
 - f* CDSR - the entire volume is reserved for processing,
 - f* CDSQ - control data sets are serialized via enqueue,
 - f* RLS - VSAM Record Level Sharing (RLS) is used to provide locking at the record level eliminating the need for data-set-level serialization
 - f* CDSQ and CDSR can be used together
 - f* RLS cannot be combined with CDSQ or CDSR
- CDS size
 - f* MCDS and BCDS can be up to 4 VSAM clusters with maximum size of 16 GB
 - f* OCDS is limited to 1 VSAM cluster with maximum size of 4 GB
 - f* None of the above can be multivolume
 - f* VSAM Extended Format data sets with the Extended Addressability attribute can be $Csize * 4GB$
 - Currently DFSMSHsm only supports VSAM/EA for CDSs accessed in RLS mode
- Control data sets serialized with CDSQ can now be VSAM/EA

40

CDS/EA CDSQ Migration/Coexistence

- Allocation of the EA CDS and copying of the contents of the existing CDS to it is a one-time process
 - f* EA CDSs must be SMS-managed and assigned to a DATA CLASS that specifies Extended Format and Extended Addressability
 - f* Use IDCAMS DEFINE CLUSTER and IDCAMS REPRO commands
- Mixing EA and non-EA clusters will be allowed since each of them is treated as separate entity
- Multicluster CDS will be able to consist of EA and non-EA clusters with this support
- Dynamic multivolume under CDSQ will be supported
 - f* "Dynamic Volume Count" feature of SMS Data Class will be used to add a new volume automatically
- All sharing systems must be at a supported level to access a CDS using EA with CDSQ

DFSMSHsm Fast Replication Support

- DFSMSHsm will invoke DFSMSdss to invoke Fast Replication Support
- Introduces new SMS 'copy pool' construct
 - f* A set of storage groups that will be processed collectively for fast replication functions
- Introduces full volume fast replication backup versions managed by DFSMSHsm
- Recovery can be performed at the volume or copy pool level but not the data set level

Copy Pool Construct

- New SMS construct Contains the names of up to 256 storage groups that should be processed collectively for fast replication operations
 - ƒ Actual volumes are obtained from each storage group during subsequent processing
 - Volumes processed could be different than the volumes that were in the storage groups at the time of the definition of the Copy Pool
 - For example, adding a volume to a storage group that is already in a Copy Pool will be processed
 - ƒ Optional parameter to specify how many backup versions of the pool that DFSMSHsm should maintain
 - ƒ Name can be up to 30 characters and the first character can be an uppercase alphabetic or national character
 - Remaining characters can be uppercase alphabetic, national, or numeric
 - ƒ An individual storage group can be in up to 50 Copy Pools
 - ƒ As many as 85 backup versions may be maintained for each Copy Pool
 - Each version requires a unique target volume for each source volume
 - Recommended to have at least 2 versions
 - If business requires n versions they should consider $n+1$ versions
 - Before a new version is created, the oldest version is invalidated and its target volumes are overwritten with the new version

43

Copy Pool Backup Storage Group Type

- New storage group type for DFSMSHsm use - Copy Pool Backup
 - ƒ In addition to Pool, VIO, Dummy, Tape, Object and Object Backup
 - ƒ A container for the volumes that DFSMSHsm selects as fast replication target volumes
- Copy Pool Backup storage group can not be assigned by storage group ACS routines
 - ƒ SCDS validation will fail
- DFSMSHsm does not look at the status of the Copy Pool
 - ƒ SMS looks at status of storage groups during data set allocation
 - ƒ Prevent DFSMSHsm from processing a volume in a Copy Pool Backup storage group by performing a vary offline on that volume
- New field in Pool storage group definition for associated Copy Pool Backup storage group
 - ƒ The Copy Pool storage group can be shared by other Pool storage groups
- An eligible target volume must satisfy the requirements of the fast replication technique invoked

44

Extend and Overflow Storage Groups

- Introduced with z/OS DFSMS 1.3
- Enables data sets to extend or be allocated onto volumes that are not part of the main source pool storage group
- Users should ensure that all Extend and Overflow storage groups are included in the storage group list for the appropriate Copy Pools
 - f* SMS does not verify that Extend and Overflow storage groups associated with main source pool storage groups have been included in the Copy Pool definition
 - f* The volumes in the Extend and Overflow storage groups also require targets

DFSMSHsm Commands

- FRBACKUP
 - f* Create a fast replication backup version for each volume in a specified Copy Pool
- FRRECOV
 - f* Use fast replication to recover a single volume or a pool of volumes from the managed backup versions
- FRDELETE
 - f* Delete one or more unneeded fast replication backup versions

Preparing for DFSMSHsm Fast Replication

- FRBACKUP optional keyword PREPARE
 - f* Validate the fast replication environment
 - f* Reduces the elapsed time of the actual Copy Pool backup
 - f* Required number of target volumes preassigned to each source volume in each storage group specified in the Copy Pool
 - Validate that there are enough target volumes for backup
 - Moves target volume selection outside fast replication window
 - f* PREPARE function performed whenever there is a change in the environment
 - Volumes in the Pool or Copy Pool Backup storage group are changed or removed
 - Number of versions to be maintained is changed
 - Storage groups defined to the copy pool have changed

47

© 2003 IBM Corporation

Creating a Fast Replication Backup

- FRBACKUP command will initiate a fast replication backup
 - f* Not supported as part of automatic backup or dump
- Invoked by:
 - f* Entering the command at the console
 - f* HSEND TSO command
 - f* Batch or programmatically via the ARCHSEND macro interface
- DFSMSHsm will invoke DFSMSdss COPY FULL function
- Command successful only after fast replication relationship established for each source volume otherwise:
 - f* Entire function is considered a failure
 - f* DFSMSHsm continues process to determine if any remaining volumes will also encounter an error
 - f* For each volume in error the target volume will be disassociated from the source volume
 - f* After all volumes are processed:
 - For FlashCopy relationships already established will be withdrawn
 - The version is marked as invalid

48

© 2003 IBM Corporation

Creating a Fast Replication Backup (continued)

- TOKEN parameter of up to 40 bytes can be specified and identifies the version
 - f* Can be used to recover a particular version
 - f* TOKEN returned as part of the List and ARCXTRCT data for the Copy Pool
 - f* DB2 will use this parameter to pass an 8 byte token that will be used during their recovery process

FlashCopy and Fast Replication

- With FlashCopy, recovery from a fast replication backup cannot be done if any volume is in a FlashCopy relationship
- Recovery of copy pool should not be attempted until all background copies have completed from last FRBACKUP
 - f* QUERY COPYPOOL command can be used to determine if any of the source volumes for a copy pool are in a FlashCopy relationship
 - f* FRRECOV VERIFY keyword can be used to prevent initiation of a recovery where all of the background copies from the backup have not completed
 - With VERIFY(Y) DFSMSHsm will first determine if any volume is in a FlashCopy relationship before initiating any volume recovery
 - If one or more volumes are in a relationship the recover will be failed
 - If an immediate recovery needs to be performed, there are two options:
 - ♦ Wait until all background copies have completed and then the latest backup version can be used for the recovery
 - ♦ Withdraw outstanding FlashCopy relationships using new FRBACKUP COPYPOOL(cpname) WITHDRAW command which invalidates the latest version of the backup but will enable an immediate recovery from prior version
- Warning: Users should issue LIST COPYPOOL command or ARCXTRCT macro to ensure they do not invalidate the only valid fast replication backup version which would prevent any recovery from being performed

Deleting Unneeded Backup Versions

- Normal processing will have new backup versions replacing older backup versions
- Two cases where processes other than normal roll-off deletion is needed:
 - f* When number of versions specified for a copy pool is decreased
 - Next time FRBACKUP COPYPOOL is issued for copy pool whose number of versions has been decreased the unneeded versions will be deleted and the target volumes made available for other use
 - f* When a copy pool is no longer needed
 - Use FRDELETE COPYPOOL to delete all backup versions
- If VERSIONS or TOKEN is specified the named versions will be deleted
- Before a version is deleted all outstanding FlashCopy relationships will be withdrawn

51

© 2003 IBM Corporation

Obtaining Fast Replication Info - QUERY

- QUERY command
- List command
- Report command
- ARCXRCT macro

52

© 2003 IBM Corporation

DFSMSHsm and Target Volumes

- Use DFSMSdss to backup target volumes
 - ƒ DFSMSHsm AUTODUMP or BACKVOL DUMP will skip the target volumes
- DFSMSHsm must prevent many of its functions from processing these new entities
 - ƒ Storage Group processing
 - Automatic functions will only process POOL type storage groups
 - Primary space management
 - Autobackup
 - Autodump
 - Interval migration
 - ƒ Volume processing
 - BACKVOL VOLUME will receive ARC1622I for target volumes
 - MIGRATE VOLUME will receive ARC1622I for target volumes
 - ƒ User commands do not support copy pool functions
 - HLIST, HQUERY, and HRECOVER

53

© 2003 IBM Corporation

Fast Replication Security

FRBACKUP

| Facility Class Resource | Authorized Function |
|------------------------------|--------------------------|
| STGADMIN.ARC.FRBACKUP | Any FRBACKUP command |
| STGADMIN.ARC.FRBACKUP.cpname | FRBACKUP only for cpname |

FRRECOV

| Facility Class Resource | Authorized Function |
|-----------------------------|---|
| STGADMIN.ARC.FRRECOV | Any FRRECOV command |
| STGADMIN.ARC.FRRECOV.cpname | FRRECOV COPYPOOL and FRRECOV TOVOLUME only for cpname |

FRDELETE

| Facility Class Resource | Authorized Function |
|------------------------------|-----------------------------------|
| STGADMIN.ARC.FRDELETE | Any FRDELETE command |
| STGADMIN.ARC.FRDELETE.cpname | FRDELETE COPYPOOL only for cpname |

54

© 2003 IBM Corporation

Function Statistic Record (FSR)

- **FRBACKUP**
 - f* New FSR type 21 for each volume that is backed up as part of the Copy Pool
- **FRRECOV**
 - f* New FSR type 22 will be created for each volume that is recovered as part of the copy pool
 - f* The new type 22 record will also be used when FROMCOPYPOOL is used
- **FRDELETE**
 - f* New FSR type 23 for each version that is deleted
- In addition to setting the standard fields (date, time, etc.):
 - f* FRBACKUP, FRRECOV, and FRDELETE
 - FSRGEN
 - FSRFVER
 - f* FRBACKUP and FRRECOV
 - FSRFNONQ
 - FSRTVOL
 - FSRFVOL

55

Support Use

- **DB2**
 - f* Storage Administrator establishes copy pool environment using required DB2 naming conventions for the copy pools
 - f* Issue FRBACKUP PREPARE against each copy pool
 - f* DB2 Administrator uses the DB2 utility to create fast replication backups of the data base
 - The DB2 utility performs the necessary serialization for the backup to be valid
- **Non-DB2**
 - f* Storage Administrator establishes copy pool environment
 - f* Issue FRBACKUP PREPARE against each copy pool
 - f* Copy pool owners issue FRBACKUP COPYPOOL command to create point-in-time backups
 - NOVTOCENQ is NOT specified because there is no outside serialization being done

56

Catalog Considerations

- Fast replication backup is a physical copy thus care should be taken to ensure that catalogs stay synchronized with the data
 - ƒ Catalogs should be copied with the data that they represent
- If it is necessary to recover an individual volume and not the corresponding catalog then be aware:
 - ƒ Data sets may exist on the volume but not in the catalog
 - ƒ The catalog may have entries for data sets on the volume that are not actually present
 - ƒ Data set information for existing data sets in the catalog may be inaccurate

57

© 2003 IBM Corporation

DB2 V8 System Level PiT Recovery

- Provide an easier and less disruptive way for fast volume-level backup and recovery
- Two new utilities:
 - ƒ BACKUP SYSTEM
 - ƒ RESTORE SYSTEM
- Backup copies can also be used for:
 - ƒ Disaster recovery
 - ƒ System cloning
- Certain activities are disabled during the backup stage
- Enhancement in the SET LOG SUSPEND command
 - ƒ As an alternative to BACKUP SYSTEM

58

© 2003 IBM Corporation

Migration/Coexistence

- Coexistence PTFs will be available for DFSMSHsm
- Systems performing recovery of fast replication volumes are at the supported level of DFSMSHsm
- Backup current SCDS before modifying the SCDS on the newer release
- Modify and validate the SCDS on the highest level system
- DFSMSHsm BCDS must have a max record size of 6544
- Identify storage groups for DFSMSHsm fast replication backup
- Determine which storage groups should be in same copy pool
- Update and validate SMS configuration
- Perform FRBACKUP COPYPOOL(cpname) PREPARE

59

RACF FACILITY Class

- Currently DFSMSHsm has authorized and nonauthorized commands
 - f* Authorized commands only available to a user authorized by storage administrator with the AUTH command
 - Can affect data sets not owned by the person issuing the command
 - f* Nonauthorized commands can be issued by any user
 - Generally only affect data sets for which end user is authorized
 - Two categories: USER and CONTROL
 - USER can issue any DFSMSHsm command except that command that authorizes other users
 - CONTROL can issue any DFSMSHsm command and AUTH command
- New RACF Facility class profiles allow nonauthorized users to issue the ABACKUP and ARECOVER commands
 - f* If certain facility classes are defined they override the AUTH level of the user
- All DFSMSHsm commands can now be protected by RACF Facility class profiles both generic and discrete
 - f* RACF Facility Class must be activated

60

Base Set of RACF Facility Classes

| | |
|--|---|
| STGADMIN.* | System level storage administrator command protection. Generic profile provides default access if other DFSMSHsm profiles are not defined |
| STGADMIN.ARC.* | DFSMSHsm command protection, generic profile for all DFSMSHsm commands |
| STGADMIN.ARC.command | DFSMSHsm authorized command protection, discrete profile for specific DFSMSHsm authorized command |
| STGADMIN.ARC.command.parameter | Discrete profile protects specific DFSMSHsm authorized command with specific parameter |
| STGADMIN.ARC.ENDUSER.* | DFSMSHsm end user command protection |
| STGADMIN.ARC.ENDUSER.h_command | DFSMSHsm end user command protection, discrete profile protects specific DFSMSHsm end user command |
| STGADMIN.ARC.ENDUSER.h_command.parameter | Discrete profile protects specific DFSMSHsm end user command with specific parameter |

RACF Authorized Commands

- STGADMIN.ARC.* can be used to protect all DFSMSHsm authorized commands
 - f* User or group requires ACCESS(READ) to issue command
 - f* ACCESS(NONE) means that user or group can't issue command
- STGADMIN.ARC.h_command or STGADMIN.ARC.h_command.parameter can be used to restrict the use of any authorized command

Migration/Coexistence

- Can coexist with previous releases of DFSMSHsm
 - Toleration maintenance will be made available

DFSMSrmm

z/OS V1R5 DFSMSrmm

- Support added by APAR OW56584 for z/OS 1.3 and incorporated in z/OS 1.5
 - ƒ Backup at Any Time
 - ƒ Duplicate VOLSER Support
 - ƒ OAM: Improved Volume Management
 - ƒ Multilevel security
- New function since z/OS 1.3
 - ƒ OW53763 5/02 Enhanced foreign tape authorization
OW55215
 - ƒ OW53989 6/02 3590-H support
 - ƒ OW55327 7/02 CONTROL-T conversion tools
 - ƒ OW55698 8/02 EDG4026I includes dsnames
 - ƒ OW54057 8/02 VTS Advanced Policy Management
 - ƒ OW55115 9/02 UNCATALOG single volume in multi-vol set
 - ƒ OW56727 10/02 Default MEDIANAME
 - ƒ OW57336 11/02 Zara/Automedia conversion tools
 - ƒ OW56172 12/02 New ABEND support for OPEN failures
OW56496
 - ƒ OW57516 01/03 Control-T R6 conversion support
 - ƒ OW56584 04/03 SPE for z/OS R3
 - ƒ OA02095 2Q03 EXPDT for data sets
 - ƒ OA02094 04/03 Report Generator Enhancements

65

© 2003 IBM Corporation

Backup at Any Time

- Prior to this support, if inventory management was already running, you could not start a backup or vice versa
- Backup can now run at the same time as inventory management or vice versa
 - ƒ When backup runs at the same time the inventory management processing tolerates the backup and waits until updates can be made to the control data set
 - Only CDS updates may have to wait as processing can continue for volumes and data sets that need no updates until the point an update is required
 - If non-intrusive backup (concurrent) is used CDS updates do not have to wait
- Journal threshold can start backup even if inventory management is running
- Option to Backup and Clear Journal without CDS backup
 - ƒ Adjust RESTORE jobs to use additional journal backups
- EDG0123D can be issued twice at startup time
 - ƒ Once for backup in progress and once for inventory management
- If you have used the RMM application definition for OPC update applications RMMBKP and RMMPOST and remove the dependency on the special resource RMM.HOUSE.KEEP

66

© 2003 IBM Corporation

Duplicate VOLSER Support


- Prior to this support, duplicate volumes were processed outside of DFSMSrmm management by using ignore processing via EDGUX100
- For private physical volumes you can define the VOL1 label parameter in addition to the unique volser
- DFSMSrmm open processing has been updated to detect duplicate volumes and use the correct CDS volume information to manage the volume
- ƒ It is possible to continue to ignore duplicate volumes but it is recommended to add them to the DFSMSrmm CDS and manage them
- Conversion extract records are updated to include the VOL1 label volser and EDGCNVT converts this new field into the volume record
- Support does not have to be enabled
 - ƒ Use ADDVOLUME or CHANGEVOLUME command with VOL1(volser) to define duplicate volser
- An NL tape cannot be a duplicate volume
 - ƒ Such a volume can be defined using any unique value for the volser and no special processing is required by DFSMSrmm to distinguish the volume from another identical volume

67

© 2003 IBM Corporation

Duplicate Volume Open Processing

```
//DD1 DD VOL=SER=ABC123
//DD2 DD VOL=SER=D00005
```



- The requested volume is the unique volser you specify in your JCL or allocation request to DFSMSrmm
- DFSMSrmm uses the Tape Mount exit to check the VOL1 label volser
- DFSMSrmm uses the requested volser to find a volume defined to DFSMSrmm and validates that the VOL1 label volser matches the VOL1 label defined to DFSMSrmm
- When DFSMSrmm believes the correct volume is mounted the VOL1 label volser will be substituted with the volser specified by user
 - ƒ Ensures that OPEN processing will accept the volume and that VERIFY processing will be successful
- At OPEN for output time when the VOL1 label is rewritten, DFSMSrmm intercepts the request in the Tape Mount exit and does reverse substitution

68

© 2003 IBM Corporation

Duplicate Volume Inventory Management

- DFSMSrmm ensures that duplicate volumes released to scratch are relabelled so the VOL1 label volser matches the volser
 - ƒ DFSMSrmm will prevent reuse of duplicates
- Automatic process with no manual intervention by user
- Librarian/user can aid this by ensuring that the release action is SCRATCH or REPLACE
 - ƒ If RETURN is specified DFSMSrmm will eventually delete the duplicate
 - ƒ Any volume returning to the scratch pool which is set pending release or which is pending release and still has a VOL1 label volser that does not match the volser number has the INIT action set automatically during expiration processing
 - Prevents scratch pool volumes having mismatching volsers
- Global confirm INIT and ERASE clear the VOL1 label volser for a duplicate volume pending return to scratch on the assumption that the librarian has taken care of the relabeling of the volume

69

© 2003 IBM Corporation

OAM: Improved Volume Management - Phase 1

- DFSMSrmm will now intercept the new OAM message CBR2165I which is issued when a volume is released from OAM control
- A parameter list will be built using the volser for EDGTVEXT which will then release the volume

70

© 2003 IBM Corporation

Multilevel Security

- Support for the RACF multilevel security option MLNAMES
 - f* Can also be controlled via parmlib OPTION COMMANDAUTH(OWNER | DATASET)
- RACF DATASET class profiles control use of:
 - f* LISTDATASET, SEARCHDATASET, CHANGEDATASET, LISTOLUME, SEARCHVOLUME, CHANGEVOLUME, DELETEVOLUME
- DFSMSrmm uses the data set name of the first file of a volume for volume commands
 - f* If no first file data set name is available and the volume is MASTER or USER status, the TAPEVOL class is used
 - Otherwise only CONTROL access to STGADMIN.EDG.MASTER allows a volume to be processed by command
- With MLNAMES or COMMANDAUTH(DSN) in use:
 - f* SEARCH: Any entry which the user has no authority to list is skipped
 - f* SEARCH or LIST: User must have READ access to MASTER
 - f* DELETEVOLUME RELEASE: User must have READ access to RELEASE or MASTER
 - f* CHANGE: DATASET or TAPEVOL class authorization only allows updates to non-restricted fields in the records and UPDATE access to MASTER is required

71

© 2003 IBM Corporation

Catalog Customer Satisfaction

© 2003 IBM Corporation

Customer Catalog Satisfaction

- Enablement/disablement of JOBCAT/STEPSCAT support according to installation choice
- Enhance program diagnosis by creating SYMREC records for select catalog errors that are detected
- Provide support to terminate a Catalog service task that shows waiting in RECALL
- Reduce the likelihood of purging BCS caches on extremely active shared catalogs
- Allow the user to indicate they want ECS AUTOADD established on the first connection to the Coupling Facility
- Allow the installation to specify when a warning should be issued about catalogs that are becoming full

73

© 2003 IBM Corporation

Selectable Support of JOBCAT/STEPSCAT

- Use of JOBCAT/STEPSCAT now optional at the installation level
 - f* Default is to no longer honor JOBCAT/STEPSCAT
- New parameter to MODIFY CATALOG command:
 - f* F CATALOG,ENABLE(JOBSTEPSCAT)
 - f* F CATALOG,DISABLE(JOBSTEPSCAT)
- Current setting retained across CAS restart
- If support is disabled and a catalog request is issued with JOBCAT/STEPSCAT will cause message IDC3009I with a return code 54 reason code 6 to be issued and the request will not be processed

74

© 2003 IBM Corporation

Creation of SYMREC Records for Diagnosis

- SYMREC records are currently written when an SVC dump is created by CAS
- There are times when catalog return code and reason code cannot uniquely identify the specific problem and now SYMREC recording will enable a user to resolve certain problems without contacting IBM support
- The use of SYMREC may also reduce the need for some SVC dumps
- Overall creation of SYMREC records can be enabled or disabled
 - f* F CATALOG,ENABLE(SYMREC)
 - f* F CATALOG,DISABLE(SYMREC)
- Cannot control individual cases when SYMREC records will be written

75

© 2003 IBM Corporation

Termination of Service Task in Recall

- MODIFY CATALOG can currently terminate a service task abnormally with options END or ABEND
 - f* But if a service task is currently active and waiting on processing in the user address space to proceed the abnormal termination will not be performed
- New keyword for MODIFY CATALOG:
 - f* F CATALOG,END(xx){,REDRIVE | NOREDRIVE}{[,FORCE]}
 - f* F CATALOG,ABEND(xx){[,FORCE]}
- User should be aware that they should only use this capability when they know that the address space or task the service task is operating on behalf of has terminated

76

© 2003 IBM Corporation

Reduce Purges of VLF Cache for Shared Catalogs

- It is possible that if there is long time between catalog references that CAS will purge the entire BCS cache before continuing due to the number of changes to the catalog
- Table containing entries has been made more efficient thereby reducing the chance of the cache being purged
- Applicable only to VLF CDSC and not ISC

77

© 2003 IBM Corporation

Automatic AUTOADD for First ECS Connector

- Currently an operator command must be issued after the first system in a parallel sysplex is IPLed to enable the ECS AUTOADD function
- This support allows ECS AUTOADD enablement without an operator command
 - f* SYS1.NUCLEUS(SYSCATxx) column 63 specify 'Y'
 - f* SYSCAT statement in LOADxx column 72 specify 'Y'
- The ability to set the AUTOADD state by the MODIFY CATALOG command is still supported

78

© 2003 IBM Corporation

Operator Warning on Catalog Usage Threshold

- Currently there is no warning when a catalog is about to run out of space
 - f* Causing data set allocation failures
 - f* Available space within a catalog must be manually monitored
- New enhancement will allow a percentage of maximum extents to be specified that when reached will cause a notification message to be issued
 - f* IEC361I CATALOG catalogname (comptype) HAS REACHED xxx% OF THE MAXIMUM
 - f* Issued once per catalog per extent and reissued for each subsequent catalog extent
 - f* comptype is either DATA or INDEX
- Invoked by issuing F CATALOG,NOTIFYEXTENT(yyy) where yyy is a value from 0-100
 - f* Default is 80
 - f* A setting of 0 or 100 will disable the function
- F CATALOG,ALLOCATED will display the percentage of allocated extents for each catalog in the list in message IEC348I
 - f* Percentage displayed is the higher of DATA or INDEX component
- F CATALOG,REPORT will display current setting of the NOTIFYEXTENT value
 - f* Message IEC359I

79

SMS Availability/Usability Enhancements

SMS Availability/Usability Enhancements

- GDS Reclaim Processing
- Save ACDS as SCDS
- High Threshold Messages
- Multi-Tiered Storage Groups
- End of Volume (EOV) Failure Messages

81

© 2003 IBM Corporation

GDS Reclaim Processing

- If a new (+1) generation of an SMS managed GDS data set is created but does not get rolled in (system crash, address space termination, etc.) the next job that attempts to create a (+1) will cause SMS to go into GDS reclaim processing
 - f* Could overlay and destroy data left in deferred roll-in state
- New enhancement allows customer to turn 'off' automatic GDS reclaim processing
 - f* Parmlib member IGDSMSxx will have a new keyword:
 - GDS_RECLAIM {YES | NO}
 - If NO is specified, manual means must be used to either delete the generation, rename it, or roll it in
 - f* SET SMS=xx to change IGDSMSxx member being used or....
 - f* SETSMS GDS_RECALIM {YES | NO} to change value
 - Not retained across an IPL
- Only applies to SMS managed data
- Different systems can have different values set
 - f* Installations that have different levels of SMS or different settings should not share GDSs across all of the systems

82

© 2003 IBM Corporation

Save ACDS as SCDS

- Relieves the tedious work of recreating the SCDS if the source SCDS and all of its backups are lost
- Now able to save ACDS as SCDS
- New command:
 - f* SETSMS SAVESCDS(scds_dsname)
- SMS will verify that the 'scds_dsname' is not the currently active ACDS or COMMDS

83

© 2003 IBM Corporation

High Threshold Messages

- There is currently no alert issued if available space in a pool storage group is above its high allocation threshold
- Messages will now be issued to the hardcopy log when a pool storage group exceeds its high allocation threshold during a new allocation or new volume extend
 - f* Cumulative space in a storage group includes all of the volumes which are online and enabled or quiesced to SMS
- Message will not be issued if there:
 - f* Is an uninitialized volume in the storage group
 - f* When extent reduction function is used by DFSMSDss/DFSMSHsm for nonvsam allocations
- Message will not be issued again until utilized space falls below 80% of the high threshold and exceeds the high threshold again

84

© 2003 IBM Corporation

Multi-Tiered Storage Groups

- The best volume from any of the supplied storage group will currently be selected for allocation
- With Multi-Tiered Storage Groups enabled the user will be able to specify a storage group order for pool storage groups
- New field added to storage class definition
 - f* "Multi-Tiered SGs" {YES | NO}
 - f* Honors storage group sequence order specified in ACS storage group selection routines
- Only volumes in the first storage group are eligible to be primary volume candidates

85

© 2003 IBM Corporation

End of Volume (EOV) Failure Messages

- Not all requesters of space on a new volume request that messages be generated to identify the exact reason for a failure
 - f* DADSM for nonvsam and VSAM EOVS
- Enhancement insures that SMS EOVS will generate messages regardless of whether the caller requested them or not
- Messages written to JOBLOG and hardcopy log

86

© 2003 IBM Corporation

Migration/Coexistence

- No specific considerations

VSAM Customer Satisfaction

VSAM Customer Satisfaction

- Enhanced Data Integrity of VSAM Data Sets with AIXs
- VSAM Message Enhancements
- VSAM Extent Reduction

89

© 2003 IBM Corporation

Enhanced Data Integrity of VSAM Data Sets with AIXs

- Prior to this enhancement, if LSR buffers were in short supply and there were numerous AIXs, VSAM would attempt to backout prior updates but cannot do it due to the insufficient supply of buffers that started the problem
 - f* Result was AIXs not in synchronization with the base VSAM cluster
 - f* "Using Data Sets" does say that there must be enough strings (PLHs) to complete a given request
- Enhancement will cause Record Management to not begin upgrade processing if sufficient resources are not available
 - f* The application will be notified and given the option of retrying the failed request

90

© 2003 IBM Corporation

VSAM Message Enhancements

- Multiple single line VSAM message WTOs are converted to a single multi-line WTO
 - f* IEC161I - OPEN
 - f* IEC251I - CLOSE
 - f* IEC252I - TCLOSE
 - f* IEC070I - EOVS
 - f* Avoids intervening messages between the multiple single line messages and allows the operator to relate the various pieces of information on each line
 - f* Allows automation programs to take action based on the composite message
- VSAM OPEN message IEC161I 052(sfi)-xxx enhanced to provide the JOBNAME of the first JOB causing a SHAREOPTIONS violation
 - f* IEC161I 052(sfi[jobname])-xxx
- VSAM messages IEC070I 032 and IEC070I 034 enhanced to provide sfi when a catalog LOCATE request fails

91

© 2003 IBM Corporation

VSAM Extent Reduction

- Limit is 255 extents
- New VSAM EOVS path will consolidate adjacent extents when extending on the same volume
 - f* Saves 1 on the extent count
- Support only for SMS managed data sets
 - f* Not supported:
 - Data sets with the IMBED attribute
 - Data sets with the REPLICATE attribute
 - Keyrange data sets
 - Page data sets
 - Catalogs
 - VVDS
 - VSAM data sets accessed in RLS mode

92

© 2003 IBM Corporation

Migration/Coexistence

- VSAM messages
 - f* Software that is dependent upon the old format of the VSAM messages need to be reviewed
 - f* Software that processes the old multiple single line VSAM messages need to be reviewed

SAM 59 Stripes

- Nonvsam Extended Format data sets will be allowed to have up to 59 stripes
 - f* Current limit is 16
- Alleviates the fact that SAM Extended Format data sets cannot extend beyond the primary volume
 - f* Provides more flexibility to spread the data over more volumes
- SAM APAR OW48983 will be required on downlevel releases in order to process > 16 stripe data sets
 - f* DFSMS/MVS V1R5
 - f* OS/390 DFSMS V2R10
 - f* z/OS DFSMS V1R3
- DFSMSHsm on downlevel releases will not be able to migrate/recall as a 59 stripe data set
 - f* Will migrate 59 stripe data set but will be recalled as a 16 or less striped data set

z/OS DFSMS 1.5 Summary

- Improved availability (Business Continuance)
 - ƒ Multi-Tiered Storage Groups
 - ƒ VSAM Extent Reduction
 - ƒ Enhanced Integrity of physical sequential data sets
 - ƒ DFSMSShsm Fast Replication
- Better performance
 - ƒ SAM 59 Stripe support
 - ƒ JOBCAT/STEPDAT selectable support
 - ƒ DFSMSShsm ABARS using CSI
- Improved System Scalability
 - ƒ Remove limit on tape media types
 - ƒ Extend limit of file sequence number
 - ƒ z/OS NFS support for legacy multivolume data sets
 - ƒ DFSMSShsm support for larger CDSs without using RLS access
- Ease of Use/Automation (Usability)
 - ƒ Prevent overlay of a GDS in deferred roll-in state
 - ƒ Enhanced messages
 - ƒ SYMREC records for diagnostics
 - ƒ Save ACDS as SCDS

95

© 2003 IBM Corporation

z/OS V1R6 DFSMS

© 2004 IBM Corporation

Table of Contents

- Session Objectives
- z/OS V1R6 DFSMS New Functions
 - DFSMSdfp
 - SMS volume selection based upon Parallel Access Volume
 - PDSE restartable address space
 - MLS SECLABEL in ACS routines
 - Miscellaneous line items
 - DFSMSdss
 - DFSMSdss Replace Unconditional
 - DFSMShsm
 - DFSMShsm Secondary Space Management Multitasking Enhancements
 - DFSMSrmm
 - DFSMSrmm as server for multiple systems
 - DFSMSrmm ISPF usability
- Session Summary
- Appendix

97

© 2003 IBM Corporation

DFSMSdfp

98

© 2003 IBM Corporation

z/OS V1R6 DFSMSdfp

- SMS volume selection based upon Parallel Access Volume
- PDSE restartable address space
- MLS SECLABEL in ACS routines
- Miscellaneous line items

SMS Volume Selection based on PAV

- This new support is intended to give the user additional control over the Volume Selection process for a SMS-managed data set depending on whether a base volume has a PAV or PAVs associated with it.
 - Allows allocations to be automatically allocated to high performance devices.
 - PAVs allow concurrent access to a device eliminating IOSQ time.
- New field named PAV Capability provided in Storage Class definition.
- Changes made to volume selection algorithm to take PAV into consideration.
- ISMF and NavQuest updated to support the new function.

Parallel Access Volume Capability

- R (REQUIRED)
 - PAV capability is required
 - Only those volumes that support this capability will be eligible
 - All other volumes will be rejected from consideration
- P (PREFERRED)
 - PAV capability is preferred
 - Volumes with this capability will be preferred over volumes that do not have this capability
- S (STANDARD)
 - Volumes without PAV capability will be preferred over volumes that have this capability
- N (NOPREFERENCE)
 - Volumes with or without PAV capability will be equally considered for volume selection
 - This is the default

101

© 2003 IBM Corporation

Current Volume Selection

- Primary volumes
 - Meet or exceed all user specified preferences and requirements
 - Volumes are enabled
 - Satisfy primary space amount without exceeding Storage Group high threshold
 - Volumes belong to a Storage Group that is enabled and able to satisfy stated volume count
- Secondary volumes
 - Volumes that do not meet one or more of the criteria for primary volumes
 - If no primary volume is available a volume is selected from the secondary list based on space, status, and features requested
- Tertiary volumes
 - Volumes which belong to a Storage Group that does not contain enough eligible volumes to satisfy volume count
 - If no primary or secondary volume is available a volume is selected from the tertiary list based on space, status, and features requested
 - Tertiary volumes are not supported for VSAM allocations

102

© 2003 IBM Corporation

Volume Selection Algorithm Changes

- N (NOPREF)
 - PAV Capability ignored so no effect on volume selection process
- S (STANDARD)
 - Volumes without PAV Capability will be preferred
 - Volumes with PAV Capability will be secondary volumes
- P (PREFERRED)
 - Volumes with PAV Capability will be preferred
 - Volumes without PAV Capability will be secondary volumes
- R (REQUIRED)
 - Volumes without PAV Capability will be rejected
 - All eligible volumes will have PAV Capability

103

© 2003 IBM Corporation

Attribute Relative Importance

Applicable to secondary and tertiary volume lists

1. Data Set Separation
2. Volume count
3. High threshold
4. SMS status
5. During EOVS processing the volume is in the Primary SG rather than the Extend SG
6. Non-overflow SG
7. IART
8. Fast Replication
9. Accessibility
10. **PAV Capability**
11. Availability
12. Extended format
13. MSR

104

© 2003 IBM Corporation

Additional Support

- ISMF
 - Storage Class Display
 - Storage Class List
 - Storage Class List Print
 - Storage Class List Sort
 - Storage Class List View
- NaviQuest
 - New NaviQuest field in sample JCL ACBJBAS1 to enable users to define, alter or display Storage Class

PAVCAP()

105

© 2003 IBM Corporation

Migration/Coexistence

- No specific migration considerations
- No specific coexistence considerations
 - Installations should be aware that a Storage Class may be defined on an up-level system to specify PAV Capability
 - Lower level systems will ignore this field
 - No toleration PTFs are required
 - An IPL with CLPA is required to start using this new support

106

© 2003 IBM Corporation

PDSE Restartable Address Space

- Improves PDSE reliability and availability
- Eliminates need to re-IPL a system due to a hang, deadlock condition, or out of storage condition
- No changes are required to JCL or programs in order to utilize this new function

107

© 2003 IBM Corporation

PDSE Address Spaces

- Address space SMSPDSE will remain non-restartable
- New address space SMSPDSE1
 - There will either be a single SMSPDSE address space or both SMSPDSE and SMSPDSE1
 - SMSPDSE will be the only address space if:
 - PDSESHARING(NORMAL) is defaulted or specified in IGDSMSxx or....
 - PDSESHARING(EXTENDED) is specified and PDS_RESTARTABLE_AS(NO) is defaulted or specified in IGDSMSxx

108

© 2003 IBM Corporation

SMSPDSE1

- Provides connections to and processes requests for those data sets that are not part of the global connections associated with SMSPDSE
 - Global connections include certain connections to PDSEs in the LNKLIST concatenation
- User programs will maintain the connection to the PDSE and its members during and after SMSPDSE1 restart
 - Restart will not cause failure of a user job, TSO session, an edit or browse of a PDSE member, or LISTPDS of a PDSE
- Termination of SMSPDSE1 will cause some ECSA to be lost (possibly up to a few megabytes)

109

© 2003 IBM Corporation

PDSE Commands

- Problem determination
VARY SMS,[PDSE | PDSE1],ANALYSIS
- Repair
VARY SMS,[PDSE | PDSE1],FREELATCH
- Restarting SMSPDSE1
VARY SMS,PDSE1,ACTIVATE
 - Use after SMSPDSE1 has been forced from the system
VARY SMS,PDSE1,RESTART
 - Terminates and then activates a new instance of SMSPDSE1

Refer to System Commands manual for more detail

110

© 2003 IBM Corporation

SMS Initialization Parameters

- Existing parameters enhanced with synonym parameters
 - LRUCYCLES synonym PDSE_LRUCYCLES
 - LRUTIME synonym PDSE_LRUTIME
 - HSP_SIZE synonym PDSE_HSP_SIZE
 - BMFTIME synonym PDSE_BMFTIME
 - MONITOR synonym PDSE_MONITOR
- New parameters
 - PDSE_RESTARTABLE_AS(YES | NO)
 - PDSE1_MONITOR ({YES | NO}[interval[,duration]])
 - PDSE1_LRUCYCLES(nnn | 240)
 - PDSE1_LRUTIME(nnn | 15)
 - PDSE1_HSPSIZE(nnn)
 - PDSE1_BMFTIME(nnn | 3600)

Refer to Initialization and Tuning Reference manual for more detail

111

Migration/Coexistence

- Migration/Coexistence
 - Toleration PTF will be available for lower level systems
 - If one system has PDSE restartable code installed then all systems must have the PTF installed even if SMSPDSE1 is not restarted

112

Multiple Level Security Label in ACS Routines

- An installation can define its own security labels representing the association between a particular security level and a set of zero or more security categories
 - Security label displayed is the minimum security you need to access a data set protected by this profile
 - Seclabel is in user's profile (discreet profile) or from the data set profile (generic profile)
- Some installations wish to segregate data of specific classifications on specific sets of volume
- New ACS read only variable &SECLABEL can now be used to make allocation decisions rather than using an allocation exit

113

© 2003 IBM Corporation

&SECLABEL

- &SECLABEL is set from:
 - User's profile (discrete profile)
 - Data set profile (generic profile)
 - ACEE pointing to a discrete profile if DD SECMODEL or PROTECT=YES parameter is specified in JCL or dynamic allocation
- Set to 'null' if SECLABEL class is not active
- If overflow or extended storage groups are also defined, the installation should ensure that security levels do not conflict
- Users with ADSP attribute are not supported

114

© 2003 IBM Corporation

Additional Support

- ISMF
 - ACS Test application updated to all specification of Seclabel
- NaviQuest
 - Test case generation from ISMF saved lists has new Seclabel field

115

© 2003 IBM Corporation

Migration/Coexistence

- ACS routines using &SECLABEL cannot be translated on lower level systems
- No specific coexistence considerations except that on lower level systems the &SECLABEL will have a 'null' value

116

© 2003 IBM Corporation

Miscellaneous Line Items

- Catalog serviceability
- DFSMSdss SNAPX
- DASD ERP
- CVAF
- VSAM

117

© 2003 IBM Corporation

Catalog Serviceability

- New parameters for F CATALOG command
 - TAKEDUMP
 - Causes the Catalog Address SPACE (CAS) to issue an SVCDUMP using the proper options to ensure all data needed for diagnosis is available
 - RESTART
 - Will prompt the user for additional information
- IEC363D IS THIS RESTART RELATED TO AN EXISTING CATALOG PROBLEM (Y OR N)?**
- If response is 'N' the restart continues
 - If response is 'Y' another prompt is issued
- IEC364D HAS AN SVC DUMP OF THE CATALOG ADDRESS SPACE ALREADY BEEN TAKEN (Y OR N)?**
- If response is 'N' an SVC dump is taken before restart
 - If response is 'Y' the restart continues

118

© 2003 IBM Corporation

DFSMSdss SNAPX Parm

- The DFSMSdss Cross Memory Application Programming Interface (ADRXMAIA) is invoked using the JCL SNAPX=x parameter on the EXEC statement
 - Used for debugging a User Interaction Module (UIM)
- Now can be invoked via user program if SNAPX is specified in the OPTPTR field

DASD ERP, CVAF, and VSAM

- DASD ERP
 - Modules have been combined into a single module and moved above the 16MB line providing VSCR
- CVAF
 - Dumps will now contain all variables improving diagnostics
- VSAM
 - Module split for enhanced serviceability
 - Dynamic LPA capability improving availability

DFSMSDss

121

© 2003 IBM Corporation

DFSMSDss Current Operation

- If a data set exists with the same name as the renamed data set on the target volume or in the standard order of search and is SMS managed the operation will fail
- A data set can also replace a data set that exists with the old name
- RENAME or RENAMEU take precedence over REPLACE
 - If a preallocated target data set exists with the new name chosen by RENAME the operation fails even with REPLACE
- If a source data set does not match rename criteria and a preallocated target data set with the source name exists on the target volume, or in the standard order of search when the source data set is SMS managed, the preallocated data set will be replaced if REPLACE is specified
- Currently users can rename a data set by specifying source data sets and corresponding new names or rename based upon high level qualifier (hlq) during COPY DATASET or RESTORE DATASET
 - RENAMEUnconditional for COPY processing
 - RENAME or RENAMEUnconditional for RESTORE processing

122

© 2003 IBM Corporation

DFSMSdss Replace Unconditional

- New function shipped with APARs OA05249 and OA05874 provides the capability to replace an existing data set with the new name while renaming a data set during a COPY or RESTORE operation
 - Data set COPY
 - Logical data set RESTORE

123

© 2003 IBM Corporation

REPLACEUnconditional

- New function that allows users to be able to rename a data set using the existing RENAME or RENAMEU keywords and replace an existing data set with the new name if that data set exists on the target volume
- REPLACEU is intended to be used in conjunction with RENAMEU and if applicable RENAME but can be specified by itself
 - If specified by itself it will have the same behavior as the REPLACE keyword
 - Will look for a preallocated data set with the source data set name and if found will be replaced
 - If not found a data set with the source data set name will be allocated
- The behavior of DFSMSdss when specifying RENAME(new name) and REPLACEU will be the same as specifying REPLACE (of the old name)

124

© 2003 IBM Corporation

Logical Data Set Copy and RESTORE

- Specifying the REPLACEU with RENAMEU means that even if a data set with the new name exists (usable preallocated target) the data set will be overwritten with the source data set
 - Usable preallocated data sets have the same attributes as the source
 - Unusable new name preallocated data sets will be scratched under the following conditions
 - Any of the following source and target data set attributes do not match:
 - ✓ CI size
 - ✓ IMBED
 - ✓ REPLICATE
 - ✓ Different key length
 - ✓ Different record length
 - ✓ SPANNED
 - The target data set is not large enough to contain all of the source data
 - The data set was not defined as reusable and the high used RBA of a target VSAM KSDS is not 0

125

© 2003 IBM Corporation

Record Level Sharing (RLS) Considerations

- Currently DFSMSdss does not save all of the information in the RLS cell in the VVDS when performing a COPY operation to a preallocated data set
 - RESTORE processing does save some RLS information
- Copy processing will now work like RESTORE processing
 - RLS information of the preallocated target data set will be saved when checking the VVRs of the preallocated target
 - The LOG, LOGSTREAMID, BWO information will be saved
 - The recovery required bit of the source data set will continue to be carried to be carried forward regardless of the preallocated target status
 - After data movement, if the preallocated target data set had RLS information, it will be put back into the RLS cell otherwise the target will have no RLS information
 - If the target was not preallocated the RLS information from the source data set will be propagated to the target

126

© 2003 IBM Corporation

Miscellaneous Considerations

- Exit 22 which allows the use of the DFSMSdss API to control what gets placed in the RLS cell is used by RESTORE processing but will not be invoked for COPY processing
- DFSMSdss when using FlashCopy will issue a DDSW withdraw for the tracks that a preallocated target data set resides on
 - For preallocated targets with the old name when REPLACE is specified as well as preallocated targets with the new name when RENAMEU and REPLACEU are specified

127

© 2003 IBM Corporation

Migration/Coexistence

- No specific migration considerations
- There will not be any toleration PTFs

128

© 2003 IBM Corporation

DFSMSHsm

129

© 2003 IBM Corporation

Secondary Space Management (SSM) Today

- Statistics cleanup
 - Restart tape copy
 - DSR and VSR records that expired and require cleanup from MCDS
- Migration level cleanup
 - Performs several scans starting from where it previously left
 - Deletion of expired migrated data sets
 - ✓ MCD records that are expired or recalled and require cleanup
 - ✓ MCD records that require ML1-to-ML2 data set movement
 - ✓ MCA and MCO records that are expired
 - One scan of the OCDS
 - ✓ TCN records that require TAPECOPY restart
 - ✓ Performed before DSR/VSR scan in MCDS
- Delete records from SDSP
- Moves migration copies from ML1 to ML2
- Expired ML1 and ML2 processing
- Runs as a single task

130

© 2003 IBM Corporation

Migration/Coexistence

- Can coexist with prior releases of DFSMSHsm
 - Toleration maintenance will be made available

DFSMSrmm

DFSMSrmm

- DFSMSrmm Client/Server support
 - DFSMSrmm can now run on a system which does not have direct access to the DASD containing the DFSMSrmm CDS
 - I/O requests to the CDS are handled over the TCP/IP network
 - Allows multiple sysplexes to have a single tape inventory
 - An object oriented interface written for the C/C++ languages is now available for the DFSMSrmm API

135

© 2003 IBM Corporation

DFSMSrmm *(continued)*

- ISPF usability
 - Variable saving and reuse
 - Listing logicals on a stacked volume
 - Line commands supported from any search results list
 - Move LIMIT input field to be on initial panel of scrollable search panels
 - New EJECT user option
 - CLIST option for all SEARCHes
 - Customizing the dialog with new line command 'U'
 - SEARCH results lists row management

136

© 2003 IBM Corporation

Client/Server Support

- Broadens the definition of RMMplex from just DFSMSrmm subsystems sharing a common CDS
 - A set of DFSMSrmm subsystems that share a common CDS either directly or through a SERVER subsystem
 - One or more DFSMSrmm SERVER subsystems
 - ✓ SERVER subsystems have direct access to the CDS
 - One or more CLIENT subsystems
 - ✓ CLIENT subsystems have no direct access to the CDS but share the CDS through the SERVER
 - One or more STANDARD subsystems
 - ✓ STANDARD subsystems have direct access to the CDS
- Allows the sharing of a single tape inventory across multiple sysplexes

137

© 2003 IBM Corporation

Client/Server/Standard Subsystems

- A SERVER subsystem is identified by the OPTION SERVER operand
 - OPTION SERVER also identifies basic TCP/IP info which allows the server to handle requests from CLIENT subsystems
 - Processes its own local requests as well as requests from clients
- A CLIENT subsystem is identified by the OPTION CLIENT operand
 - OPTION CLIENT also identifies basic TCP/IP info which allows the client to send requests to a server
 - Any parmlib options not required for a client are ignored
 - Client can only connect to one server at a time
 - Keep one or more parmlib members for alternate server capability
- A standard subsystem does not have OPTION CLIENT or OPTION SERVER specified and is the default
- TCP/IP V4 required

138

© 2003 IBM Corporation

Processing Requests

■ CLIENT

- Most requests are processed by communicating with the server
- Requests that can be processed by the client are processed locally
- When multiple requests are being processed a FIFO queue is maintained
- Operator command 'Q A' displays the tasks and a summary of the queues

■ SERVER

- Local requests are processed unchanged
- Client requests are accepted and processed synchronously while client waits
 - There is no queue of client requests
- Request queues are only maintained for local requests
- Operator command 'Q A' displays local requests as well as accepted client requests

139

© 2003 IBM Corporation

Client/Server Coexistence

- All client/server systems must be at least at z/OS 1.6
 - Recommended that the server system(s) be upgraded first
 - As long as toleration is installed on all systems in an RMMplex they can successfully support client/server processing.
 - All dates and times are local times and any date and times displayed are exactly as they are stored in the CDS
 - No conversion from server time zone to client time zone is performed

140

© 2003 IBM Corporation

RMM API Command Classes

- New interface allows high-level languages such as C/C++ and Java to issue commands
- Input is and RMM command string
- Output can be either SFIs or XML
 - Output for SFIs is in EBCDIC
 - Values include binary, character, decimal
 - Output for XML is converted to character in unicode format as define in the XML Schema file for the RMM resources
 - Each XML object returned from the getBufferXml method of the RmmCommand class contains only the data and tags to define the data
 - The document rmmxml.xsd is a new file that is shipped and is referenced from each XML object
- The RMM API is called RmmApi
 - Header file containing the class definitions is contained in part EDGXHCLU
- Use the RmmApi class to prepare the environment for using the RmmCommand classes to run DFSMSrmm TSO subcommands via the API

141

© 2003 IBM Corporation

ISPF Usability

- Variable saving and reuse
 - Variables used for data entry are saved and restored in the shared variable pool so that they can be saved and reused across sessions
 - Initial value is REUSE ON
 - To switch the setting you can use the fastpath command REUSE ON/OFF
 - Variables are always saved independent of the setting but if REUSE OFF is set variables will not be retrieved and all values must be entered
- Listing logical volumes on a stacked volume
 - The 'l' line command issued on a stacked volume and will list volumes
 - Data sets will not be listed along with the volumes
- Existing and new line commands supported from any search results list
 - VL – List volume chain from the data set search results
 - IL – List data set chain from the volume search results
 - CA – Confirm all actions and moves outside of action summary lists
 - CH – Change from the action summary list
 - CS - Confirm scratch supporting new manual scratch release action for volumes

142

© 2003 IBM Corporation

ISPF Usability (continued)

- Move LIMIT input field to be on initial part of scrollable search panels
 - Moved to the first viewed part of the scrollable area on panels
 - LIMIT value now set to '*' if the LIMIT field is null or blank
- New EJECT User option
 - Now can specify and change the default EJECT station to CONVENIENCE or BULK
 - Specified in Dialog User Options panel or fastpath command EJECT BULK/CONVENIENCE with the initial value set to EJECT CONVENIENCE
- CLIST option for all SEARCHes
 - New field on the first scrollable screen for all search panels prompts for CLIST
 - If YES is specified a new panel will be displayed prompting for relevant values

143

© 2003 IBM Corporation

ISPF Usability (continued)

- Customizing the dialog with new line command 'U'
 - Enables a user to provide locally provided line command support and calls exec EDGRLCL
 - This exec can be used to implement any local extensions to the search results line commands
- Search results lists row management
 - Capability to delete rows for records that are deleted
 - Lists effected:
 - Data set
 - Product
 - Volume (but FORCE and REMOVE only)
 - VRS

144

© 2003 IBM Corporation

z/OS 1.6 Summary

| Better Utilization of Resources | Ease of Use | Improved Security | Continuous Availability | Enhanced Performance |
|--|---|--|--|--|
| <ul style="list-style-type: none"> ✓ DFSMSrmm Provides Tape Management Support for multiple Systems ✓ Volume Selection Based on PAVs | <ul style="list-style-type: none"> ✓ DFSMSrmm ISPF usability ✓ DFSMSdss REPLACE UNCONDITIONAL | <ul style="list-style-type: none"> ✓ MLS SECLABEL in ACS Routines | <ul style="list-style-type: none"> ✓ PDSE support for multiple restartable address spaces | <ul style="list-style-type: none"> ✓ DFSMShsm Secondary Space Management Multitasking |

145

© 2003 IBM Corporation

References

- URLs
 - <http://www.storage.ibm.com/software/index.htm>
 - <http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/storsw>
- Publications
 - SC26-7329 DFSMS Migration
 - SC26-7339 DFSMS: Using Data Sets
 - SC26-7331 DFSMSdftp Storage Administration Reference
 - SC26-7326 DFSMS Access Method Services for Catalogs
 - SC26-7352 DFSMS Installation Exits
 - SC26-7337 DFSMS Macro Instructions for Data Sets
- Redbooks
 - SG24-6105 VSAM Demystified

146

© 2003 IBM Corporation