



This presentation provides an introduction to the DFSMShsm™ fast replication function.

Agenda

- Introduction to fast replication
 - ▶ FlashCopy® overview
 - ▶ SMS constructs
- Creating a backup copy
- Recovering a backup copy
- Recommendations

This presentation begins with an overview of FlashCopy and the SMS constructs used by the function. It then describes the various features of DFSMSHsm fast replication and how it can be used to create and recover point-in-time backup copies. In closing, a few recommendations will be given.

Introduction

- DFSMSHsm invokes volume-level *fast replication* to create backup versions for sets of volumes
 - ▶ Backup hundreds of volumes in a matter of minutes
 - ▶ DFSMSHsm manages copying the disk copies to tape
 - ▶ Developed jointly with DB2®
 - BACKUP SYSTEM utility
- Generic term of “*fast replication*”
 - ▶ Multiple Point-in-Time copy functions supported by DFSMSdss
 - FlashCopy (Any vendor disk that supports FlashCopy interface)
 - SnapShot
- Recovery can be performed for the entire volume set, individual volumes or individual data sets

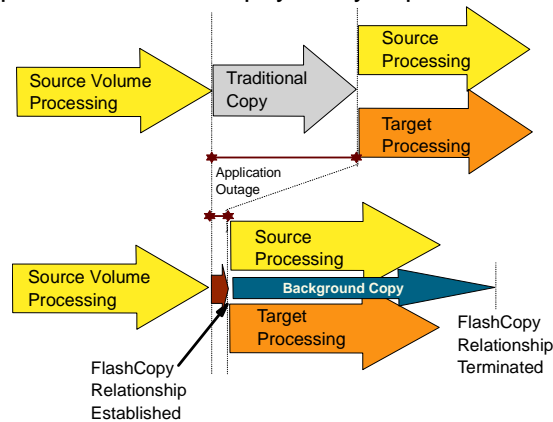
DFSMSHsm invokes volume-level fast replication to create point-in-time backup versions for sets of volumes. This enables DFSMSHsm to backup hundreds of volumes in a matter of minutes. DFSMSHsm can also create tape copies from the disk copies that are generated. DFSMSHsm manages both the disk and tape copies. This solution was developed jointly with DB2 and is an integral part of the DB2 BACKUP SYSTEM utility.

The generic term ‘fast replication’ is used to indicate that DFSMSHsm supports any point-in-time backup copy technique supported by DFSMSdss™. Today, DFSMSHsm supports FlashCopy and SnapShot. This means that this function can be used on any vendor disk that supports the FlashCopy interface.

Besides being able to recover the entire set of volumes that was backed up, DFSMSHsm can also recover individual volumes and even individual data sets.

FlashCopy

- FlashCopy is a two-phase Backup
 - ▶ Initialization phase: Establishes source-to-target relationship and immediately returns control
 - ▶ Copy phase: Tracks are physically copied in the background



4

DFSMShsm fast replication

© 2008 IBM Corporation

FlashCopy is a two-phase backup. The first phase is the initialization phase. During the initialization phase, the storage controller establishes a source-to-target relationship between the source volume and the volume that is the target of the copy. After the initialization completes, control is returned to the initiator. Updates to the source volume can begin immediately.

The second phase is the copy phase. In this phase, the storage controller performs a background copy of the tracks from the source volume to the target volume. If a track on the source volume is updated before that track has been copied to the target, then the storage controller will immediately copy that track to the target volume before actually writing the update to the source.

Supported FlashCopy functions

- **Full copy** (V1R5)
 - ▶ Full background copy is performed
- **No copy** (V1R8)
 - ▶ Only *changed* tracks on source are copied to target volume
 - ▶ Target volume is used as a cache to create a dump copy
 - ▶ DFSMSHsm does not support Flashback
- **Incremental copy** (V1R8)
 - ▶ Each Incremental backup copy only copies the tracks that changed on the source volume since the last incremental copy was created

DFSMSHsm supports three FlashCopy options: Full Copy, No Copy and Incremental Copy.

Full copy, which is the default, copies every track of the source volume to the target volume. When the copy is complete, the target volume will be identical to the source volume, except for the volser. The volser remains unchanged so that the target volume can remain online. This option requires a minimum of one dedicated target volume for each source volume.

When the No Copy option is used, there is no background copy performed by the storage controller. Only tracks that are updated on the source are copied to the target volume. The No Copy option enables the target volumes to serve as a cache from which a tape copy of the point-in-time disk copy can be created. After the tape copy has been created, the relationships between the source and target volumes is withdrawn. DFSMSHsm does not support disk recovery when the No Copy option is used. This option does not require a dedicated target volume for each source volume. Target volumes can be used by multiple source volumes. This option is available in z/OS DFSMSHsm V1R8 and higher.

The Incremental option reduces the amount of time required to create the background copy. With the Incremental option, the storage controller keeps track of which source tracks were updated since the last incremental copy. When a new copy is created, only the tracks that changed are copied in the background. This option requires one dedicated target volume for each source volume. This option is available in z/OS DFSMSHsm V1R8 and higher.

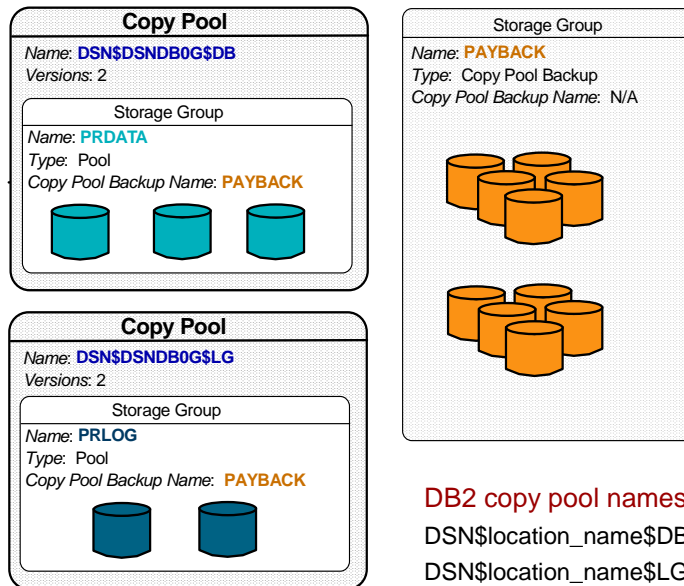
SMS enhancements

- “Copy pool” SMS construct
 - ▶ Defines which *storage groups* should be processed collectively for point-in-time functions
 - ▶ Contains from 1 to 256 storage groups
 - Backup processed at storage group level because it is a track-level function
 - ★ Storage group volumes dynamically retrieved at time of each backup
- “Copy pool backup” storage group type
 - ▶ Defines which volumes DFSMSHsm may select as target volumes for point-in-time backup versions
 - ▶ Defines candidate target volumes for DFSMSHsm fast replication
 - ▶ Cannot be assigned for allocation by SMS ACS selection routines
 - ★ Protects data from being overwritten

The SMS ‘copy pool’ construct defines which storage groups should be processed collectively for point-in-time functions. A copy pool may contain from 1 to 256 storage groups. Storage groups may be contained within multiple copy pools. When DFSMSHsm creates a fast replication backup copy, it is at the copy pool level. That means that every volume in every storage group defined to the copy pool will be backed up at the same time. The backup is processed at the storage group level since the fast replication function is a track-level operation. DFSMSHsm dynamically retrieves the storage groups in the copy pool and each volume in those storage groups at the time of the backup. This ensures that volumes that have been added to a storage group since the last backup will be included in the new backup copy.

Instead of requiring you to create source/target pairings, DFSMSHsm automatically selects target volumes for each source. The SMS ‘copy pool backup’ storage group type defines which volumes DFSMSHsm may select as target volumes for point-in-time backup versions. SMS prevents new data set allocations to volumes within copy pool backup storage groups. This protects data from being overwritten during the fast replication operation.

SMS enhancements (continued)



This chart contains a picture of the copy pool environment. It shows two copy pools and a single copy pool backup storage group. The first copy pool is named 'DSN\$DSNDB0G\$DB' and has a single storage group, PRDATA, assigned to it. The second copy pool is named 'DSN\$DSNDB0G\$LG' and has a single storage group, PRLOG, assigned to it. This example uses sample names as if the copy pools were a part of a DB2 environment. Each copy pool is configured to keep two backup versions on disk. Storage groups PRDATA and PRLOG indicate that DFSMSshm should select the target volumes from copy pool backup storage group PAYBACK. Copy pool backup storage group PAYBACK has ten target volumes. That is enough target volumes for two backup versions for both copy pools.

Storage groups PRDATA and PRLOG are not required to specify the same copy pool backup storage group.

Creating a fast replication backup

- DFSMSHsm command - **FRBACKUP**

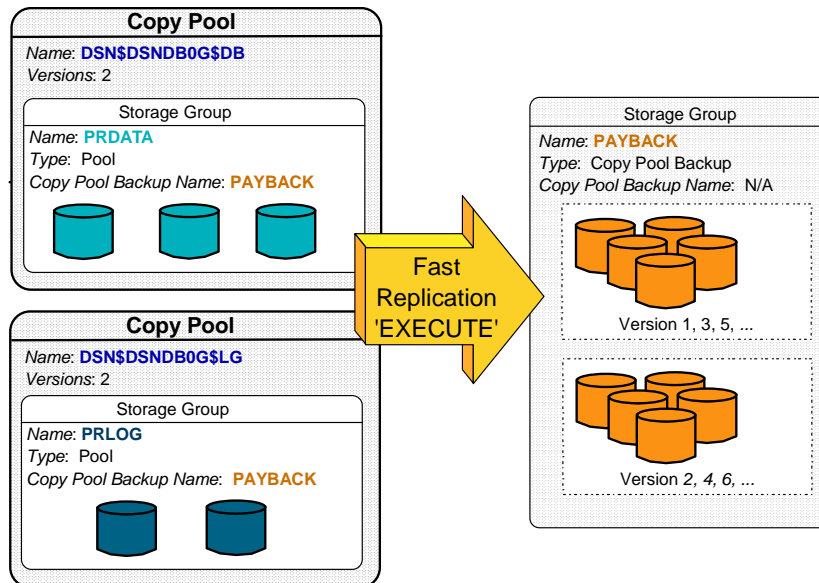
FRBACKUP COPYPOOL(cpname)

- Command is considered successful only after a fast replication relationship has been established for every source volume
- If one or more volumes fail:
 - ▶ Version is marked as a failure
 - ▶ Failed version will be the target of the next FRBACKUP command
 - ★ Version number does NOT increment until successfully created

A fast replication backup copy is created using the DFSMSHsm FRBACKUP command. A required parameter on the FRBACKUP command is the name of a copy pool. When the command is issued, DFSMSHsm retrieves the specified copy pool definition from SMS and creates a fast replication backup copy for every volume in every storage group defined to the copy pool. The command is only considered successful if a fast replication relationship is successfully established for every volume.

If one or more volumes fail, the backup version is marked as a failure and all of the relationships that were successfully established are withdrawn. DFSMSHsm will target the same set target volumes when the FRBACKUP command is reissued, except for those volumes that failed. The version number is not incremented until a new backup version is successfully created.

Creating a fast replication backup (continued)



The picture on this chart shows how DFSMSHsm 'round-robins' through the target volumes when more than one disk version is being maintained. In this example, DFSMSHsm selects one set of target volumes for the first version. DFSMSHsm then selects another set of target volumes for the second version. Since two versions were specified, DFSMSHsm will overlay the first version when it creates the third version, and so forth. Maintaining two or more disk versions guarantees that at least one disk version will always be valid and available for recovery.

Creating a fast replication backup (*continued*)

- **Serialization**

- ▶ Individual volume VTOCs are serialized as they are flashed
- ▶ Responsibility of *caller* to protect source volume against data set level updates until FlashCopy relationship established
- ▶ If all volume VTOCs need to be serialized concurrently, caller must do so
 - **NOVTOCENQ** keyword indicates that VTOCs have already been serialized

When creating a fast replication backup version, DFSMSHsm serializes each volume VTOC as it processes it. DFSMSHsm does *not* serialize the entire set of volumes at one time. Nor does DFSMSHsm prevent data set level updates during the fast replication backup. It is the responsibility of the invoker of the fast replication function to quiesce the source data if that is needed. If the invoker has done so, then the NOVTOCENQ optional keyword should be specified on the FRBACKUP command.

Fast replication backup tape support

- DFSMSHsm manages copying FlashCopy target volumes to tape
 - ▶ By command, immediately after FlashCopy version is created
 - ▶ By command, some time after the FlashCopy version is created
 - ▶ During Automatic Dump window
- Tape copies are 'Dump' copies
 - ▶ DFSMSDss Full-volume physical dump
 - ▶ Existing DFSMSHsm Dump function used to create copies
 - Dump Classes used to define policies
- Maintain up to 85 versions
 - ▶ Up to five copies per version
- Support available on z/OS V1R8


Besides managing fast replication disk copies, DFSMSHsm creates and manages fast replication copies on tape. DFSMSHsm can create these tape copies immediately after the disk copy is created, some time after the disk copy is created, or automatically during a predefined Automatic Dump window.

The tape copies are created using the DFSMSDss full-volume physical dump function in the existing DFSMSHsm dump function. The DFSMSHsm Dump Class construct is used to define the policies used to create these tape copies. Up to 85 tape versions can be created, with up to five copies for each version.


The DFSMSHsm tape support for fast replication was introduced in z/OS V1R8.

Example

Step 1:
Version 1 is created with FlashCopy to disk followed by Dumping target volumes to tape.

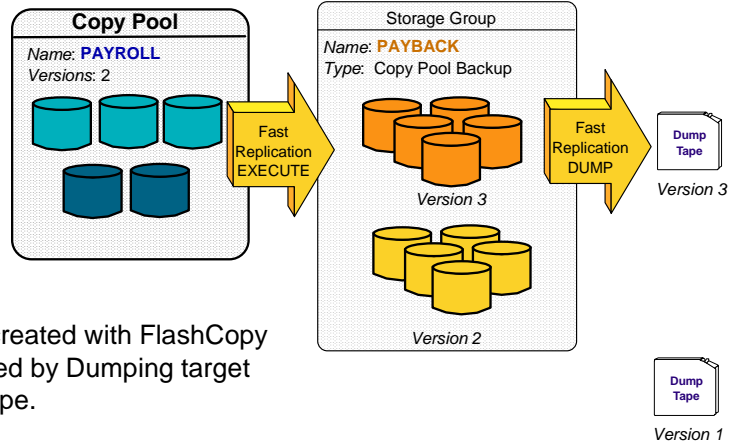


Step 2:
Version 2 is created with FlashCopy to disk only.



The next few slides have a pictorial example of how the function works. In the first step, the FRBACKUP COPYPOOL(payroll) DUMP command is used to create a point-in-time copy of all the volumes in copy pool PAYROLL. Since the optional DUMP keyword was specified, after the disk copy is successfully established, DFSMSHsm immediately creates a tape copy from the target volumes. In the second step, another point-in-time disk version is created. After step two, there are two backup versions. The first version has both a disk and tape copy and the second version has only a disk copy.

Example (continued)

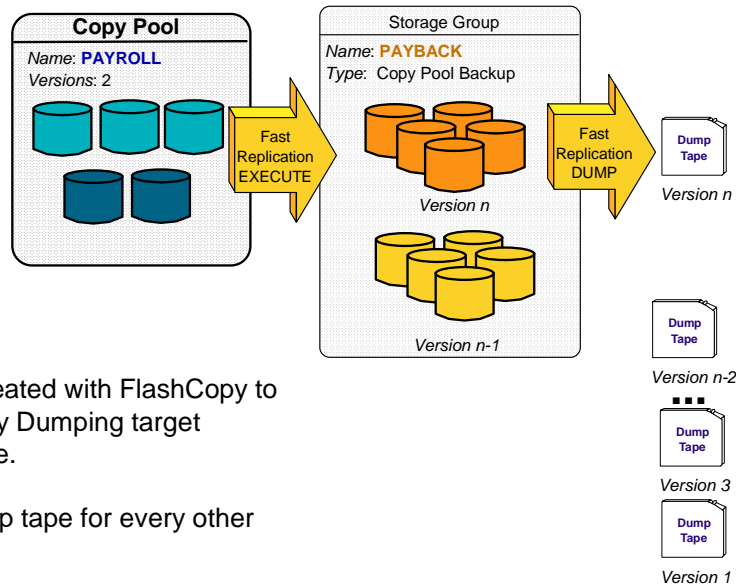


Step 3:
Version 3 is created with FlashCopy
to disk followed by Dumping target
volumes to tape.

Notice that Version 1 disk copy was
overlaid to create Version 3

For the third step, both a disk copy and tape copy are created. The creation of the third disk version overlays the disk copy for version 1. After this step, there are three versions. Version 1 is now only on tape, version 2 is only on disk, and version 3 is on both disk and tape.

Example (continued)



Step *n*:
Version *n* is created with FlashCopy to disk followed by Dumping target volumes to tape.

There is a dump tape for every other version

As this process continues, the two most recent backup versions will be on disk, with one of those versions also having a tape copy. Additionally, there will be a tape copy for every other version that has rolled-off of disk.

Withdrawing a backup version

! NEVER Withdraw FlashCopy relationships outside of DFSMSHsm control

- ▶ Will cause backup version to be invalidated, but DFSMSHsm will process it as valid
- ▶ Could result in data integrity exposure
- Use the **FRBACKUP COPYPOOL(cpname) WITHDRAW** command
 - ▶ Will withdraw all current FlashCopy relationships
 - ▶ Applies to most recent VALID version of the specified copy pool
 - ▶ The version state will be changed to FAILED
 - Only if relationships still exist
 - ▶ Any in-progress dumps are failed and no further dumps are created
- Can only be performed at the copy pool level

A FlashCopy relationship established through the FRBACKUP command must never be withdrawn outside of DFSMSHsm. Doing so will invalidate the backup version, but DFSMSHsm will still process it as valid. This could result in a data integrity exposure. If a FlashCopy relationship needs to be withdrawn, then use the FRBACKUP COPYPOOL(cpname) WITHDRAW command. This command will withdraw all the active FlashCopy relationships for the most recent VALID version for the specified copy pool. If any relationships are withdrawn, then that version will be changed to FAILED. Any in-progress dumps of that version will be failed and no further dump can be initiated for that version. The WITHDRAW command can only be performed at the copy pool level.

Recovering a fast replication backup

- Target volumes cannot be used as source volumes
 - ▶ Dump conditioned volumes
- DFSMSHsm **FRRECOV** command must be used to recover target disk volume back to the source
 - ★ Data is immediately available after the FlashCopy initialization completes
- ! **FRRECOV cannot be performed until all background copies have completed**
 - **QUERY COPYPOOL(cpname)** will report back which volumes still have an active background copy
 - ★ Advantage for FlashCopy Incremental – only changed tracks are copied in the background so recovery can be performed sooner

Fast replication target volumes cannot be used directly as source volumes since they are dump conditioned volumes. A dump condition volume is a volume that looks identical to the source volume except that its volser is different and the volser that it is a copy of is contained within Track 0. The DFSMSHsm FRRECOV command must be used to recover a target disk volume back to the source. When fast replication recovery is performed, the data on the source volume is immediately available after the FlashCopy initialization completes.

The FRRECOV command cannot be performed for a copy pool until all background copies for the source volume have completed. The FRRECOV command will fail if any of the source volumes are in an active FlashCopy relationship. Use the QUERY COPYPOOL(cpname) command to determine which volumes have an active background copy. One of the advantages of the FlashCopy Incremental function is that the background copy completes sooner since it is only copying changed tracks. This enables a recovery to be done sooner than if a full background copy was being performed.

Recovering a fast replication backup

- **FRRECOV COPYPOOL(cpname) VERIFY(Y)**
 - ▶ Recovers all of the volumes from the named copy pool
 - ▶ Only supports recovery from *disk*
 - Recovery from dump tape is single threaded
- **FRRECOV TOVOLUME(volser)**
 - ▶ Recovers an individual volume
 - ▶ Use when one or more volumes fail during copy pool recovery
 - ▶ Prefers disk recovery if both disk and tape are available
- **FRRECOV DSNAME(dsname1, dsname2, ...) REPLACE**
 - ▶ Recovers one or more data sets
 - ▶ DSNAME is only available in V1R8 or higher
 - ▶ Default is recovery from disk, if both disk and tape are available

DFSMSHsm can recover an entire copy pool, an individual volume or individual data sets. The FRRECOV COPYPOOL(cpname) VERIFY(Y) command will recover each volume that was backed up as a part of the copy pool. Recovery performed at the copy pool level is only supported for recovery from disk. This is because DFSMSHsm dump tape recovery is only single threaded.

To recover a single volume, use the FRRECOV TOVOLUME(volser) command. This command can be used to retry individual volumes that failed during FRRECOV COPYPOOL(cpname) processing. The recovery can be performed from disk or tape. If both disk and tape are available, then DFSMSHsm will recover from disk.

Use the FRRECOV DSNAME(dsname) REPLACE command to recover one or more data sets. Wildcards may be specified to recover multiple data sets with a single command. The recovery can be performed from disk or tape. If both disk and tape are available, then DFSMSHsm will recover from disk. If enabled, the data set recovery from disk will use fast replication. This capability was introduced in z/OS V1R8.

Recovery limitations

! Data Set Recovery

- ▶ At the time of the recovery, data sets *must be* cataloged on the same volumes that they resided on at the time of the backup
- ▶ No support for:
 - Deleted data sets
 - Data sets that have moved since the time of the backup
- ★ Does support data sets that expanded to new volumes

! Recovery From Tape

- ▶ DFSMShsm has a single task for recovery from a dump tape
- ▶ Use native DFSMSdss commands

! DO NOT Migrate data within copy pool

- ▶ Fast Replication is a track-level operation
- ▶ Migrating / Recalling data within a copy pool will cause data loss during recovery

There are several limitations associated with recovery that should be noted.

First, in order to recover an individual data set, the data set must be cataloged on the same volumes that it resided on at the time of the backup. There is no support for data sets that have been deleted or moved since they time that they were backed up. DFSMShsm does support the recovery of data sets that expanded to new volumes since the time of the backup.

Secondly, volume-level recovery from tape is single threaded, so only a single volume can be recovered at a time. To bypass this restriction, native DFSMSdss commands can be submitted to process volume-level tape recoveries in parallel. The output from the DFSMShsm LIST COPYPOOL(cpanme) command can be used to create the DFSMSdss commands.

Finally, data within a copy pool should never be migrated with the DFSMShsm migration function. This is because fast replication is a track-level operation that does not work in conjunction with the DFSMShsm migration inventory. Migrating and Recalling data within a copy pool will cause data loss during recovery.

Recommendations

- **Define catalog within copy pool**
 - ▶ Otherwise, recovery will cause catalog and data to be out-of-synch
 - ▶ Unallocate all catalogs in the copy pool before recovery
- **Patch to determine why a target volume is not being selected**
 - ▶ Warning – Lots of messages!
 - ▶ PATCH .FRGCB.+9 BITS(.1.....)
- **BCDS record size *must* be 6544**
 - ▶ In V1R9, FRBACKUP command will fail if BCDS < 6544

In closing, a few recommendations:

First, the catalog or catalogs that contain data sets that are within a copy pool should themselves be defined on a volume within the same copy pool. Also, those catalogs should only contain data sets that are within the copy pool. They should not contain data sets that do not reside within the copy pool. If these rules are not followed, recovery will cause the catalogs and data to be out-of-synch. Also, all catalogs within a copy pool should be unallocated before recovering the copy pool. Otherwise, the volumes on which the catalogs reside will fail due to a serialization error.

If you are having difficulty setting up your environment such that DFSMSHsm cannot pair a target volume to each source volume, use the 'PATCH .FRGCB.+9 BITS(.1.....)' command. This patch indicates that DFSMSHsm should issue a message for every target volume that it determines is ineligible for a source volume. It should be noted that this patch will generate a lot of messages.

Finally, when using the DFSMSHsm fast replication function, the maximum BCDS record size must be defined as 6544. In V1R9, the FRBACKUP command will fail immediately if the record size is less than 6544. On V1R8 and below, the FRBACKUP may fail with a CDS I/O error.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_IEA-DFSMSHsm-FastRepl.ppt

This module is also available in PDF format at: [../IEA-DFSMSHsm-FastRepl.pdf](http://IEA-DFSMSHsm-FastRepl.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2 DFSMSdss DFSMSHsm FlashCopy IBM z/OS

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.