#SHAREorg

SHARE

zGM, XRC, PPRC, GM, GC, MM, FC, CC, VCC: Introduction to the Alphabet Soup of Copy Services

Lisa Gundy
IBM Corporation
lisat@us.ibm.com

Tuesday, March 11, 2014:
3:00 PM – 4:00 PM

SHARE in Anaheim

Legal Disclaimer

NOTICES AND DISCLAIMERS

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 intellectually 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

- **Disaster Recovery Considerations**
- **Copy Services Decoder Ring**
- **Understanding the Technology**
  - Point in Time Copy
  - PPRC
  - z/OS Global Mirror (XRC)
- **Copy Services Topologies / Solutions**

Disaster Recovery Considerations

- Business Objectives for Disaster Recovery
  - **Recovery Time Objective (RTO)**
    - How long can you afford to be without your systems
  - **Recovery Point Objective (RPO)**
    - When it is recovered, how much data can you afford to lose
  - Network Recovery Objective
    - How much of the network must be restored
- Business Considerations for Disaster Recovery
  - Recovery requirements dependent upon people and automation
    - Striking the correct balance
  - Capacity (MIPs and gigs) at recovery site
    - sized for production that will run there
  - Distance between data centers
- Disasters may cause multiple companies to recover, putting stress on the commercial business recovery services
- These Objectives and Considerations must be established selecting right technology

Tiers of Disaster Recovery

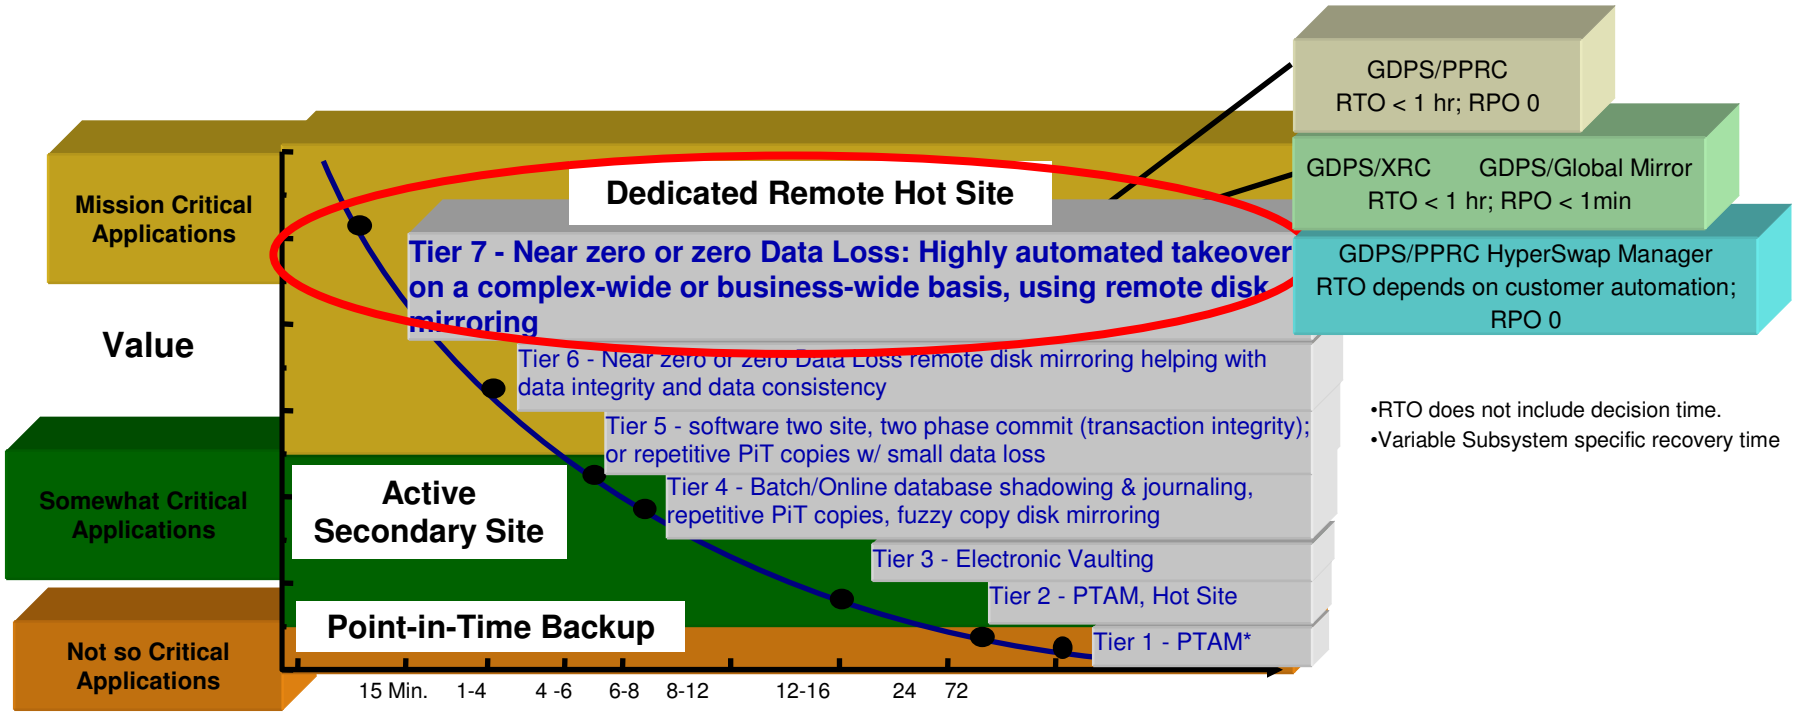

- RTO does not include decision time.
- Variable Subsystem specific recovery time

Tiers based on Share Group 1992
*PTAM = Pickup Truck Access Method

Best D/R practice is blend tiers of solutions in order to maximize application coverage at lowest possible cost . One size, one technology, or one methodology does not fit all applications

IBM Copy Services Decoder Ring

- IBM Copy Services Terminologies
  - **High Availability**
    - data availability for production
  - **Disaster Recovery**
    - data availability or recoverability in the event of a disaster
  - **Metro** – local distances
  - **Global** – extended distances
  - **Mirror** – consistency of data at secondary
    - Can be synchronous or asynchronous
  - **Copy** – data is copied, consistency not always ensured
    - In PPRC is continuous
      - *Mirror is consistent, copy is not*
    - Also 'Point in Time' (PiT) copy mechanisms (FlashCopy, Snapshot, Concurrent Copy, Virtual Concurrent Copy)

IBM Copy Services Decoder Ring

- Geographies Terminologies
  - **2-Site**
    - Can be single region/geography or two regions/geographies
      - *2-site / Single region with HA*
      - *2-site / Two regions for D/R*
  - **3-Site**
    - Typically Two regions
      - *2-site at production region with HA*
      - *Third site out of region for D/R*
  - **4-Site**
    - Typically Two geographies
      - *2-site at one region with HA*
      - *2-site at second region with HA*
      - *Asynchronous mirror between geographies for D/R or site switch*
    - Also used for Active/Active

Understanding the Technology

- Point in Time Copy
  - FlashCopy
  - Concurrent Copy
  - Virtual Concurrent Copy
- PPRC
  - Metro Mirror (Synchronous PPRC)
  - Global Mirror (Asynchronous PPRC Mirror)
  - Global Copy (PPRC Extended Distance Copy)
- XRC
  - z/OS Global Mirror (XRC, zGM)
  - z/OS Metro/Global Mirror (XRC and Synchronous PPRC)
    - Also z/OS Global Mirror Incremental Resync (with GDPS)

IBM Disk Storage Advanced Copy Services

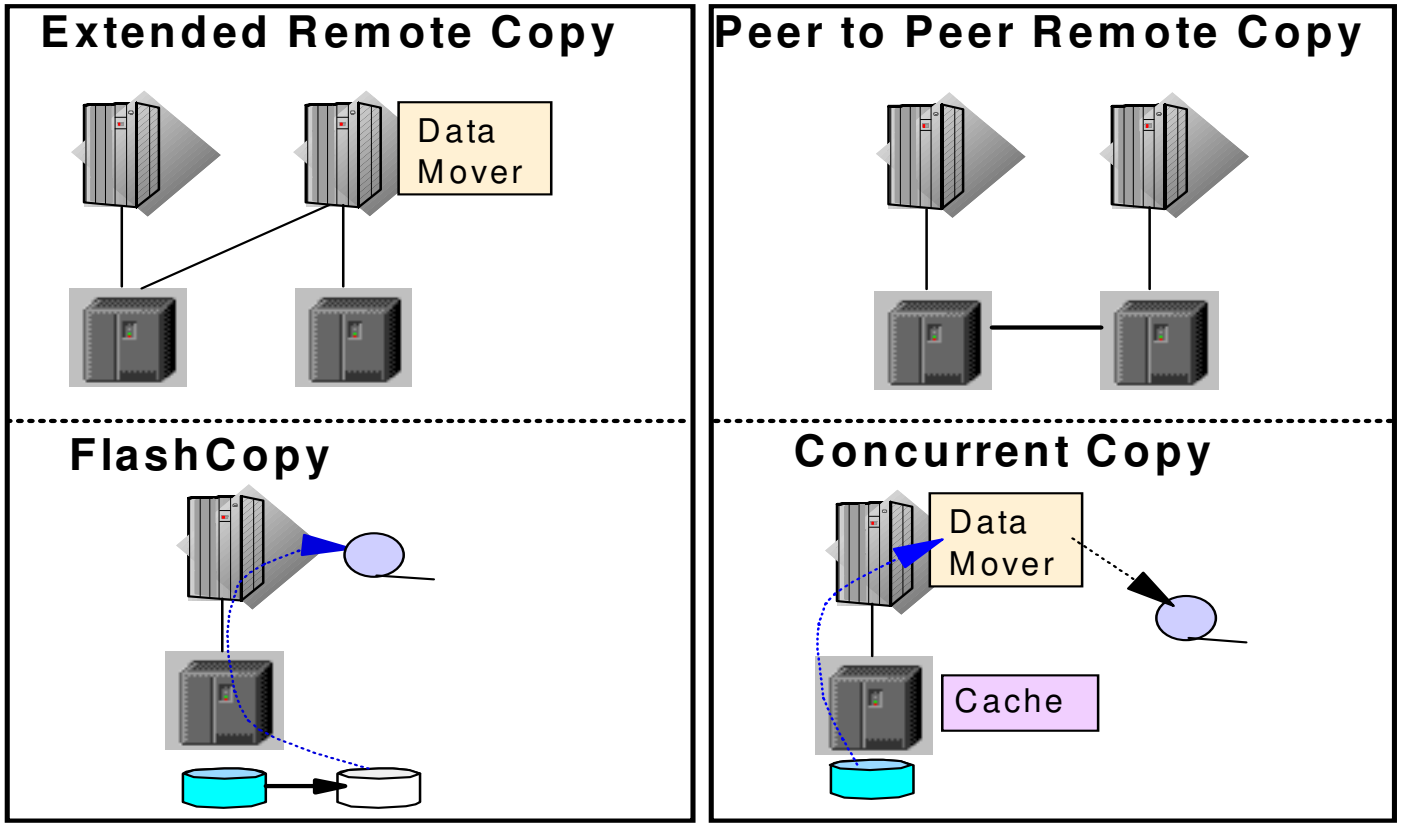

FLASHCOPY

FlashCopy

- **Point in Time copy**
- **Establishes logical copy in seconds**
- **Source and target quickly available for full read/write**

Options
- Full volume FlashCopy
- Dataset FlashCopy (z/OS only)
- FlashCopy consistency groups
- Incremental FlashCopy
- Inband FlashCopy
- Space Efficient FlashCopy
- Fast Reverse Restore
- Remote Pair FlashCopy

Uses
- Online backup
- Tape backup
- Moving datasets
- Practice Copy
- Safety Copy

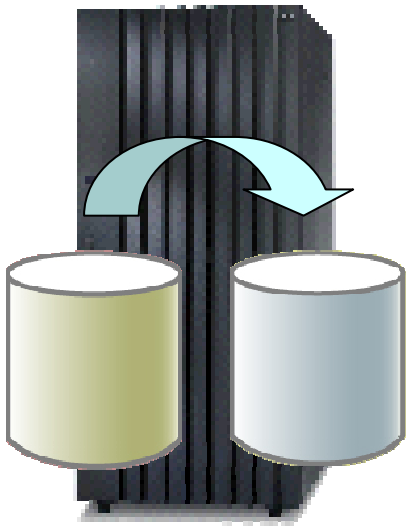

FlashCopy

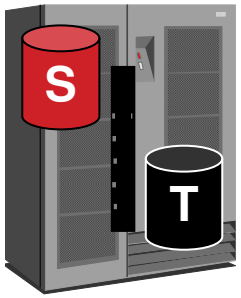

PiT copy technology on the ESS
When a FlashCopy is issued the copy is available immediately
A bitmap tracks the relationship between source and target tracks

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

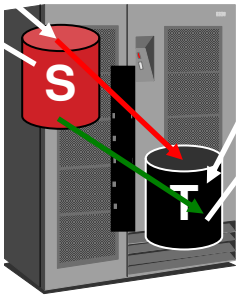

Writes to the source may cause a copy on write if the track has not been copied to the target

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

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

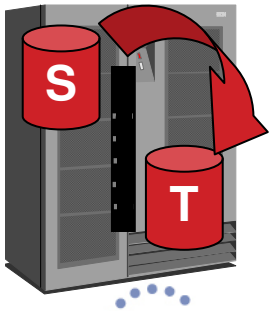

Dataset Level FlashCopy

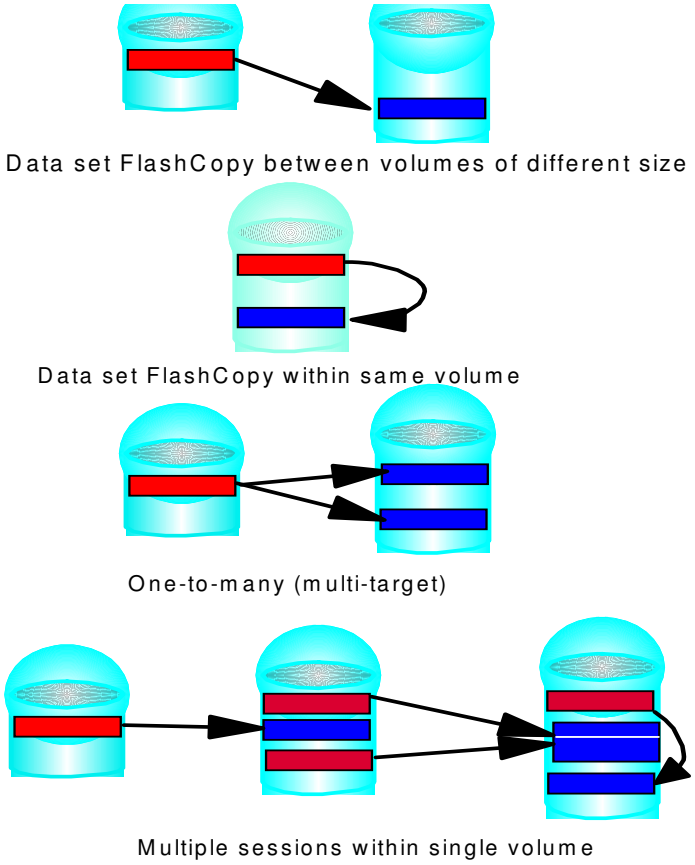

- All dataset types
- Source and Target may be the same volume
- Any track can be the source of up to 12 target relationships
- Any track can only be a source OR a target

Miscellaneous FlashCopy Options

- Incremental FlashCopy
- Space Efficient FlashCopy
- Remote Pair FlashCopy for PPRC
  - Allows FlashCopy between Metro Mirror primaries, FlashCopy is mirrored and performed between secondaries
- No Background Copy
- NOCOPY to COPY conversion
  - Allows FlashCopy to be established with NOCOPY, and converted to a COPY when there would be less impact
  - Starts background copy for existing NOCOPY relationship
- Deleted Dataset Withdraw (DDSW)
  - Initiates background copy for source relationships, withdraws target relationships

CONCURRENT COPY

Concurrent Copy

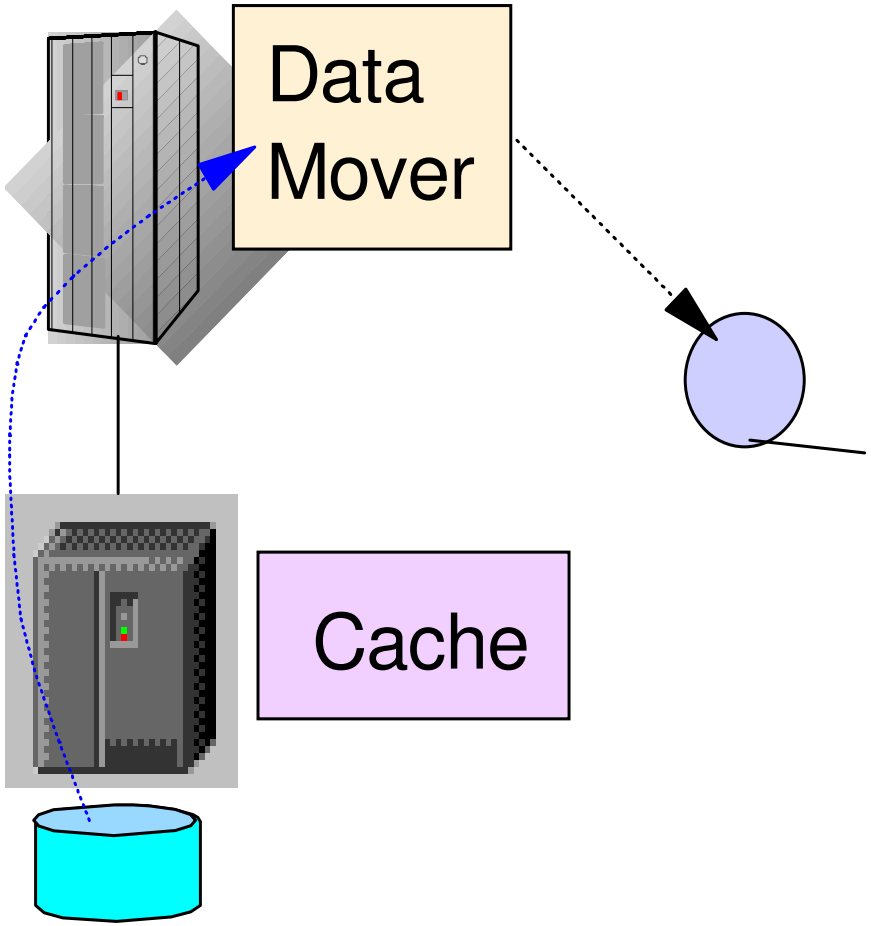

- Point in time copy
- Generate copy/dump while dataset is being updated
- Volume and dataset level

Virtual Concurrent Copy

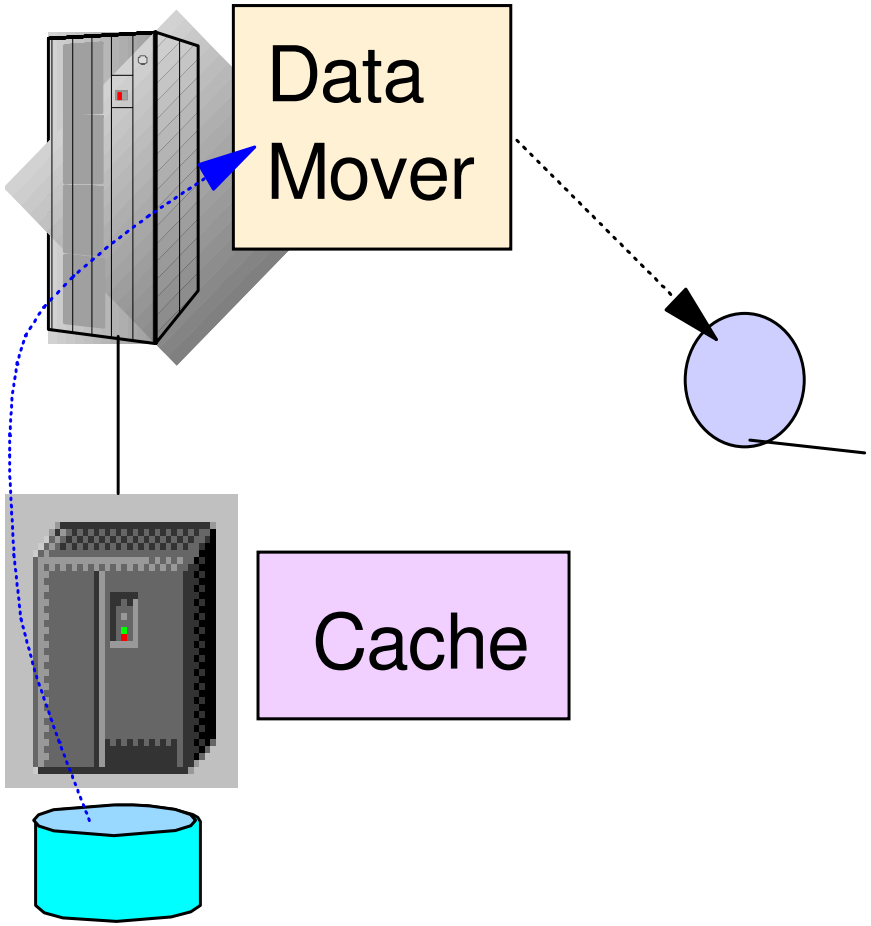

- Point in time copy
- Generate copy/dump while dataset is being updated
- Volume and dataset level
- User identifies Working Space Datasets (WSDSs) for the Data Mover to use
- WSDSs allocated on FlashCopy eligible devices
- FlashCopy used to virtually move data faster

Concurrent Copy Phases

- CC session initialization
- Copying the data
- Intercepting writes
- CC session termination

Concurrent Copy Considerations

- Appropriate time to start backup
- Many jobs can cause auxiliary storage shortage
- 16 sessions per devices
- 64 sessions per LSS

PPRC

Peer to Peer Remote Copy (PPRC

- Metro Mirror
- Global Copy
- Global Mirror

PPRC (Metro Mirror and Global Copy)

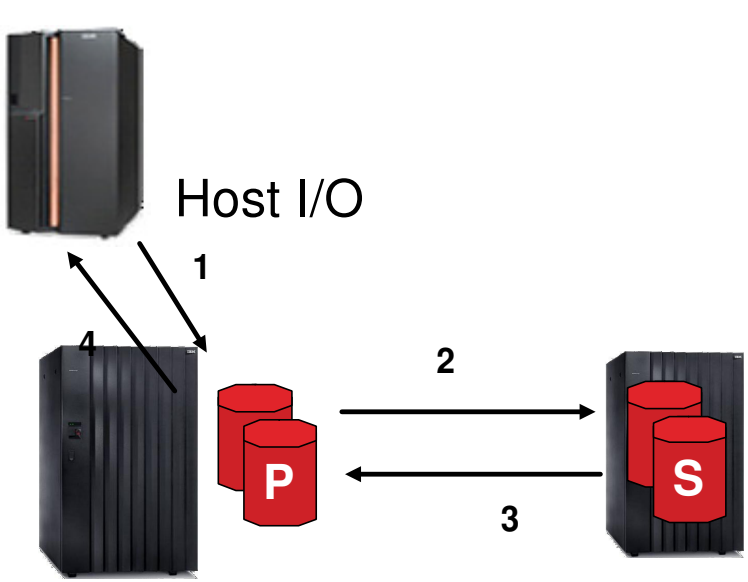

Metro Mirror

•Writes are performed synchronously

1) Write to primary volume.

2) The primary disk subsystem initiates an I/O to the secondary disk subsystem to transfer the data.

3) Secondary indicates to the primary that the write is complete.

4) Primary acknowledges to the application system that the write is complete.

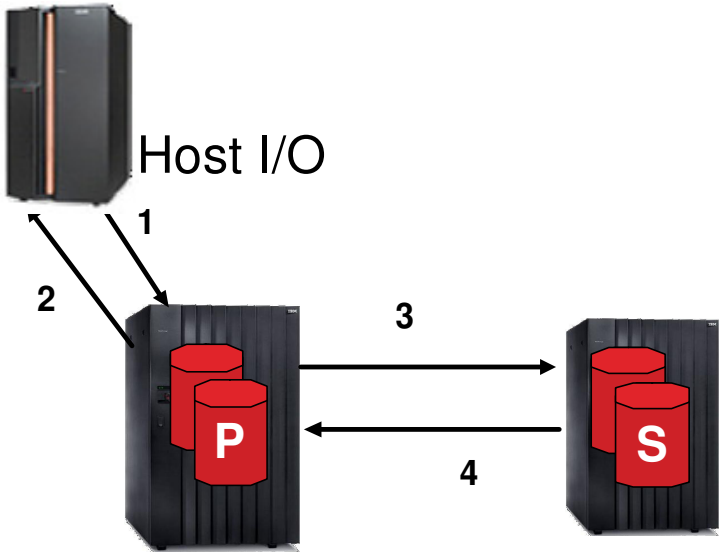

Global Copy

•Writes are performed asynchronously

1) Write to primary volume.

2) Primary acknowledges to the application system that the write is complete.

At some later time:

3) The primary disk subsystem initiates an I/O to the secondary disk subsystem to transfer the data.

4) Secondary indicates to the primary that the write is complete.

5) Primary resets indication of modified track.

PPRC-XD data transmission

- Each LSS has one or more tasks sending data that needs to be copied
- The tasks run through the devices in sequence sending any tracks that need to be copied
- The number of tasks are influenced by the number of DAs and numbers of PPRC links

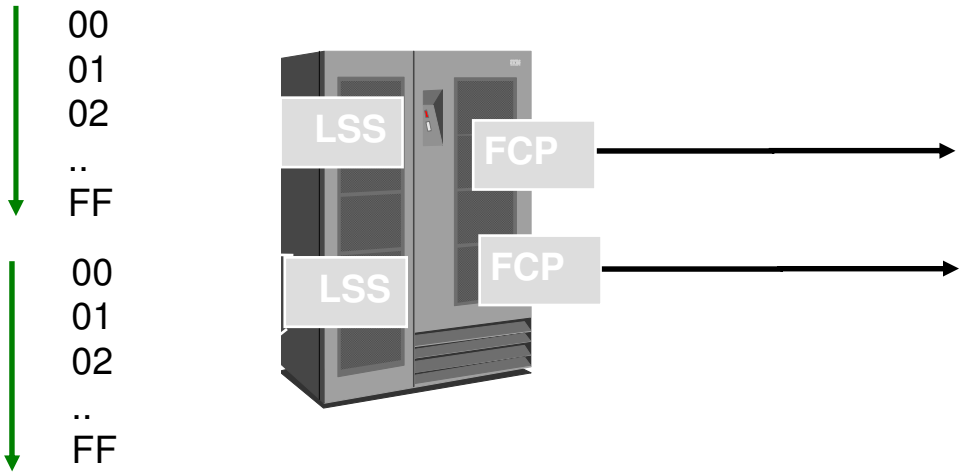

PPRC communication

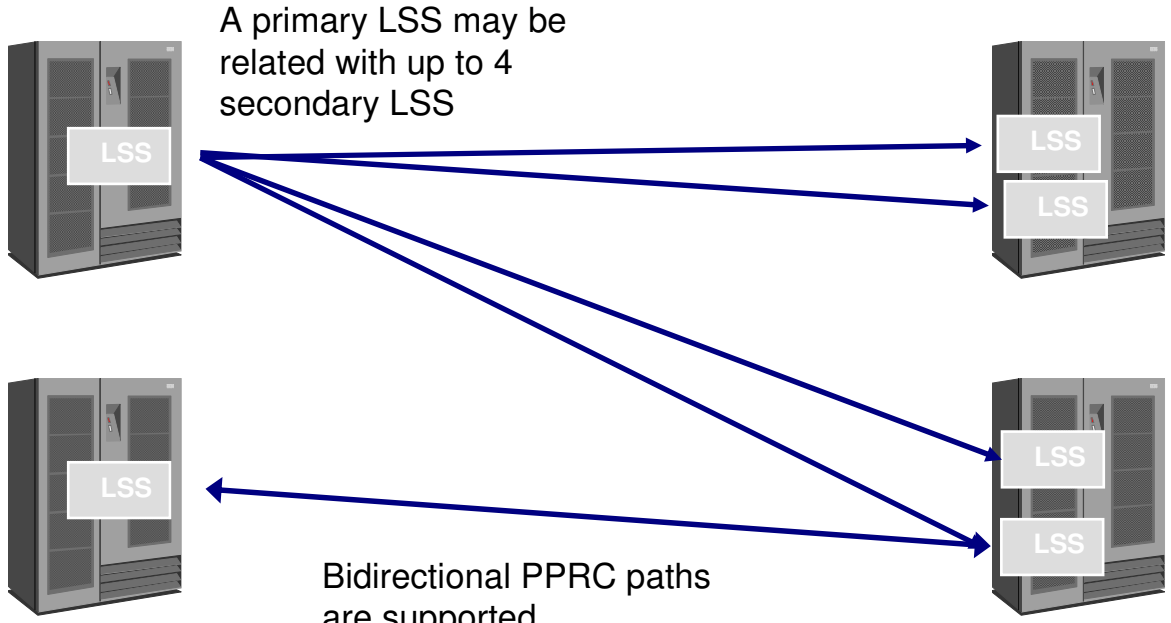

Failover / failback

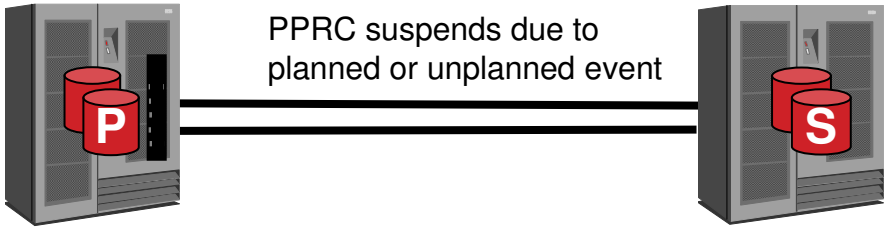

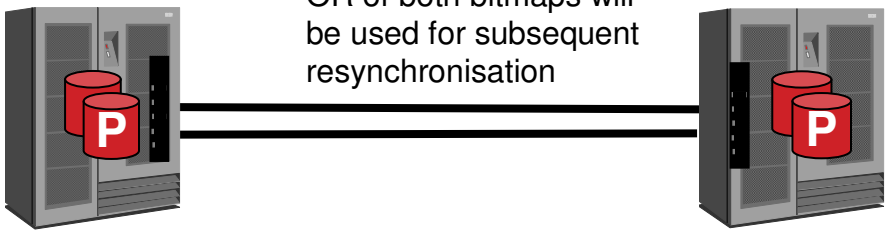

Failover command issued to allow access to secondary devices and placed them in primary suspended state.

Failback command can be using in either direction to resynchronise the PPRC relationship

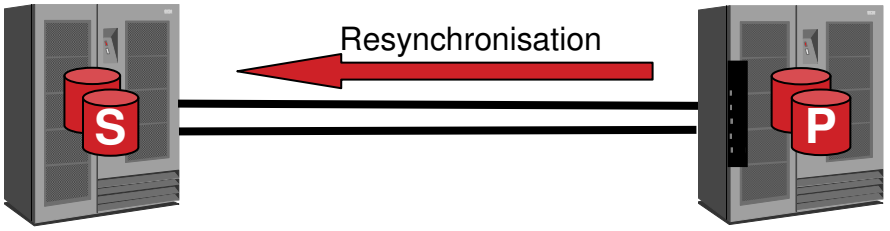

Global Mirror illustration

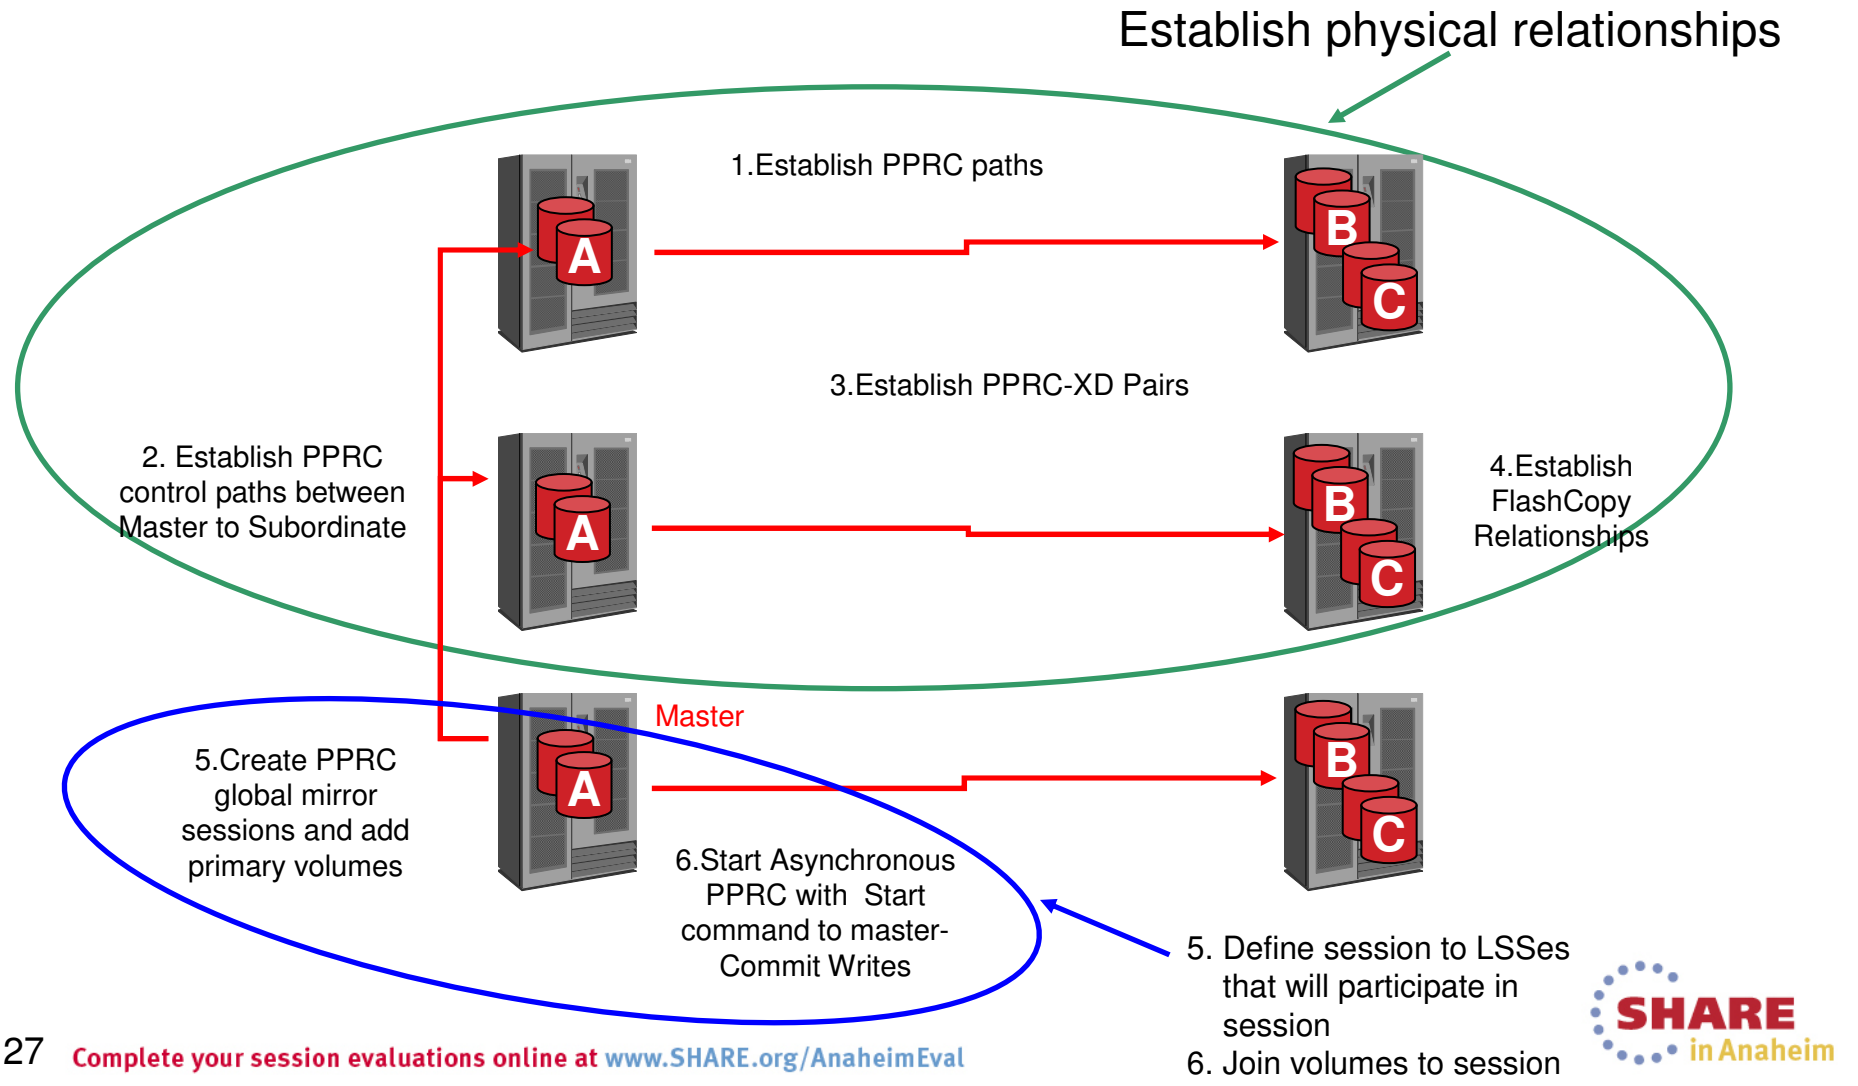

z/OS Global Mirror (XRC)

XRC System – Simple View

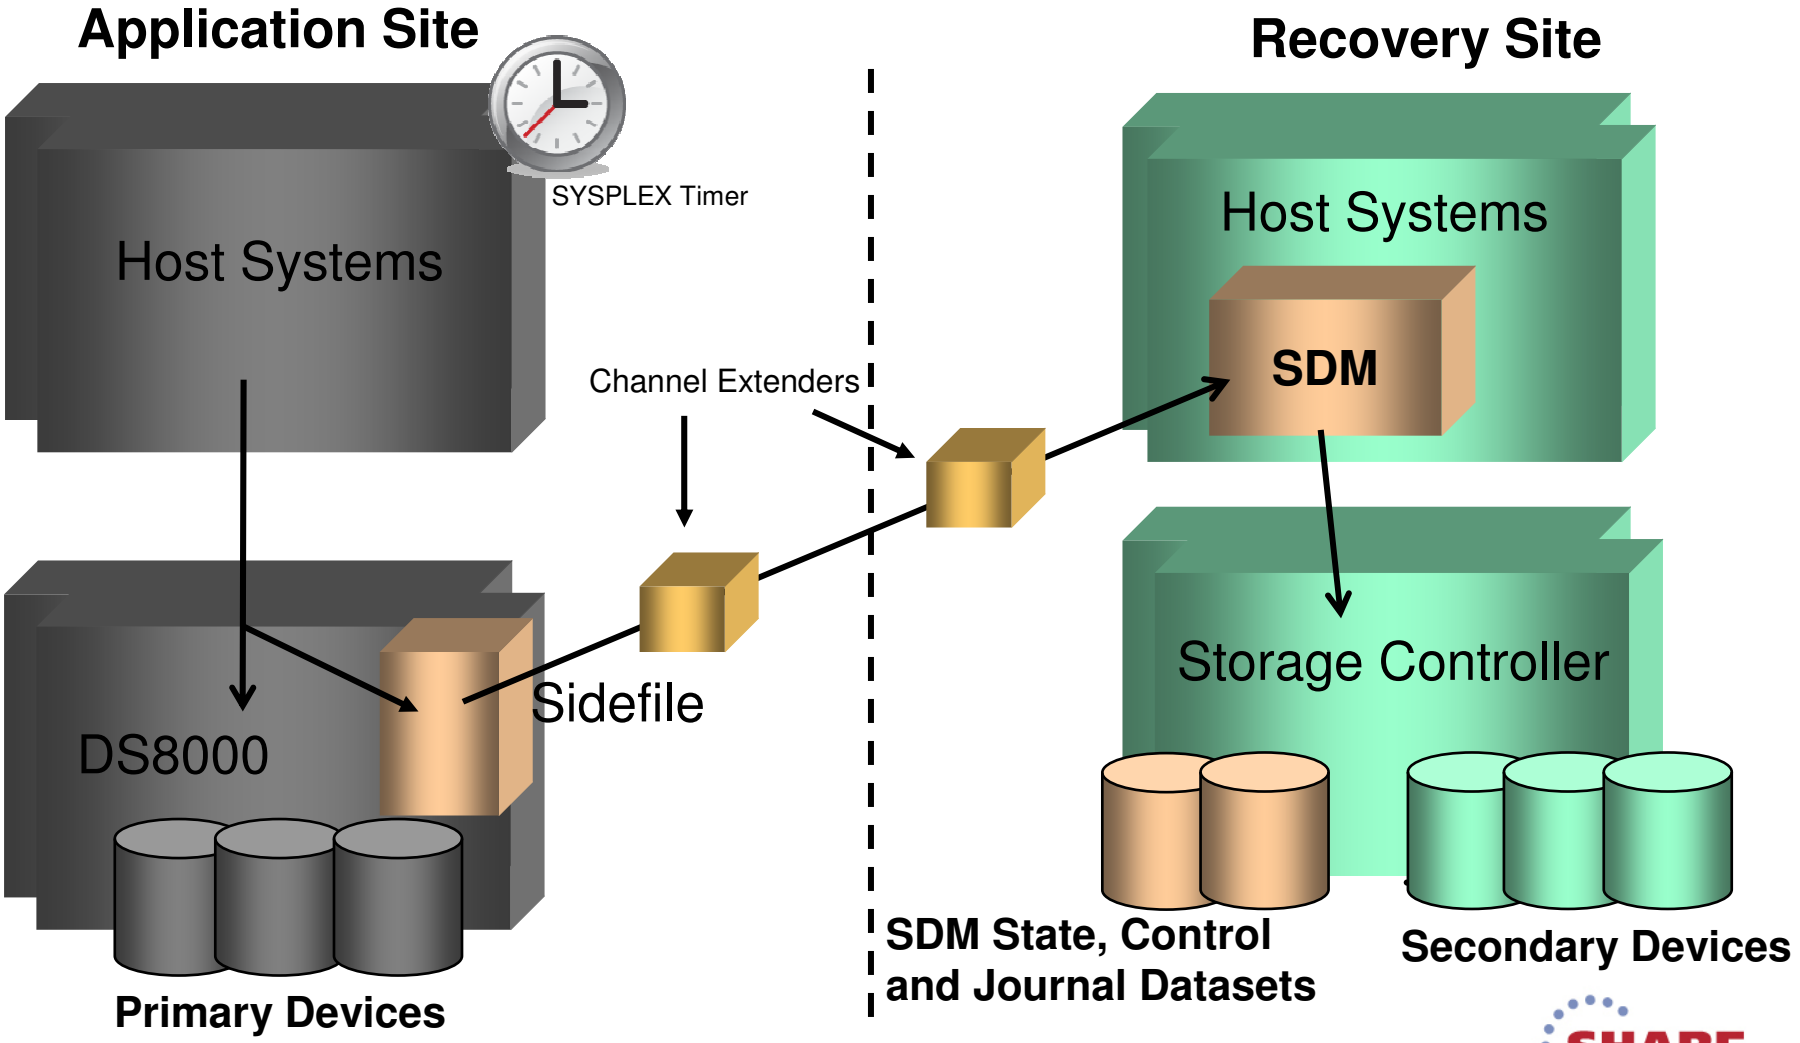

XRC Operation

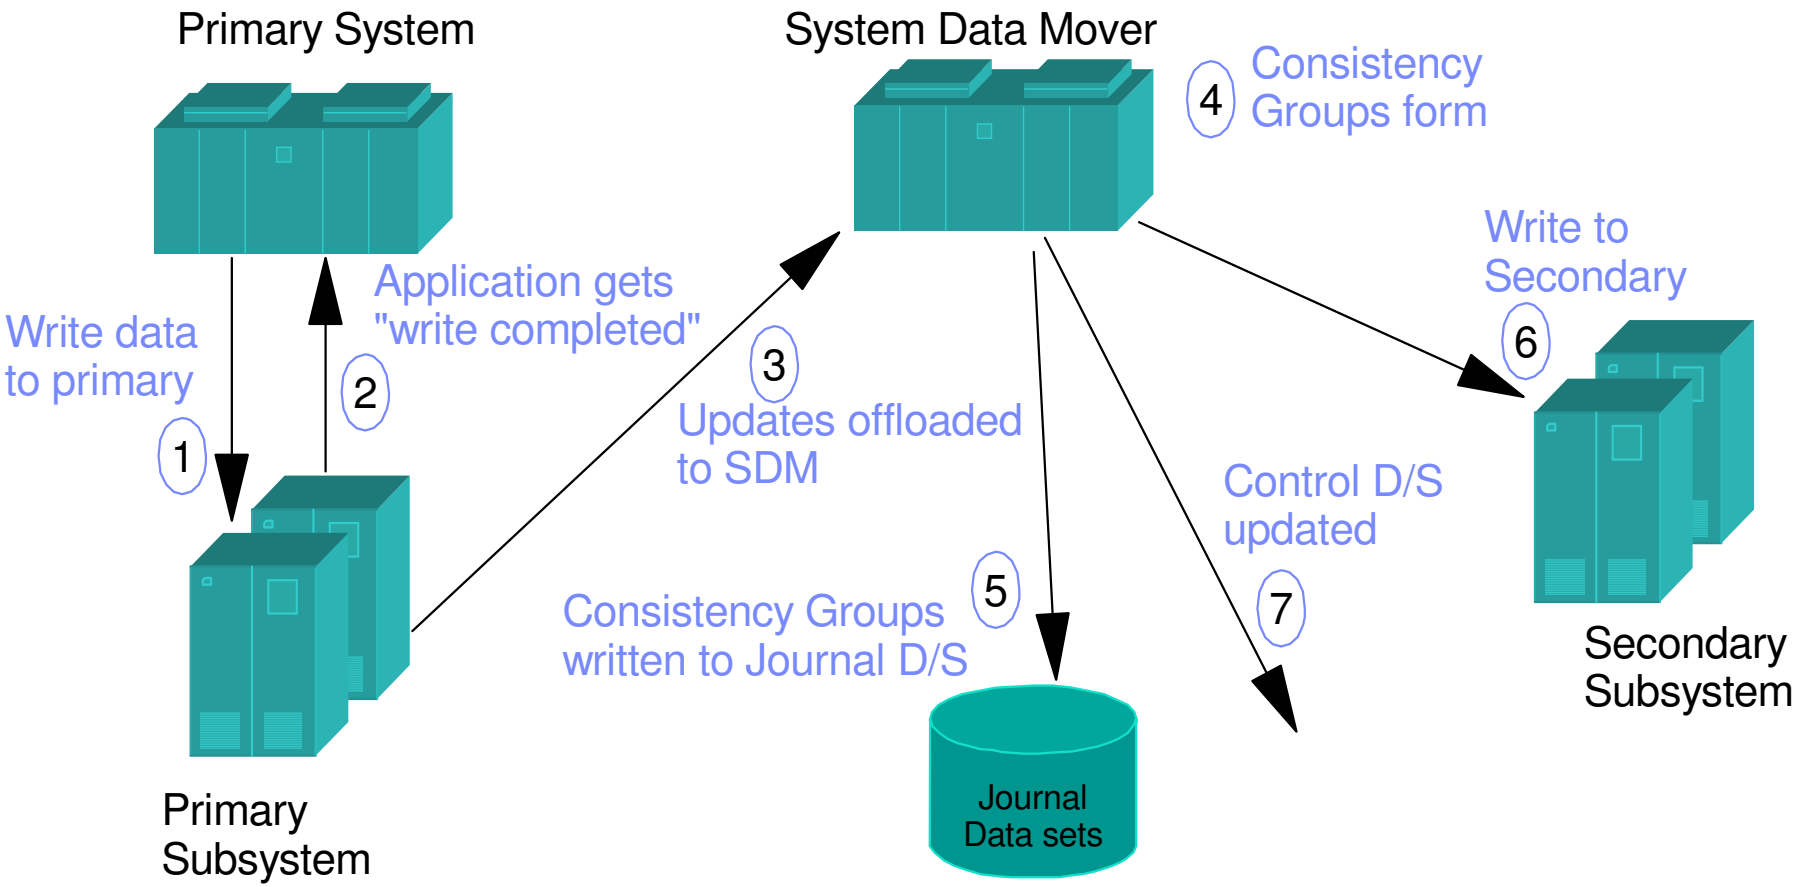

IBM Copy Services Topologies

IBM Copy Services Solution Topologies

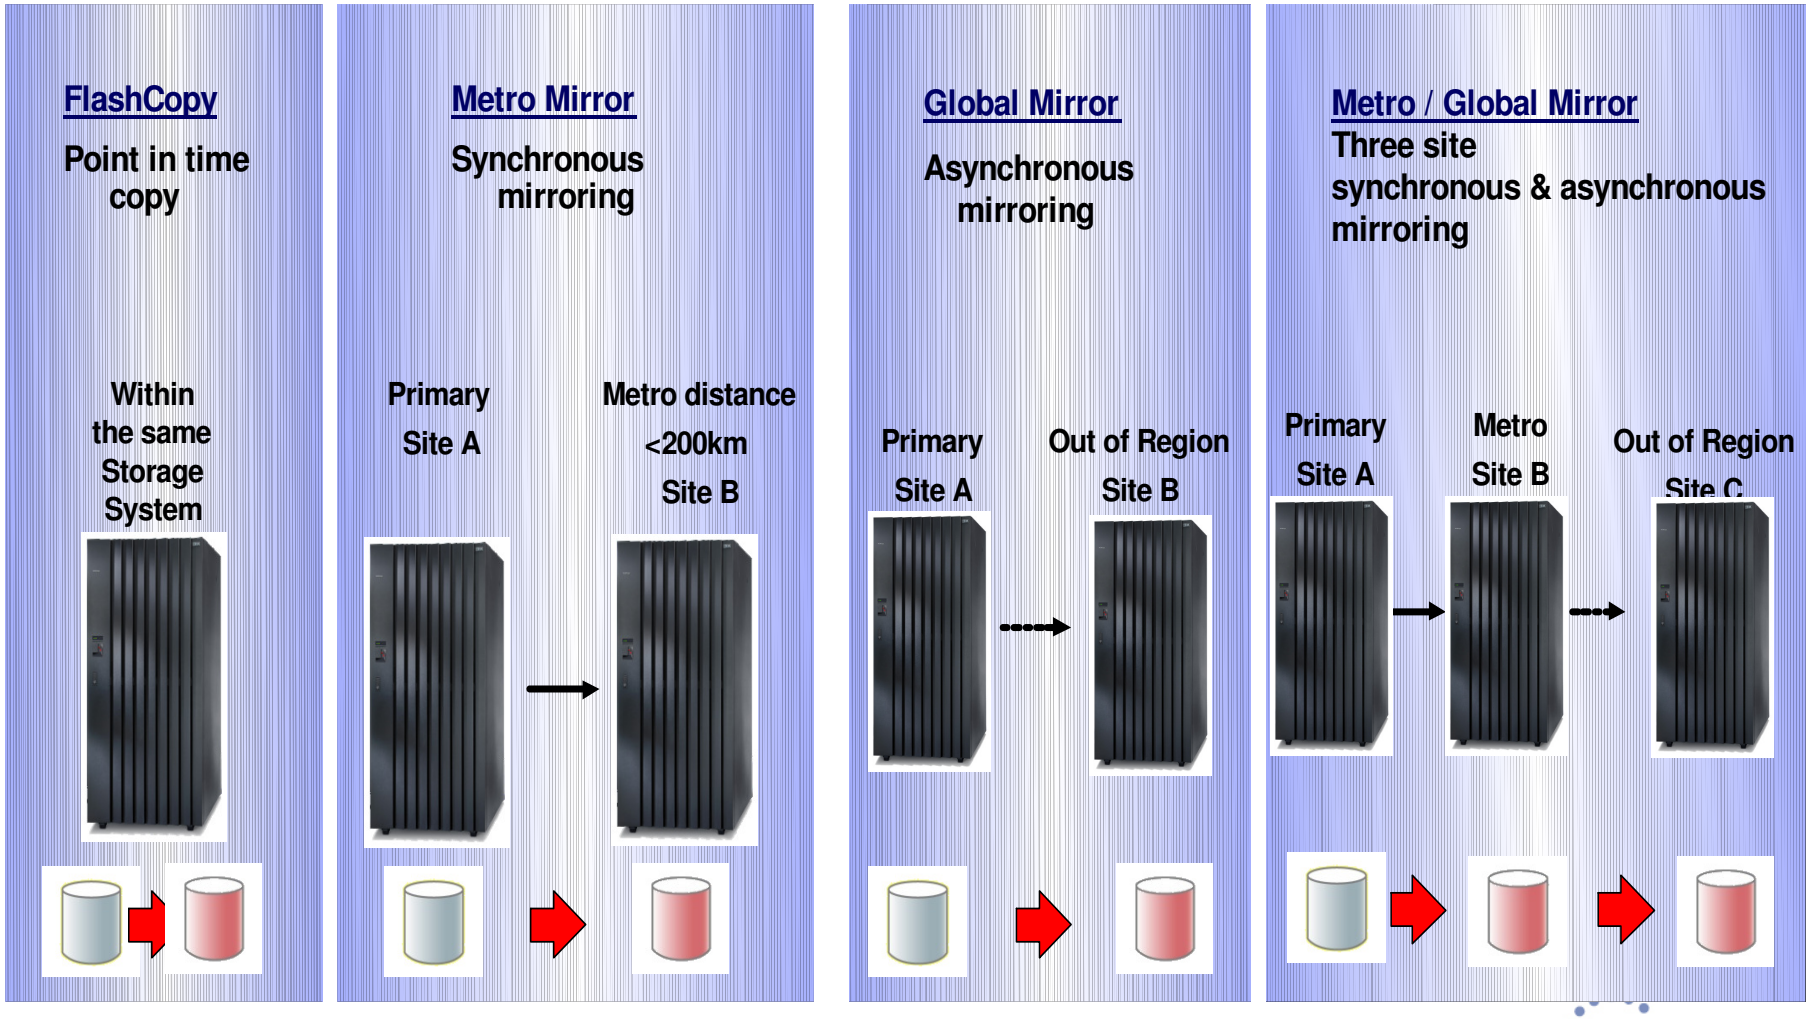

IBM Copy Services Solution Topologies

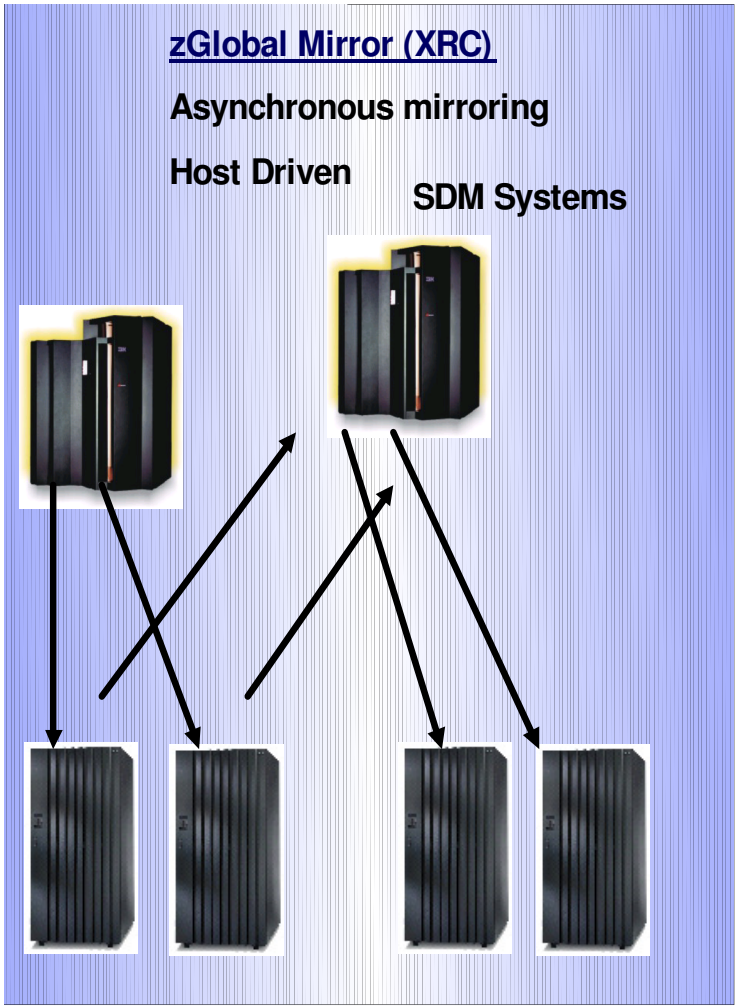

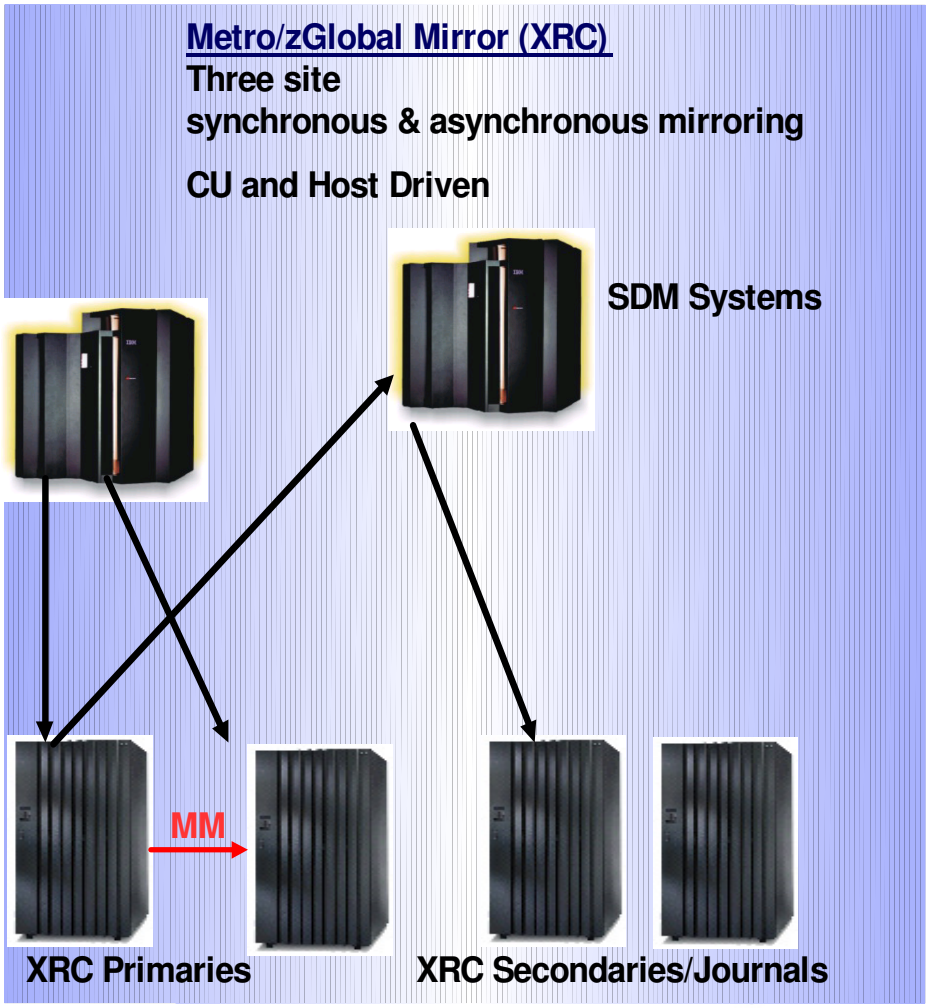

4-site topology with Metro Global Mirror

Global Copy from Global Mirror secondary devices already used for datacentre migrations

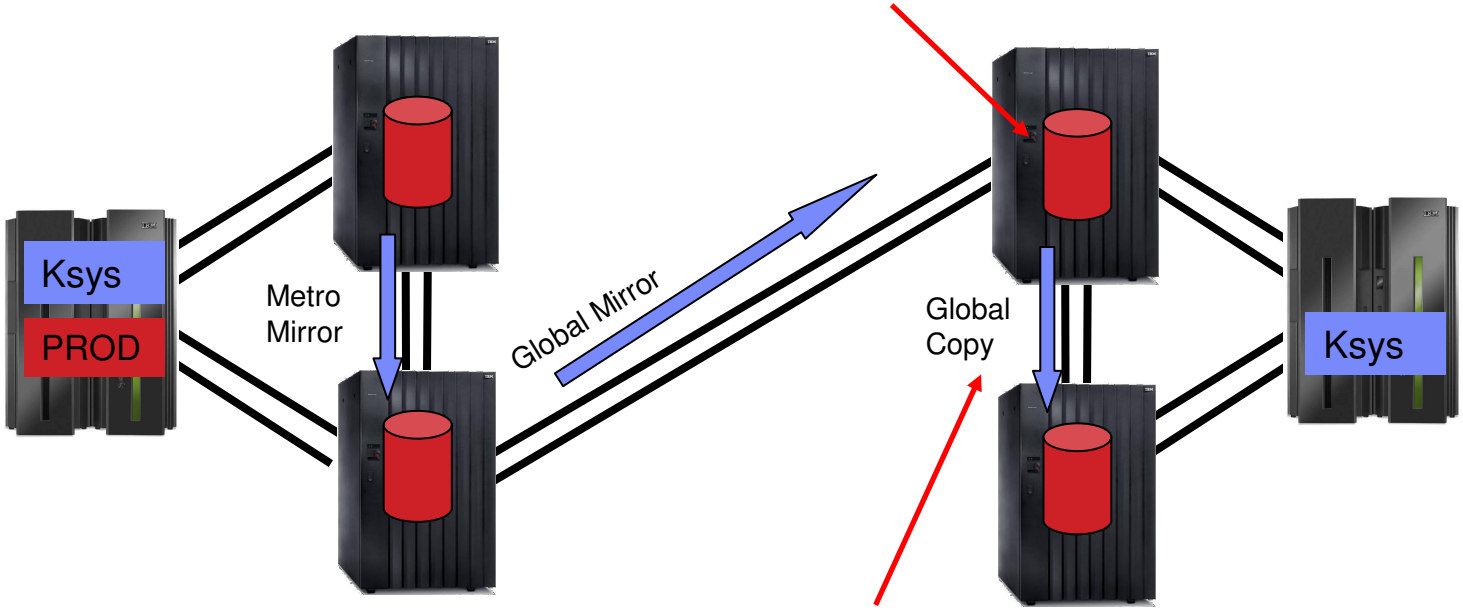

Global Copy in secondary site converted to Metro Mirror in case of disaster or planned site switch

DFSMS Interfaces

- DFSMS SDM provides primitives that are used by IBM products and can be used by ISV programs (or built into CLISTs or REXX execs)
  - Application Programming Interfaces
    - ANTRQST API
    - ANTTREXX
  - TSO Commands
- ICKDSF provides primitives that are used by IBM products and can be built into CLISTs for z/OS and z/VM devices
- DFSMS DSS provides Fast Replication exploitation via JCL and Panel interfaces to perform dataset services, including FlashCopy
- DFSMS HSM provides Fast Replication exploitation in their Copy Pool function
  - Part of DB2 CDP

Summary

- **Disaster Recovery Considerations**
- **Copy Services Terminologies**
- **Copy Services Technologies**
  - Point in Time Copy
    - FlashCopy
    - Concurrent Copy
  - PPRC
    - Metro Mirror
    - Global Copy
    - Global Mirror
  - z/OS Global Mirror (XRC)
- **Copy Services Topologies**
- **Questions?**

Reference Materials

- Publications:
  - SC35-0428: DFSMS Advanced Copy Services
  - SC35-0423: DFSMSdss Storage Administration
- Redbooks
  - SG245680: IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services with IBM eServer zSeries
  - SC26-7445: IBM TotalStorage Enterprise Storage ServerUser's Guide