#SHAREorg

Continuing the understanding of IBM Copy Services: Peer-to-Peer-Remote-Copy (PPRC) and Point in Time Copy (FlashCopy) for High Availability (HA) and Disaster Recovery (DR)

Lisa Gundy

IBM Corporation

lisat@us.ibm.com

Wednesday, March 12, 2014:

11:00 AM – 12:00 PM

Session 15077

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

- HA and DR Considerations
- The Role of Peer to Peer Remote Copy
- The Role of FlashCopy
- Solutions / Configurations based on Business Requirements

Business Continuity 101

HA and DR Considerations

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

Aspects of Availability

High Availability

Fault-tolerant, failure-resistant infrastructure supporting continuous application processing

Protection of critical business data

Recovery is predictable and reliable

Continuous Operations

Non-disruptive backups and system maintenance coupled with continuous availability of applications

Operations continue after a disaster

Costs are predictable and manageable

Disaster Recovery

Protection against unplanned outages such as disasters through reliable, predictable recovery

Customer Business Objectives

- Determine business continuity objectives:
  - Recovery Time Objective (RTO) –
    How long can you afford to be without your systems?
  - Recovery Point Objective (RPO) –
    How much data can you afford to lose / recreate?

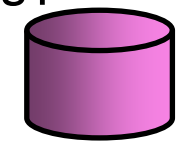

- Select technology(s) to meet business objectives

SYNCHRONOUS Remote Copy

METRO MIRROR

Continuous data availability.
Use when:

- ***Response time impact*** **is acceptable**
- **Within** ***metro distance***
- ***No data loss*** **is the objective**
- ***Fastest recovery time*** **is required**

ASYNCHRONOUS Remote Copy

GLOBAL MIRROR

Extended distance disaster recovery. Use when:

- ***Smallest response time impact*** **to applications is required**
- ***Extended distance*** **disaster recovery is the objective**
- ***Minimal data loss*** **is acceptable**

Lessons Learned About IT Survival

- Repeated Testing before a disaster is crucial to successful recovery after a disaster
  - TTWYR – Test The Way You Recover
  - RTWYT – Recover The Way You Test
- After a disaster, everything is different

- Repeated Testing before a disaster is crucial to successful recovery after a disaster
  - **TTWYR – Test The Way You Recover**
  - **RTWYT – Recover The Way You Test**
- After a disaster, everything is different
  - **Company will benefit greatly from well-documented, tested, available and automated … recovery procedures**

- Failover capacity can be obtained by
  - Prioritizing workloads
  - Exploit new technology: Capacity Back Up (CBU)
- Data backup planning and execution must be flawless
  - Disk mirroring required for <12hr RTO (need 2x capacity)
  - Machine-readable data can be backed up; not so for paper files
- Check D/R readiness of critical suppliers, vendors

Automation: Critical for successful rapid recovery and continuity

- The benefits of automation:
  - Allows business continuity processes to be built on a reliable, consistent recovery time
  - Recovery times can remain consistent as the system scales to provide a flexible solution designed to meet changing business needs
  - Reduces infrastructure management cost and staffing skills
  - Reduces or eliminates human error during the recovery process at time of disaster
  - Facilitates regular testing to help ensure repeatable, reliable, scalable business continuity
  - Helps maintain recovery readiness by managing and monitoring the server, data replication, workload and the network along with the notification of events that occur within the environment
- Monitor and manage for planned and unplanned events

Automate - Automate - Automate

Peer to Peer Remote Copy (PPRC)

Metro Mirror (synchronous PPRC) Overview

Metro Mirror

- Synchronous remote data mirroring
  - Application receives "I/O complete" when both primary and secondary disks are updated
- Typically supports metropolitan distance
- Performance impact must be considered
  - Latency of 10 us/km
- z/OS and open data supported

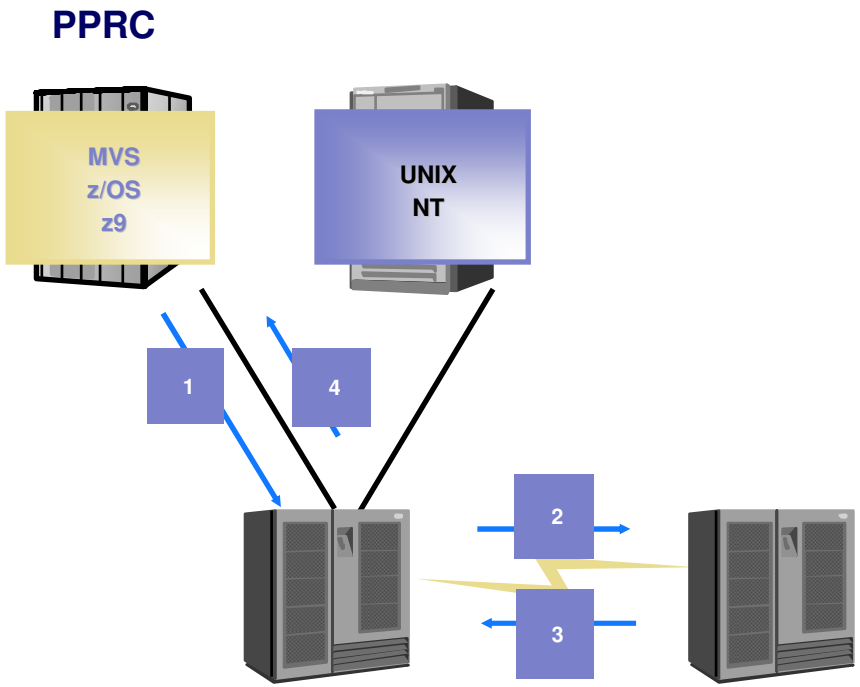

PPRC communication

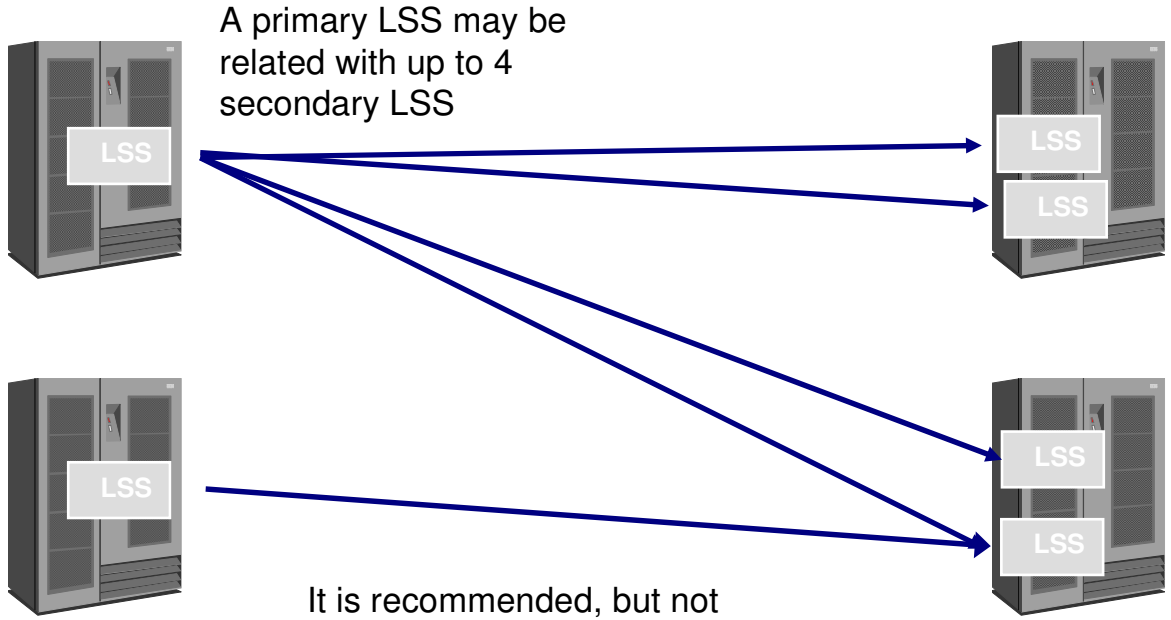

PPRC FREEZE

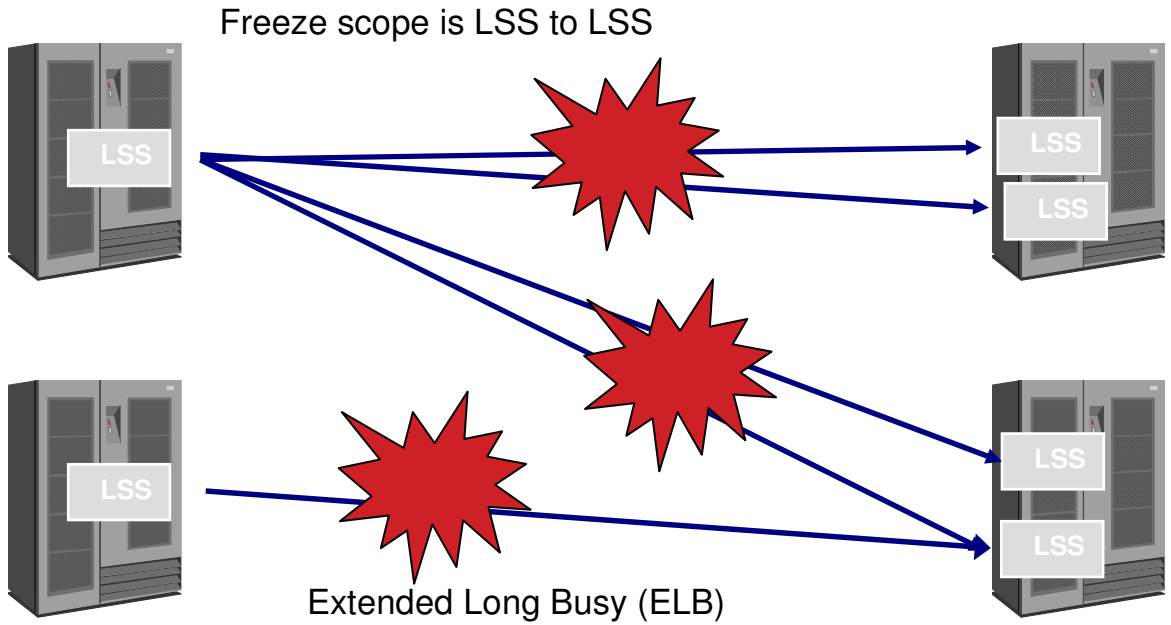

HyperSwap – the Technology

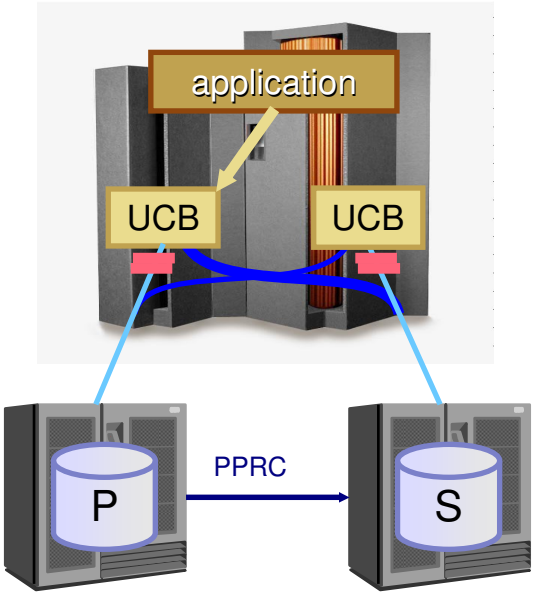

- Substitutes PPRC secondary for primary device
  - No operator interaction - GDPS-managed
  - Can swap large number of devices - fast
  - Includes volumes with Sysres, page DS, catalogs
  - Non-disruptive - applications keep running

Brings different technologies together to provide a comprehensive application and data availability solution

Planned Swap

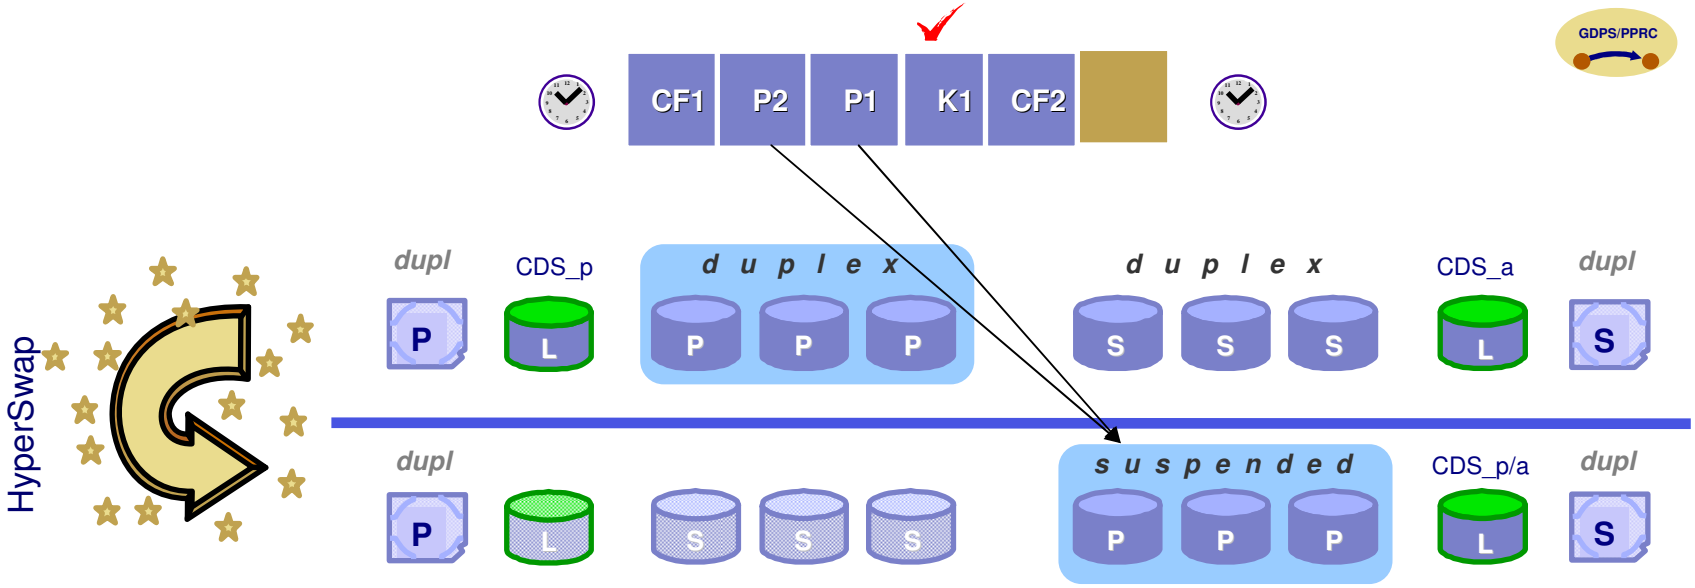

Without HyperSwap

PLANNED ACTION INITIATED

shutdown systems, remove systems from Sysplex, reverse PPRC (suspend PPRC), restart systems

1-2 hrs (approx)

With HyperSwap and FO/FB

15 Seconds! (6545 vol pairs, 19.6 TB, 46 LSS's)!

PPRC Failover, swap the primary & secondary PPRC UCBs, systems continue

P1, P2 remain active throughout the procedure

Unplanned Swap

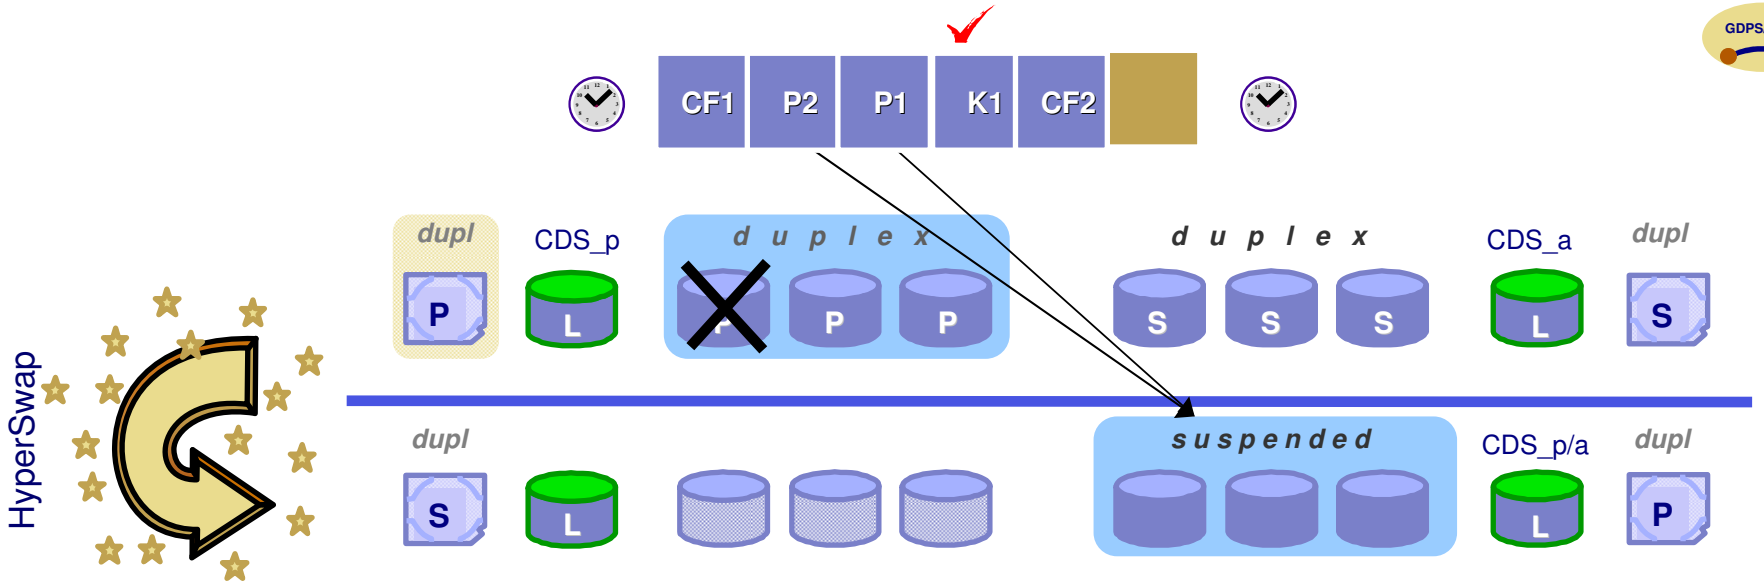

Without HyperSwap

TAKEOVER PROMPT — Decision Time — TAKEOVER STARTED

FREEZE

systems quiesced

if Freeze & Stop/Cond

//

remove systems from Sysplex

restart systems

30 – 60 Minutes or more

With HyperSwap and FO/FB

13 Seconds! (6545 volume pairs, 19.6 TB, 46 LSSs)

Only changed data needs to be copied to restore to original configuration

PPRC Failover, swap the primary & secondary PPRC UCBs, systems continue

P1, P2, remain active throughout the procedure

Global Mirror Overview

GM (Global Mirror)

- Asynchronous remote data mirroring
- Unlimited distance support
- Performance impact negligible
- Copy consistency: managed autonomically by Master Control Server in master storage server
- Up to 16 ESSs in GM session (w/RPQ)
- Supports System z and OPEN Systems data

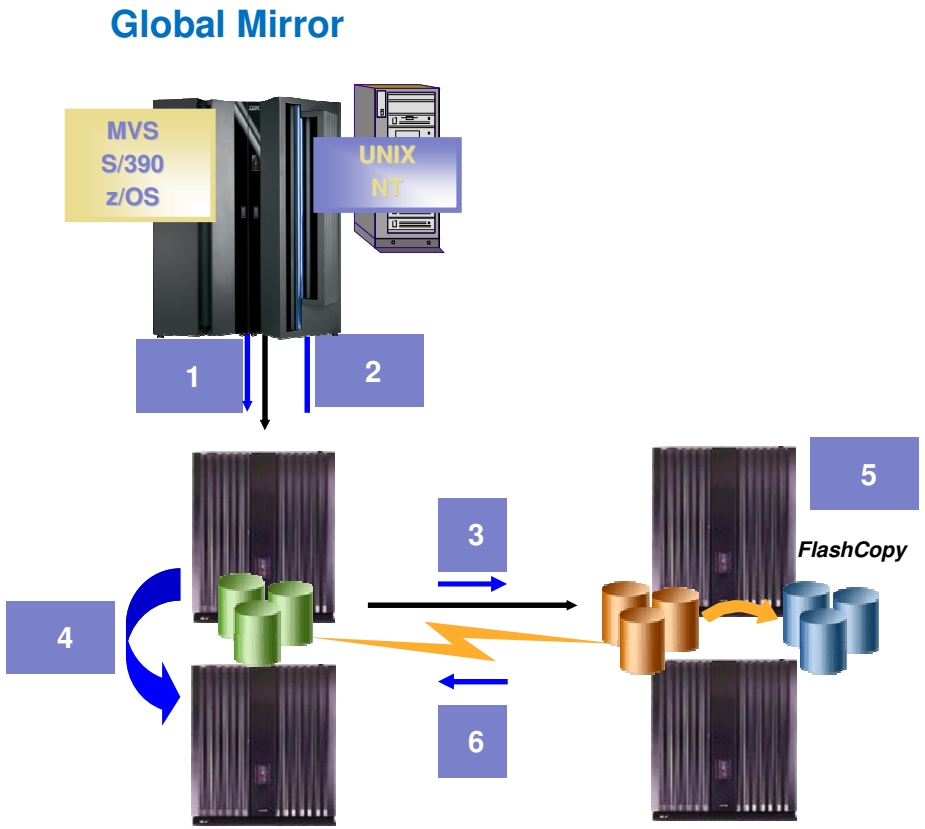

Global Mirror – Site 1 Failure

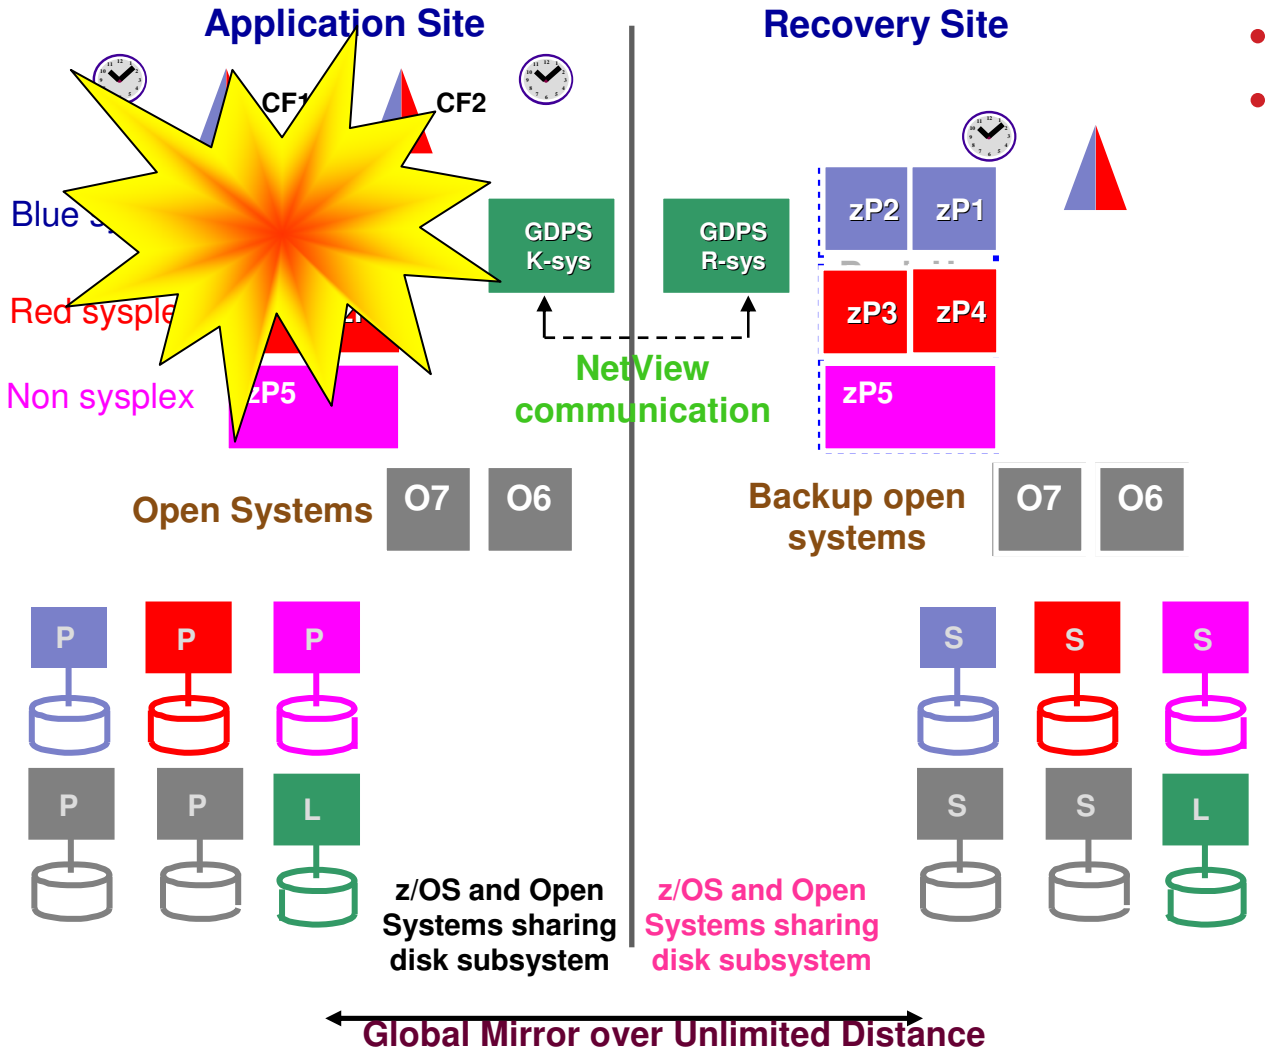

- **RTO < 1 hour**
- **RPO < 1 minute**
  - (depends on bandwidth)

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

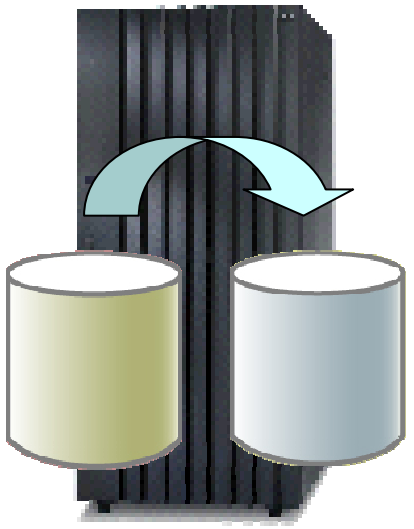

Incremental FlashCopy

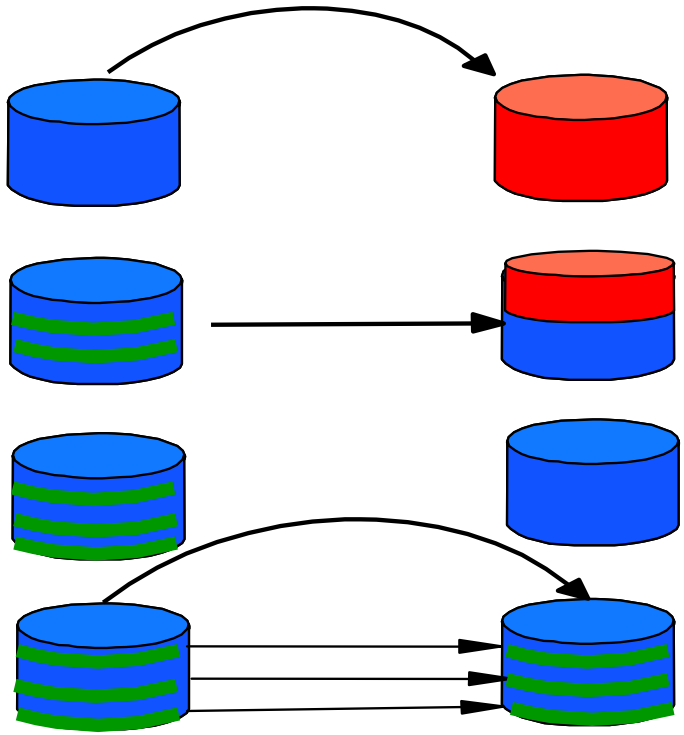

Full volume FlashCopy with incremental

After Flash Copy, Background Copy partially complete - update on source

Background Copy Complete

Next increment, only changed tracks copied Incremental

Inband FlashCopy (continued)

- Initiates FlashCopy from PPRC Secondary
- No need for UCB for FlashCopy volumes

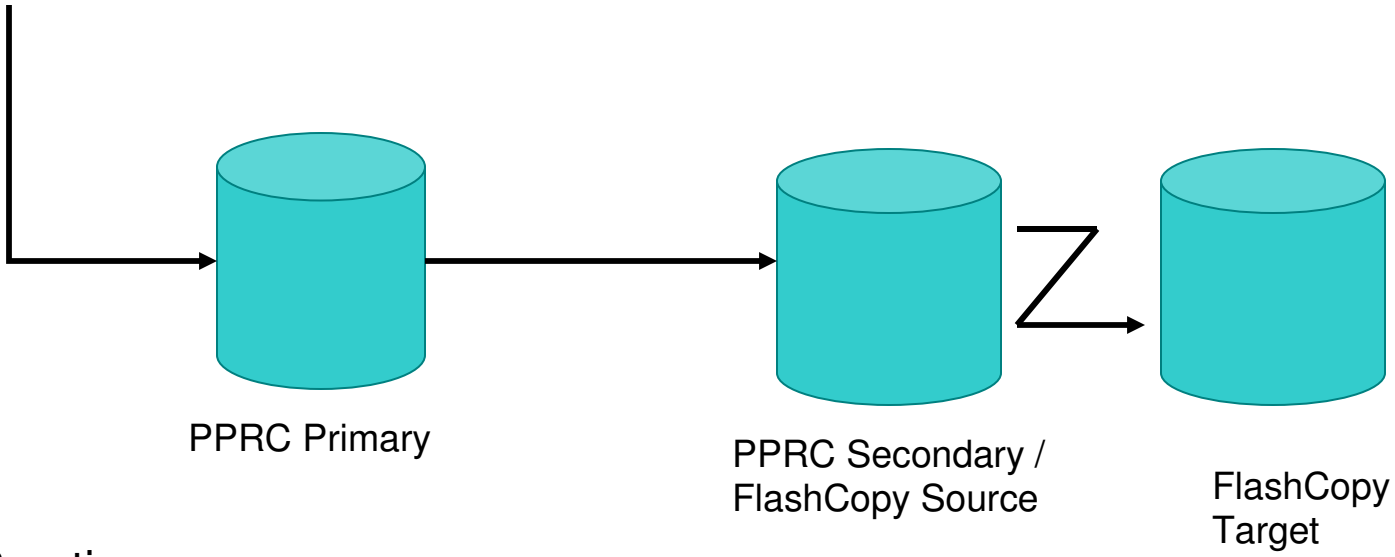

- Practice copy
- Safety copy
- Full volume only

FlashCopy Fast Reverse Restore

FlashCopy relationship exists, target T0
Stop updates to the A volume

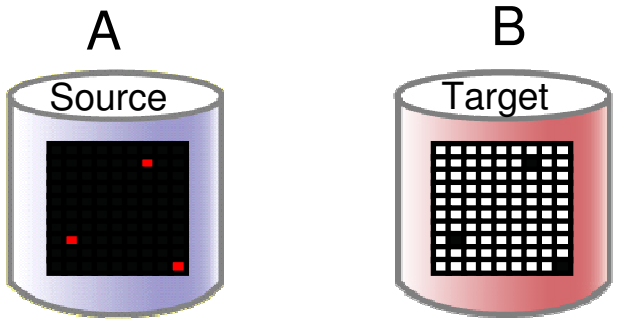

Perform a Fast Reverse Restore B>A to 'reset' A back to T0

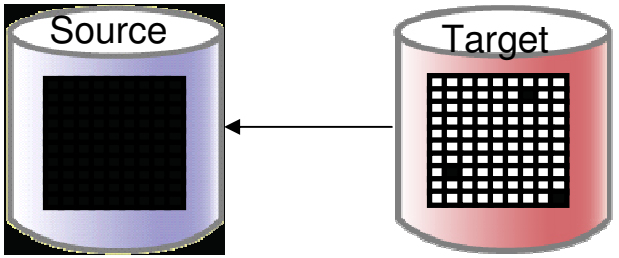

Once the background copy B>A is complete, A is now T0
Can FlashCopy A back to B

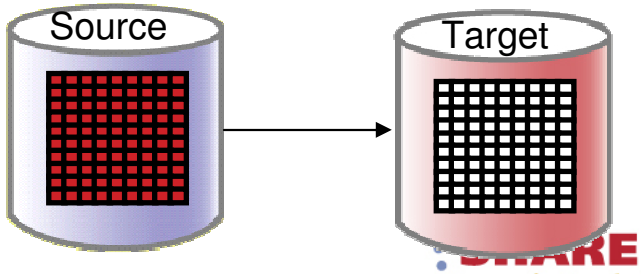

Remote Pair FlashCopy (PPRC Preserve Mirror)

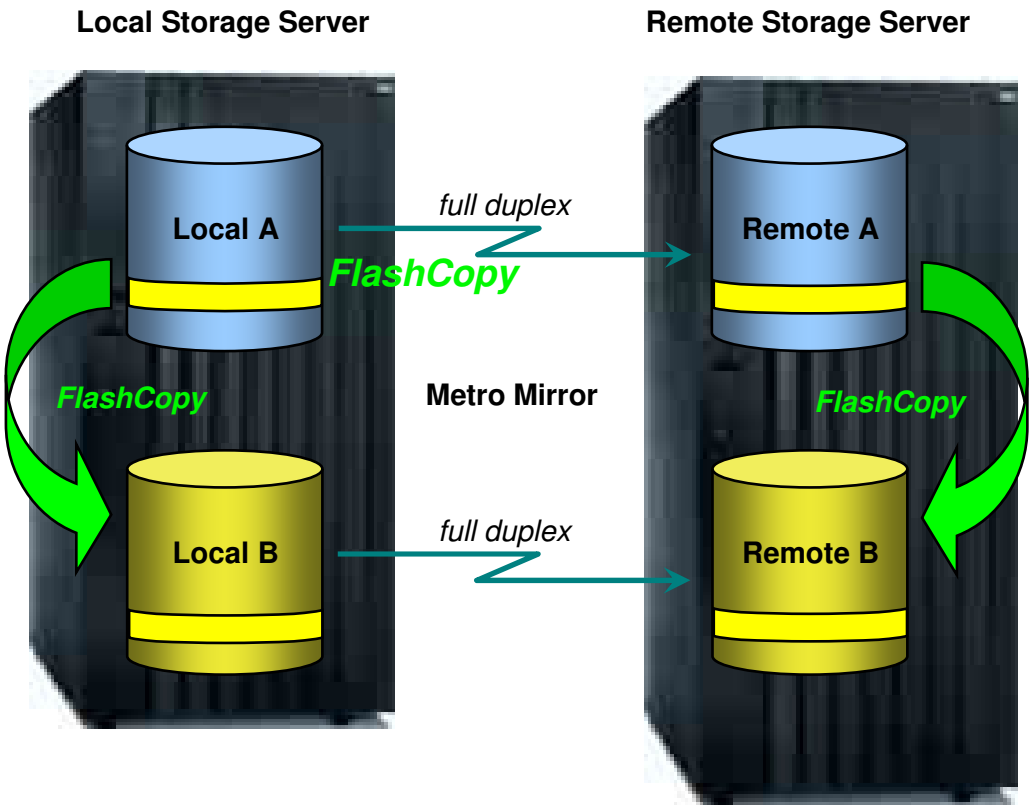

Metro Mirror and FlashCopy

MM 'Practice' Copy

- FlashCopy with 'FREEZE' option for consistency
  - Data dependency consistent (power off consistency)
- Can use Inband or direct addressability
- Can use Incremental to reduce tracks to be copied

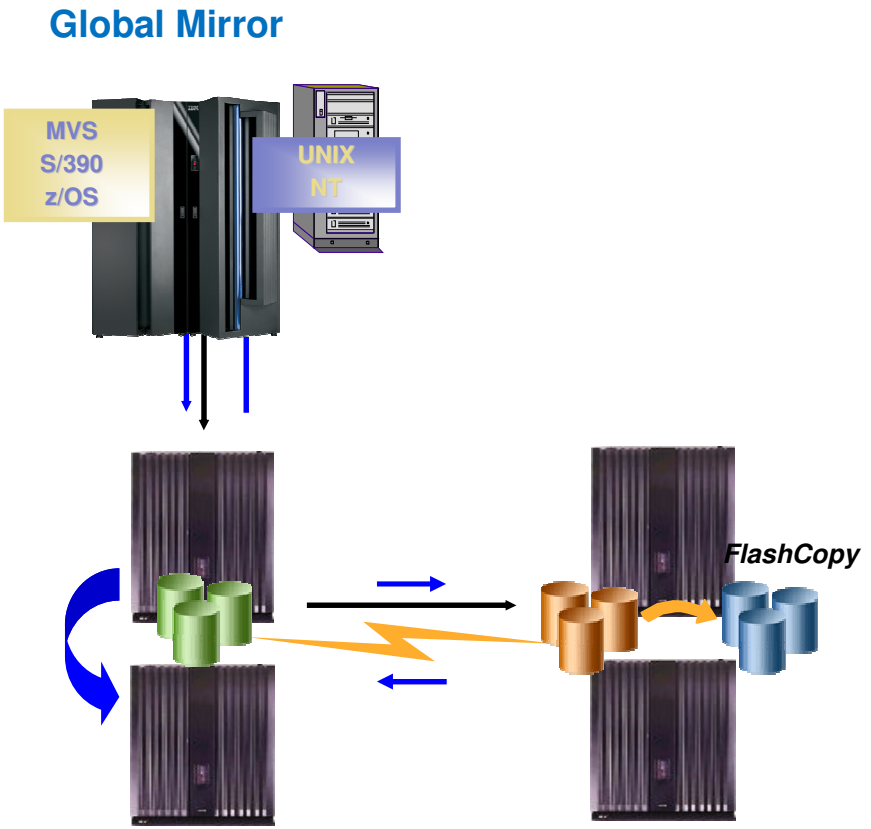

Global Mirror and FlashCopy

GM 'Practice' Copy

- Global Mirror Pause with Consistency
  - Session will pause, and suspend all pairs
- Secondaries are consistent
- Perform FlashCopy at secondary site
- Resume session
  - Pairs will be resync'ed and session resumed
- Can use Incremental to reduce tracks to be copied

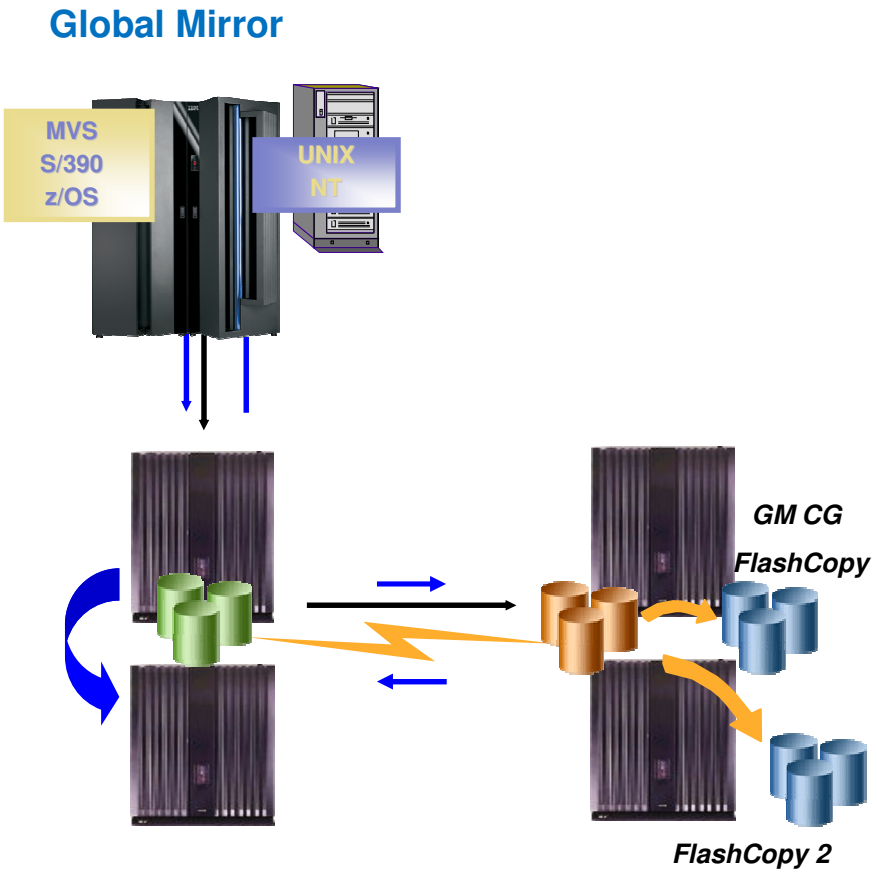

Solutions / Configurations

What are customers doing today ?

Continuous Availability of Data within a Data Center	Continuous Availability / Disaster Recovery within a Metropolitan Region	Disaster Recovery at Extended Distance	Continuous Availability Regionally and Disaster Recovery Extended Distance
Single Data Center Applications remain active Continuous access to data in the event of a storage subsystem outage	**Two Data Centers** Systems remain active Multi-site workloads can withstand site and/or storage failures	**Two Data Centers** Rapid Systems Disaster Recovery with "seconds" of Data Loss Disaster recovery for out of region interruptions	**Three Data Centers** High availability for site disasters Disaster recovery for regional disasters
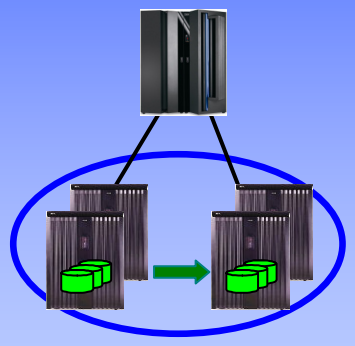	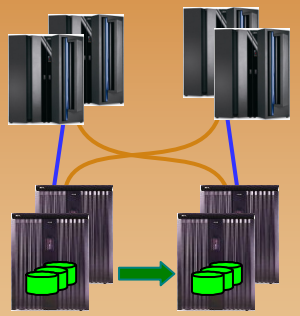	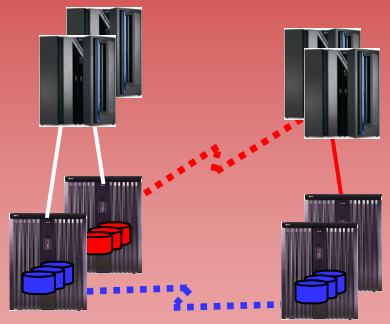	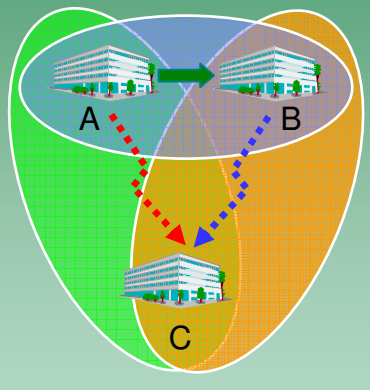
HyperSwap RPO=0 & RTO=0	**Metro Mirror** RPO=0 & RTO<1 hr	**Global Mirror** RPO secs & RTO <1 hr	**Metro/Global Mirror**

4-site topology with Metro Global Mirror

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

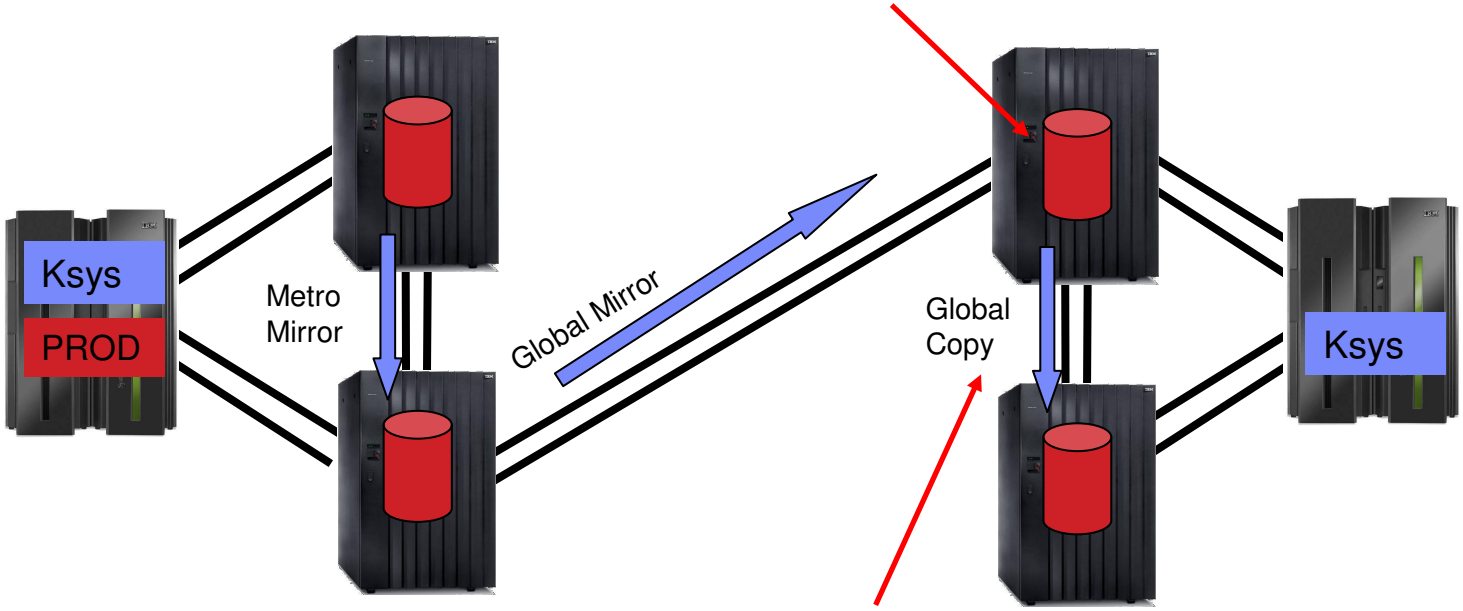

Global Copy in secondary site converted to Metro Mirror in case of disaster or planned site switch

Additional Information

- **Questions?**