


DFSMS Replication to Meet Your Disaster Recovery Needs

Session # 10955

Jeff Suarez

IBM Corporation

jrsuarez@us.ibm.com



Legal Disclaimer

NOTICES AND DISCLAIMERS

Copyright © 2012 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

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. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

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. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

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 NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

Legal Disclaimer

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily 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. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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.

Trademarks



The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BookManager*	Enterprise Storage Server*	IP PrintWay	RMF
CICS*	ES/9000*	Language Environment*	S/370
DB2*	FlashCopy*	Lotus*	S/390*
DB2 Universal Database	GDPS*	Multiprise*	Tivoli*
developerWorks*	HiperSockets	MVS	TotalStorage*
DFSMSdfp	IBM*	Notes*	WebSphere*
DFSMSdss	IBM eServer	OS/390*	z/Architecture
DFSMShsm	IBM e(logo)server*	Parallel Sysplex*	z/OS*
DFSMSrmm	IBM logo*	RACF*	zSeries*
DFSORT	IMS	RAMAC*	
Domino	InfoPrint*		

Intel is a trademark of the Intel Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

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

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

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

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Agenda

- **Business Continuity Overview**
- **Copy Services Functions**
 - **FlashCopy**
 - **Metro Mirror**
 - **Global Mirror**
 - **z/OS Global Mirror**
- **DFSMS Interfaces**
- **Recent Copy Services Enhancements**
 - **Space Efficient FlashCopy**
 - **IBM Remote Pair FlashCopy**
- **APARs of Interest**

Business Continuity

Business Continuity is not simply IT Disaster Recovery... it is a management process that relies on each component in the business chain to sustain operations at all times.

- **Effective Business Continuity depends on ability to:**
 - **Reduce the risk of a business interruption**
 - **Stay in business when an interruption occurs**
 - **Respond to customers**
 - **Maintain public confidence**
 - **Comply with requirements:**
 - **Audit**
 - **Regulator/Legislative**
 - **Insurance**
 - **Health and Safety**



People

Facilities

Business
Processes

Infrastructure

Applications


... An end-to-end Business Continuity program is only as strong as its weakest link

IBM Copy Services Terminologies

FlashCopy

Point in time copy



Within the same Storage System



Metro Mirror

Synchronous mirroring


Primary Site A
Metro distance <300km
Site B



Global Mirror

Asynchronous mirroring



Primary Site A
Out of Region Site B





Metro / Global Mirror

Three site synchronous & asynchronous mirroring


Primary Site A
Metro Site B
Out of Region Site C



IBM Copy Services Terminology's




FlashCopy overview



PiT copy technology on the disk subsystem
 When a FlashCopy is issued the copy is available immediately
 A bitmap tracks the relationship between source and target tracks

Read and write activity are possible on both the source and target devices




Writes to the source may cause a copy on write if the track has not been copied to the target

Reads of tracks on the target that have not been copied from the source will be redirected to the source

Several options available for FlashCopy including

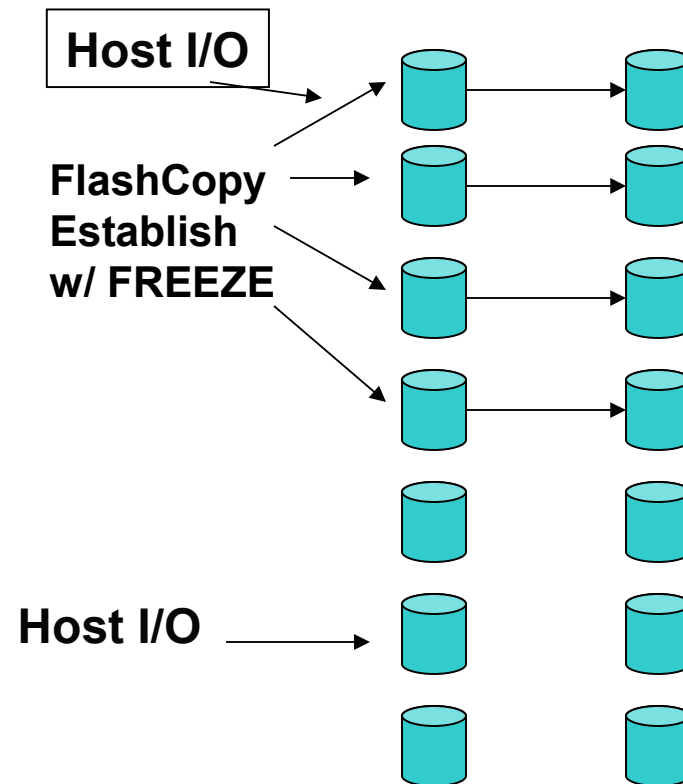
- Incremental FlashCopy
- Consistent FlashCopy
- Multiple FlashCopy relationships
- Dataset level FlashCopy
- Space Efficient FlashCopy
- Remote Pair FlashCopy

An optional background copy process will copy all tracks from the source to the target which will end the relationship

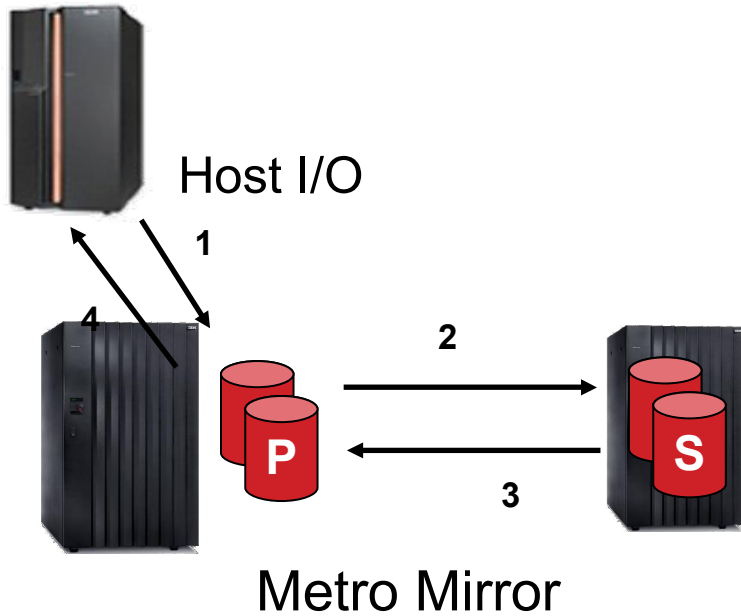


FlashCopy Consistency Groups

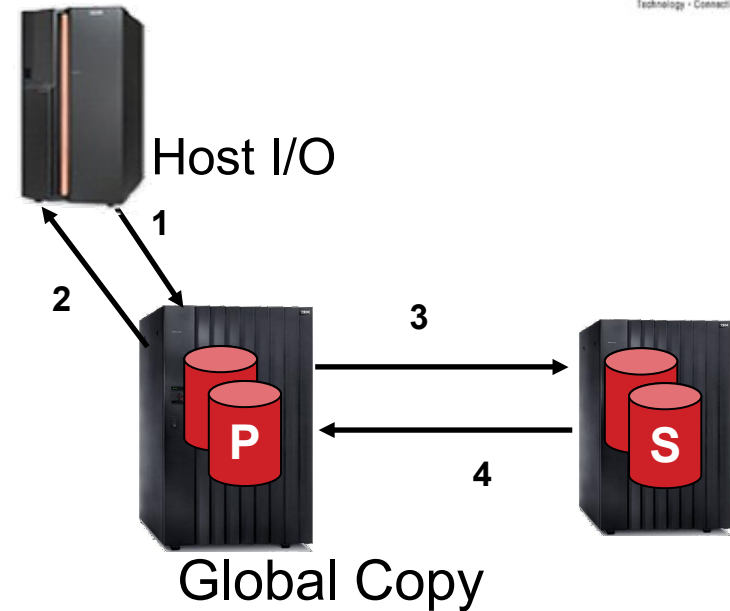
- Data consistent copies across boundaries
- Reduce or eliminate need to quiesce applications
- Long busy on establish with FREEZE
- LSS thawed with ACTION=THAW or 2 minute timer expired
- Application must be aware of data dependency



Metro Mirror and Global Copy



- Write to primary volume.
- The primary disk subsystem initiates an I/O to the secondary disk subsystem to transfer the data.
- Secondary indicates to the primary that the write is complete.
- Primary acknowledges to the application system that the write is complete.




- Write to primary volume.
 - Primary acknowledges to the application system that the write is complete.
- At some later time:
- The primary disk subsystem initiates an I/O to the secondary disk subsystem to transfer the data.
 - Secondary indicates to the primary that the write is complete.
 - Primary resets indication of modified track.

Global Mirror architecture

- Bitmaps used to hold consistency group information on primary disk subsystem
- PPRC paths used for disk subsystem to co-ordinate consistency groups among multiple disk subsystem
- Global Copy used to transmit consistency group between primary and secondary
- FlashCopy used to save consistency groups on secondary disk subsystem


Consistency group
Co-ordination and formation



Multiple Global Mirror


- Allows for finer granularity of management and recovery
 - Separate workloads with different RPO requirements
 - Ability to failover only failing servers or applications
- Can now have one or more test session along with one or more production sessions on the same control unit(s)
- Builds upon the existing Global Mirror technology and microcode
- 32 active sessions per control unit
- Works with 3-site Metro Global Mirror (MGM)
- One session per LSS
 - This must be taken into consideration when planning how to divide up volumes into separate Global Mirror sessions
- Multiple session supported on DS8700 only
 - DS8300/DS8100 can only support a single active session
- Intermix of DS8700 and DS8300/DS8100 supported
 - DS8300/DS8100 can be a subordinate for a single or target for many
 - If used as a target, multiple sessions cannot be run bi-directionally
- OA30618 (SDM)

Examples of Intermix



DS8300/DS8100 as a GM target

- Multiple sessions can active but only in one direction (DS8700 to DS8300)




DS8300/DS8100 as a GM subordinate

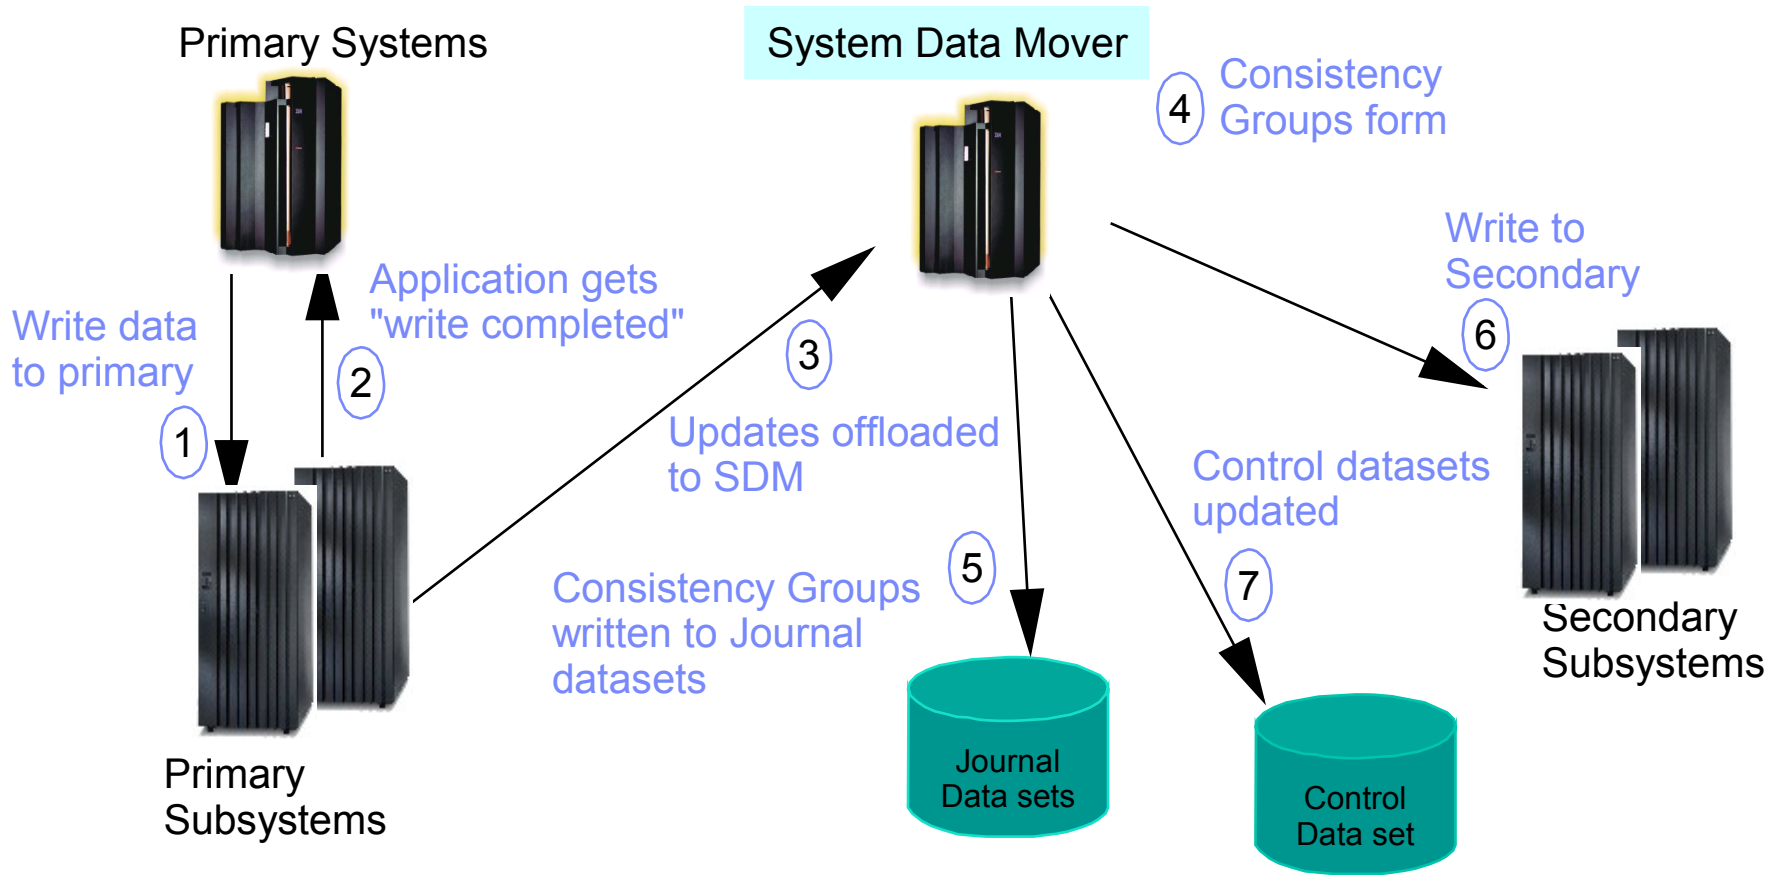
- The DS8700 can run multiple sessions but can only use the DS8300 as a subordinate for one session

Metro Global Mirror

- Combines synchronous and asynchronous PPRC
 - Allows for local continuous availability and out of region disaster recovery
- RPO of 0 for “A” site failure
 - Can recover to either B or D
 - Zero RPO implies automation to ensure no production updates if mirroring stops
- Potential RPO of seconds for “A” and “B” twin site failure
 - Depends on workload and bandwidth between B and C



z/OS Global Mirror (XRC) Concepts



XRC enabled for zIIP

- XRC processing eligible to be dispatched on zIIP
 - Approximately 97% of XRC processing is eligible (under lab conditions)
 - Actual amount of offload depends on amount of CPU consumed, number of ordinary CPs, zIIPs, and CEC level limitations on CP/zIIP ratio
- Enabled via the zIIPENABLE(YES) setting in XRC PARMLIB
- Available with:
 - z/OS V1.10 and above
 - IBM System Storage DS8000, or any storage controller supporting DFSMS SDM

DFSMS Copy Services Interfaces

PPRC Manager and FlashCopy Manager Introduction

- PPRC Manager FlashCopy Manager are two IBM software products that are designed to allow the user to exploit the underlying IBM technology without requiring a understanding of the details involved in using the technology.
 - z/OS based
 - ISPF based for configuration processes
 - Operates on full volumes
- There are two parts to the use of PPRC or and Flash Copy technology
 - Defining the Configuration and building JOBS – ISPF panel based functions
 - Manipulating the configuration – Batch JOB based functions
- Batch JOB based manipulations allow the user to execute the operations individually or to imbed the JOBS into a more complex business application
- Scope of applications – single DS8100 thru very large configurations
- Use technologies that systems programmers and storage administrators use daily (ISPF, ISMF etc.).
- Use a Batch Job based methodology to allow the manipulation tasks to be easily integrated into a more complex business solution job stream.
- Minimize the level of understanding of the details of the underlying PPRC or FlashCopy process for both configuration development and manipulation of the environment.
- For FlashCopy, eliminate SYSGEN dependencies & duplicate label problems with IPLs.

PPRC Manager and FlashCopy Manager Prerequisites



- z/OS 1.7(and up) with ISPF
- REXX library or the REXX alternate Run Time library
 - The REXX Alternate Run Time Library is included with z/OS 1.9
- Appropriate microcode licenses for the IBM Storage subsystem hardware
 - IBM 1750, IBM 2105-F20, IBM 2105-800 or IBM 2107
- Both products are priced
 - One time charge (site license)
 - Subscription and Support option is available
- Can be ordered via Shop zSeries
 - Catalog reference information
 - Package type: z/OS –CBPDO or ServerPac
 - Group: MVS –Miscellaneous/Other
 - PPRC Manager - 5635-PPM
 - FlashCopy Manager – 5635-FCN

ICKDSF FLASHCPY and PPRCOPY

- ICKDSF FLASHCPY
 - First implemented to be only used within the scope of Async PPRC, but restriction has since been lifted.
 - May only be invoked via batch JCL
 - Only supports full volume
 - Provide volume serialization, if online
 - Users are able to establish, query, and withdraw relations
 - Issues withdraw during INIT
- ICKDSF PPRCOPY
 - Supports all functions to setup and run PPRC, for both SYNC and ASYNC (ie Metro Mirror, Global Mirror)
 - Provides volume serialization, if online
- Supports multiple Operating Systems
 - MVS, CMS (CKD only), VSE/ESA (CKD only), Stand-Alone

System Data Mover

- DFSMS SDM provides primitives that are used by IBM products and can be used by ISV programs (or built into CLISTs or REXX execs)
- ANTMAIN/ANTAS000 address spaces must be active
- SDM does not provide data management services (allocation, catalog, serialization)
 - Datasets to be accessed from target volume need user intervention or use of interface like DFSMSdss
- TSO commands
 - All FlashCopy and PPRC functions supported
- Application Programming Interface
 - ANTRQST macro
- REXX Interface
 - Formally Supported in z/OS V1R12

System Data Mover

- **Point-in-time Products**

- FlashCopy
- Concurrent Copy
- Virtual Concurrent Copy (IBM) / Snapshot (OEM)

- **Continuously Mirrored Products**


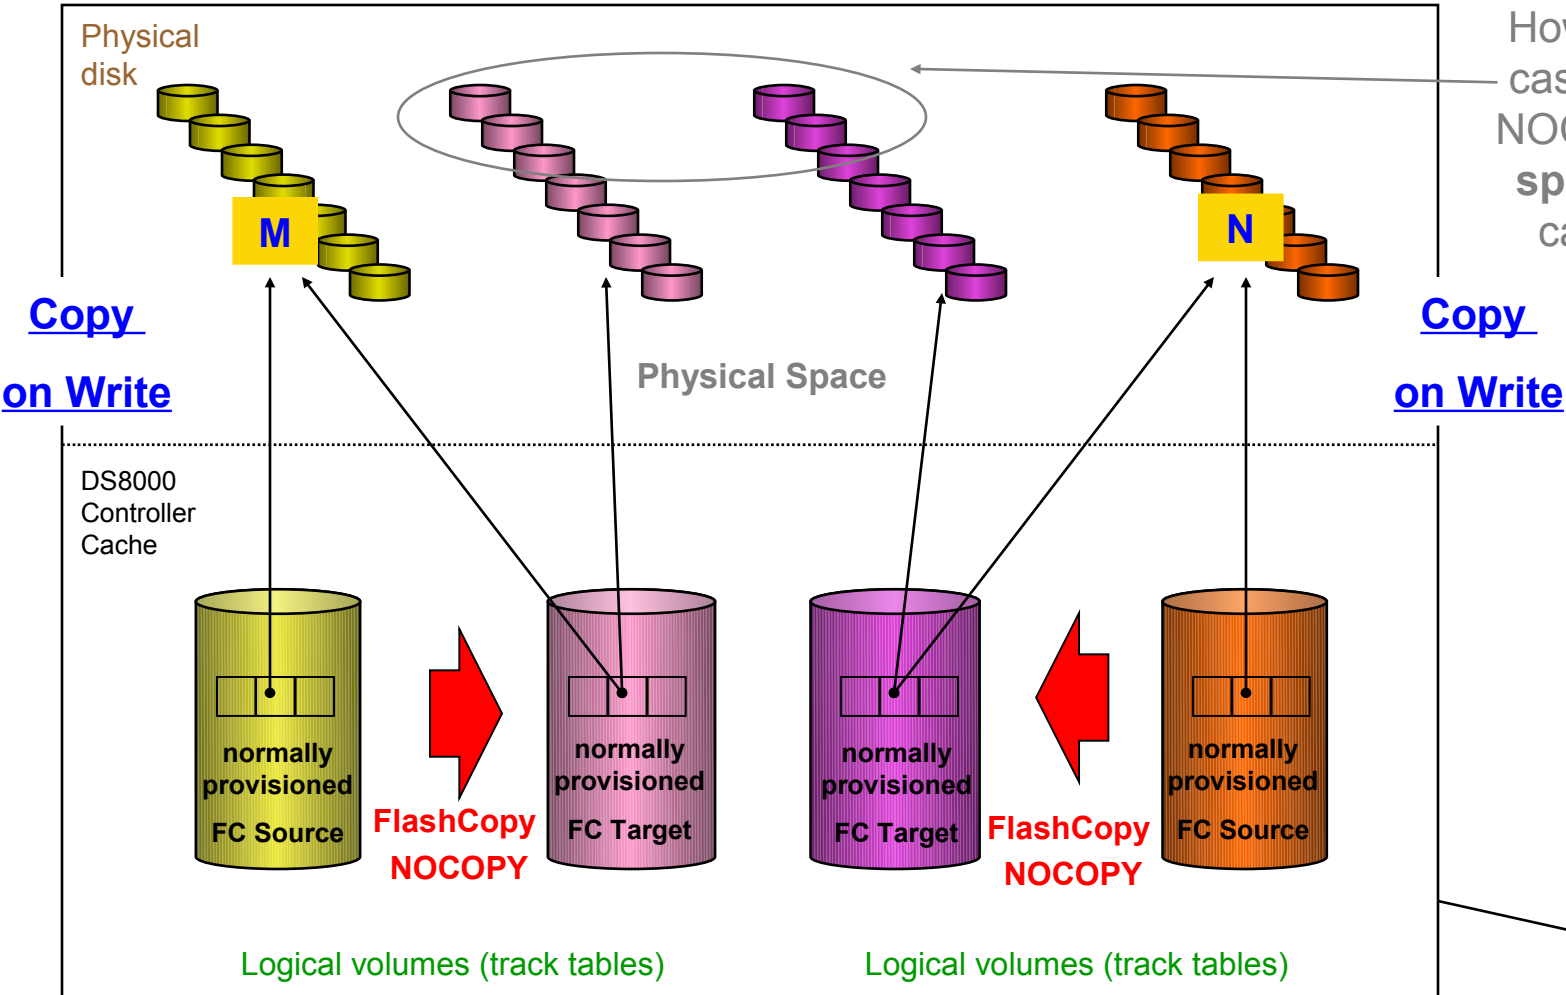
- PPRC
 - Metro Mirror for ESS, aka (synchronous) Peer-to-Peer Remote Copy (PPRC)
 - Global Mirror for ESS, aka (asynchronous) Peer-to-Peer Remote Copy
 - Global Copy for ESS, aka PPRC-Extended distance (PPRC-XD)
 - Metro/Global Copy for ESS, aka synchronous PPRC combined with PPRC-XD
 - Metro/Global Mirror, aka synchronous PPRC combined with Global Mirror
- XRC
 - Global Mirror for zSeries, aka Extended Remote Copy (XRC)

DFSMSdss FlashCopy

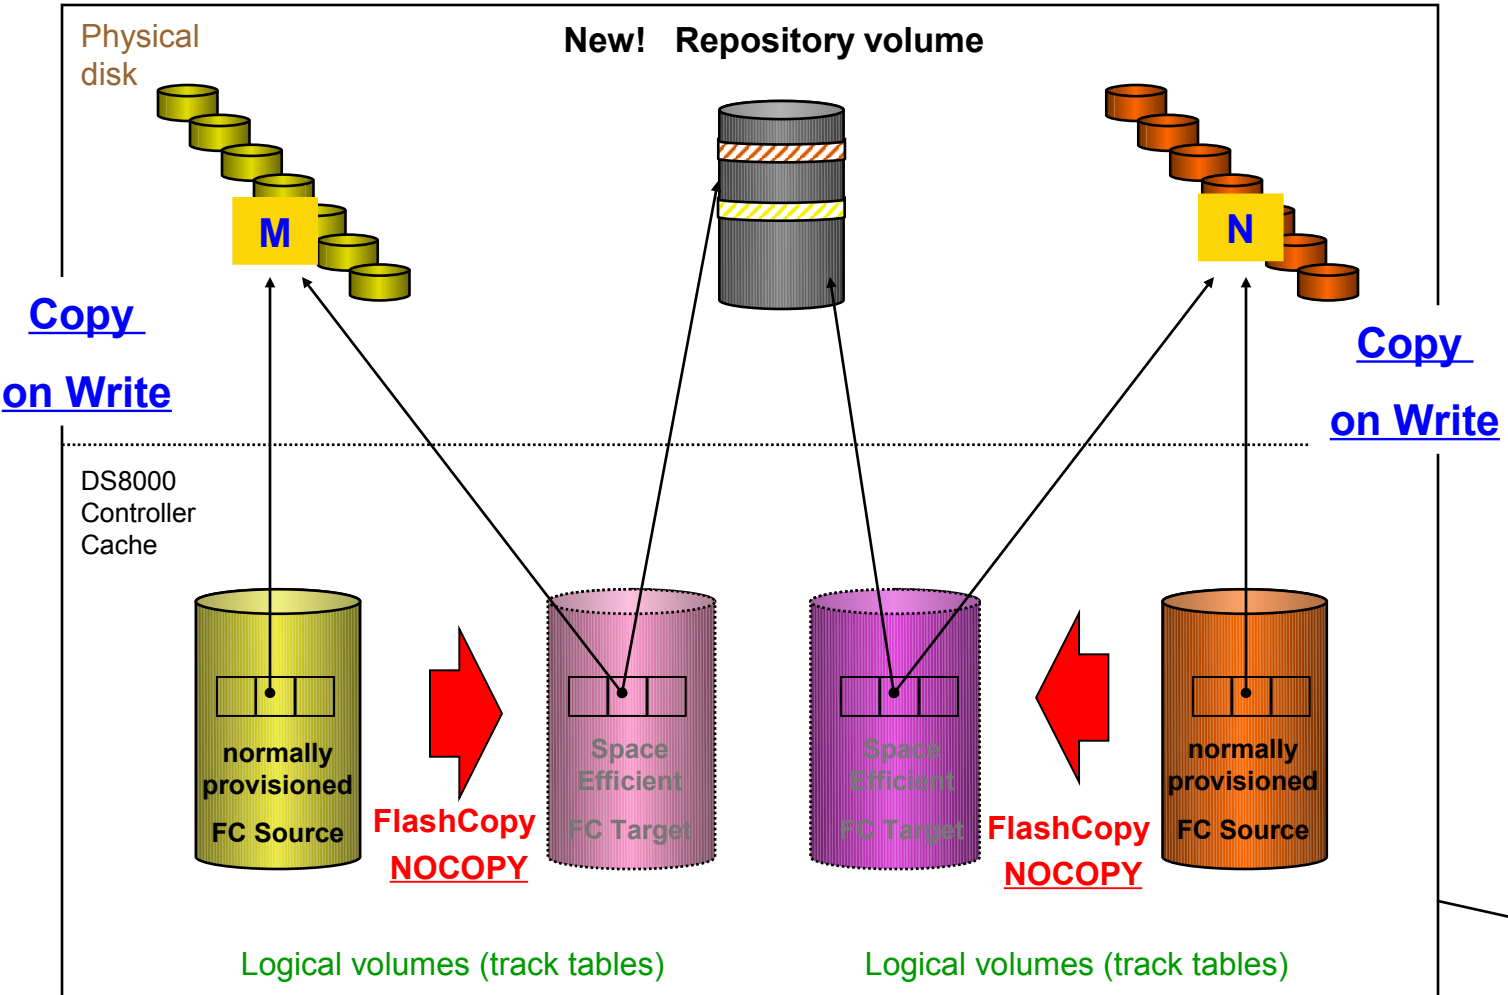
- **Functions that exploit FlashCopy**
 - COPY
 - Full volume
 - *DUMPCONDITIONING available to provide online volume clone to be used for backup purposes*
 - Data Set level
 - *Logical or Physical*
 - *Physical allows copying back a set of files using FC on a volume basis*
 - *DataSet DUMPCONDITIONING*
 - *Can Rename data sets to a temporary name at COPY time, and at DUMP time have the ability to rename them back to its original name*
 - DEFrag and CONSOLIDATE
 - Default is to use FlashCopy if possible, can revert to using CC,VCC, and traditional I/O
 - Updates catalog and performs serializations
- Volume and data set level

Review: FlashCopy NOCOPY

However, in many cases, FlashCopy NOCOPY **physical space** for targets can go unused




FlashCopy SE to Space Efficient volumes




Goal:
Smaller amount of total physical GB for Repository

Note: Repository has a defined
1) Physical Size
2) Logical Size



Establish FlashCopy to PPRC Primary prior to Remote Pair FlashCopy




Remote Pair FlashCopy

- Advantages
 - Low link/resources utilization
 - PPRC pair does not become “unsynchronized” and it is not useful for DR or hyperswap
 - ➔ FlashCopies do not disrupt PPRC sync state


State/Configuration Restrictions

- PPRC pairs may be in Full Duplex, Pending, or Suspend state
 - Pending and Suspend capability available in 6.2
- Secondary of FLC source and target are in same box
- Local or remote targets may not be Space Efficient (SE) volumes
- No cascading configuration MGM/MGC (target volumes)
- ‘Classic’ FlashCopy restrictions exist at remote
 - E.g., Because FlashCopy is established at remote, cannot have other relationship conflicts

FlashCopy to PPRC Primary using Preserve Mirror



Withdraw Preserve Mirror FlashCopy relation



Establish Preserve Mirror options

- None – Use “old way” (primary goes duplex pending, and transfers all target tracks)
- Preferred – Use Preserve Mirror (inband FLC to secondary) if possible (e.g., secondary FLC source and target are in same box), otherwise, use “old way”
- Required – Fail FLC establish if it is not possible to establish FLC on primary without causing FLC target on primary to go duplex pending. (i.e., cannot use “old way”)
 - Utilizing ADRUIXIT to force PM Required option is a trend

ADRUIXIT Example

```

ADRUIXIT CSECT
ADRUIXIT AMODE 31
ADRUIXIT RMODE 24
    STM 14,12,12(13)    SAVE REGISTERS
    USING ADRUIXIT,15    ADDRESSABILITY TO ADRUIXIT
    USING ADRUFOB,1     ADDRESSABILITY TO ADRUFOB
    SR 2,2              ZERO REGISTER 2
    CH 2,UFFUNCT       CHECK ENTRY TYPE
    BNE FUNCENT        BRANCH TO FUNCTION ENTRY
    SR 3,3              PARM CHANGE ENTRY, SAVE RC 0
    B FINISH           FINISHED
FUNCENT LH 2,UFBDOFF   GET OFFSET TO UFOFUNCT
    AR 2,1              CALCULATE ADDRESS OF UFOFUNCT
    USING UFOFUNCT,2    ADDRESSABILITY TO UFOFUNCT
    NI UFO8FLGS,X'FF'-(UFOFMPRE+UFOFPMNON+UFPMPREQ)
    OI UFO8FLGS,UFOFMPREQ PRESERVE MIRROR PRESMPREQ
    LA 3,4              SAVE RETURN CODE 4
    DROP 1              DONE USING 1 FOR ADRUFOB
    DROP 2              DONE USING 2 FOR UFOFUNCT
    DROP 15            DONE USING 15 FOR ADRUIXIT
FINISH LR 15,3         SET RETURN CODE
    L 14,12(,13)       RESTORE REGISTER 14
    LM 0,12,20(13)    RESTORE REGISTERS 0 THRU 12
    BR 14              RETURN
ADRUFOB INCLUDE ADRUFOB CONTROL BLOCK
END

```

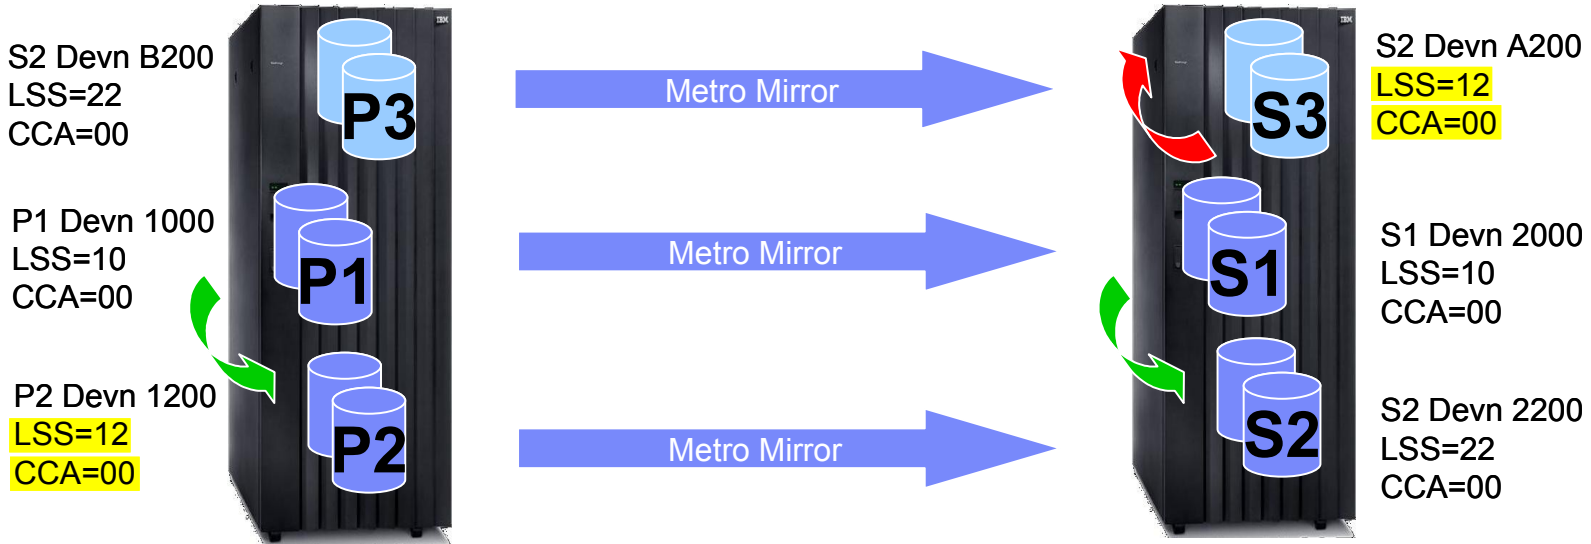

APARs of Interest

- OA38606: Performance Improvement during writes
 - Allows concurrent writes to a dataset extent
 - HW serializes write processing when volumes are PPRC primaries
 - Should see an improvement in I/O during DSS COPY (nonFC) and RESTORE processing
 - Must also install OA38579 (EXCP...requires IPL)
- OA35034: 1 TB Support
 - Support converts VSAM track allocated VSAM data sets to cylinder allocation
 - When source has >16777214 tracks and was track allocated
- OA39039: Lifting ADR918I RC26 for LDSs (DB2, zFS)
 - Prevented the use of FC when source HURBA=HARBA and targets HARBA > source HARBA

APARs of Interest

- OA37422 (SDM): FlashCopy during UCB Swap Exposure
 - Small timing window in which a FlashCopy replication command executing during a disk swap operation could be redirected to the other disk subsystem (i.e., the PPRC partner disk subsystem) and incorrectly overwrite data on this disk subsystem that was not intended to be overwritten.
 - OA37417 is enablement APAR
 - OA37420 (IOS)
 - OA37421 (FlashCopy Manager)
 - PM46645 (ICKDSF)
 - OA37423, OA37424 OA37425 (AOM, Device Support, DEVMAN)

APARs of Interest



- Logical Subsystem (LSS) number and Channel Connection Address (CCA) for source and target volumes are specified in the FlashCopy command
- The green arrows represent the Remote Pair FlashCopy operation that would result if the command is executed by the primary disk subsystem prior to HyperSwap
- The red arrow represents what would happen if the execution of the command spans a HyperSwap, ultimately being performed by the secondary disk subsystem

Reference Materials

- Introduction to FlashCopy Manager for z/OS. IBM Storage and Networking Symposium, New Orleans, LA, USA. July 25-29, 2005
- SHARE : Session 3084 Medco Clones DB2 Environments Using IBM FlashCopy and Mainstar's VCR. Boston, MA, USA. August 22, 2005
- Publications:
 - PPRC Manager
 - User's Guide and Reference – G325-2633
 - Program Directory – GI11-2905
 - FlashCopy Manager
 - User's Guide and Reference – G325-2632
 - Program Directory – GI11-2904
 - DFSMS Advanced Copy Services - SC35-0428:
- Redbooks Redpaper
 - REDP-4065-01: IBM System Storage FlashCopy Manager and PPRC Manager Overview
 - SG245680: IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services with IBM eServer zSeries
 - SC26-7445: IBM TotalStorage Enterprise Storage Server User's Guide

Reference Materials

- Publications:
 - SC35-0428: DFSMS Advanced Copy Services
 - SC35-0423: DFSMSdss Storage Administration
 - GA22-7499: z/OS V1R12 Migration
- Redbooks
 - SG245680: IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services with IBM eServer zSeries
 - SC26-7445: IBM TotalStorage Enterprise Storage Server User's Guide

Questions?