GDPS® Overview

The IBM e-Business Availability Solution

gdps@us.ibm.com

Trademarks

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

AIX*
CICS*
DB2*
e-business logo*
Enterprise Storage Server*
ESCON*
FICON
FlashCopy*
GDPS*
HyperSwap
IBM*
IBM eServer*
IBM logo*
NetView*
OS/390*
Parallel Sysplex*
S/390*
Sysplex Timer*
Tivoli*
TotalStorage*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
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.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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**
  - Business Continuity Objectives
  - Need for Data Consistency
  - GDPS overview
- **Continuous Availability (CA) of Data within a Single Site**
  - Disk Maintenance and Disk Failures with HyperSwap
- **Metropolitan Distance CA/Disaster Recovery (D/R) Solution (2 sites)**
  - Configuration Options
  - Planned and Unplanned Site Reconfiguration with HyperSwap
  - Open LUN Management
- **Metropolitan Distance D/R Solution (2 sites)**
  - GDPS/PPRC Storage Manager
- **Unlimited Distance D/R Solution (2 sites)**
  - GDPS/XRC (Unlimited Distance)
- **CA/DR Solution (3 sites)**
  - z/OS® data only
- **Summary**
  - Reference Customer Experiences
  - Future Cross Platform CA/DR Solutions

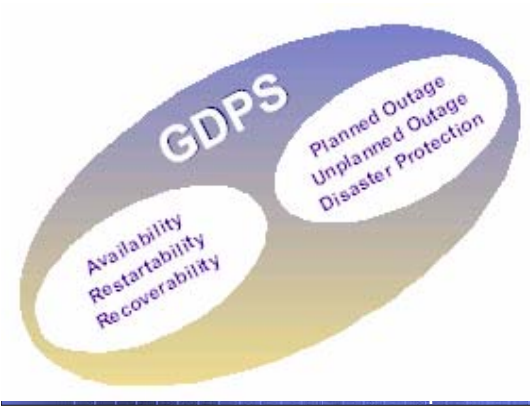

Business Continuity Overview

- ✓Business Continuity Objectives
- ✓Tiers of Disaster Recovery
- ✓Synchronous vs. Asynchronous
- ✓Need for Time Consistency
- ✓What is GDPS

Business Continuance Objectives

- **Determine your Objectives for Business Continuance (by application)**
  - *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 recreate?**
  - *Network Recovery Objective (NRO)*
    - **how long to switch over network?**

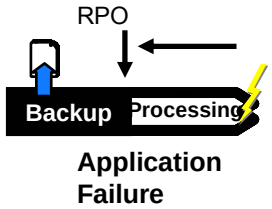

- **Determine cost / recovery time curve**
  - **If I spend a little more, how much *faster* is Disaster Recovery?**
  - **If I spend a little less, how much *slower* is Disaster Recovery?**

Determining the cost vs. RTO recovery curve is the key to selecting proper solution(s)

Tiers of Disaster Recovery: Level Setting GDPS

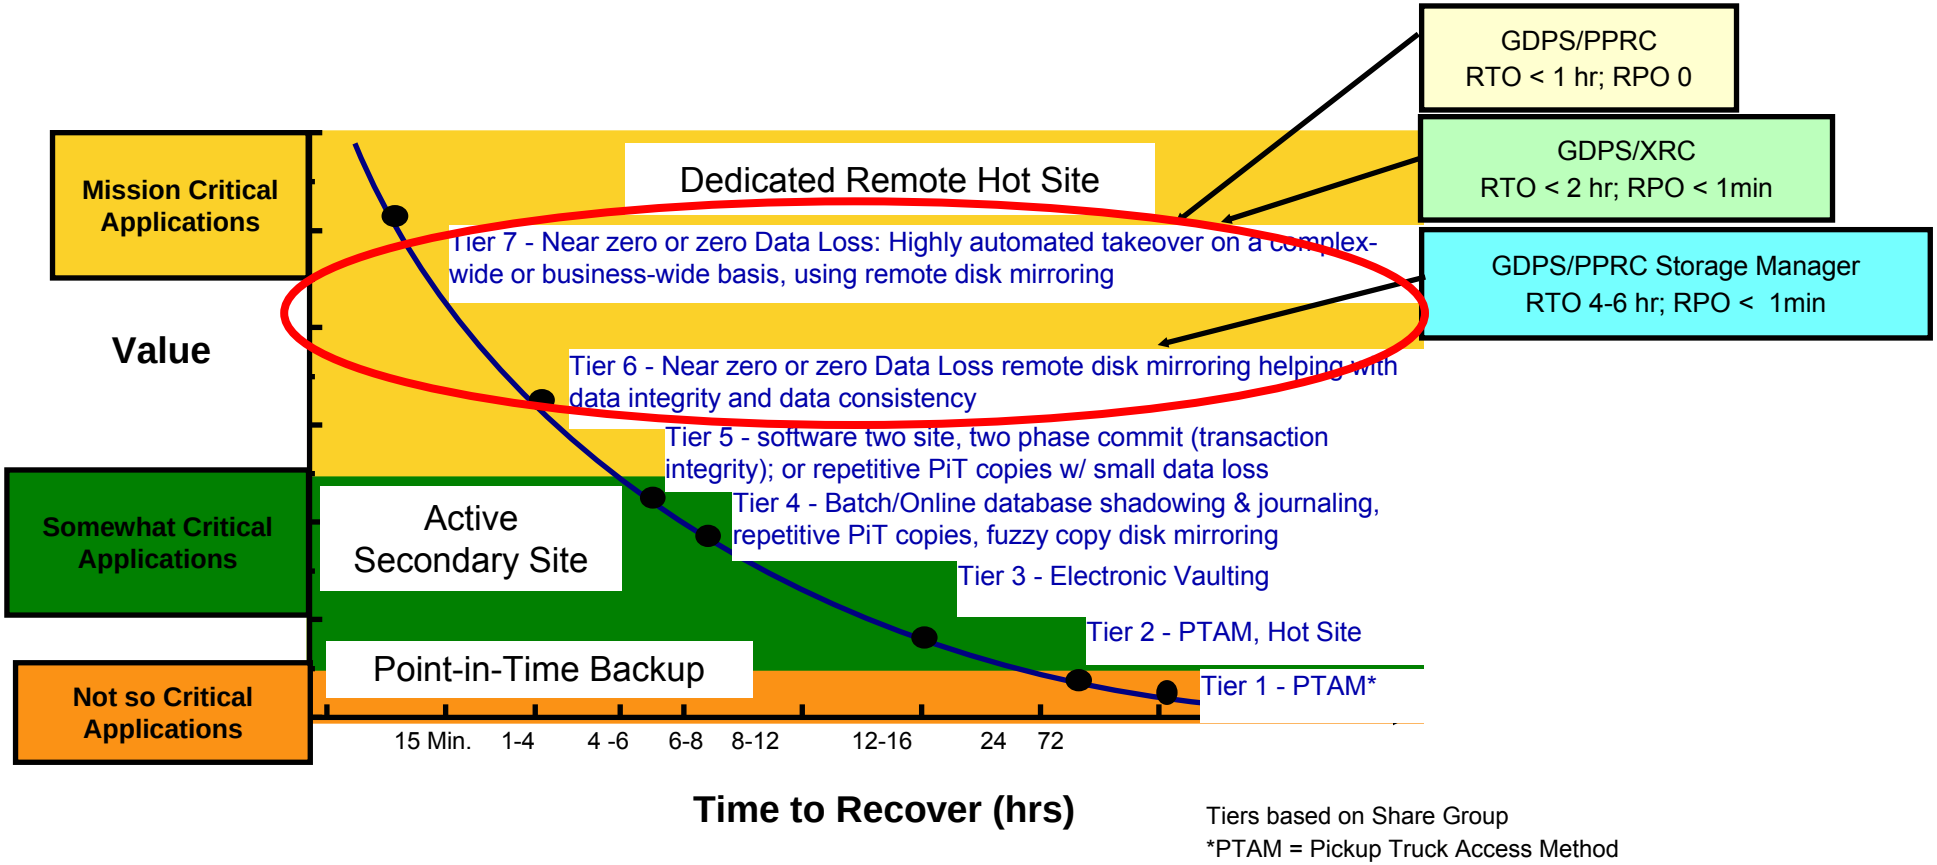

Tiers based on Share Group
*PTAM = Pickup Truck Access Method

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

Business Impact Analysis - Synchronous vs. Asynchronous

Business Impact Analysis

- **Maximum acceptable *response time impact***
- **Maximum acceptable *transaction loss by business process* (RPO)**
- **Distance between production and recovery sites**

- **SYNCHRONOUS** remote copy:
- Use when ***response time impact*** is acceptable
- Use when distance is short
- Use when no data loss is ***the objective***
- ***Often best choice for fastest recovery***
- ***Tradeoff:***
  - ***Meet goal of No data loss & potential CA vs***
  - ***Application impact & short distance***

- **ASYNCHRONOUS** remote copy:
- Use when ***smallest possible impact*** to primary site performance is ***required***
- Use when unlimited distance is ***the objective***
- ***Use when minimal loss of data is acceptable***
- ***Tradeoff:***
  - ***Negligible appl. impact & unlimited distance vs***
  - ***Minimal data loss***

Today's Business Continuity Objectives Demand Rapid Database Availability

- Achieve Application and Database Restart
  - Consistent, repeatable, fast
  - Database Restart: To start a database application following an outage without having to restore the database
    - This is a process measured in minutes
- Avoid Application and Database Recovery
  - Unpredictable recovery time, usually very long and very labor intensive
  - Database Recovery:
    - Restore last set of Image Copy tapes and apply log changes to bring database up to point of failure
    - This is a process measured in hours or even days

What is GDPS?

- **Automation that manages application and data availability in and across sites**
  - **Monitors systems, disk & tape subsystems**
  - **Builds on (multi-site) Sysplex and data mirroring technologies**
  - **Manages planned and unplanned exception conditions**
    - **System maintenance / failure**
    - **Site maintenance / failure**
- **User interface through**
  - **Panels - status and planned actions**
  - **Scripts - planned and unplanned actions**

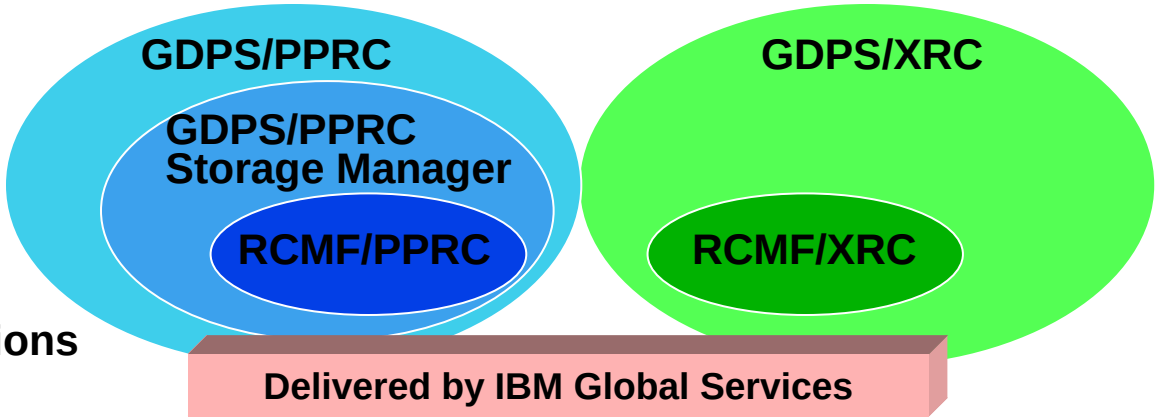

Designed for Continuous Application & Data Availability
Single point of control
Delivered through IBM Services

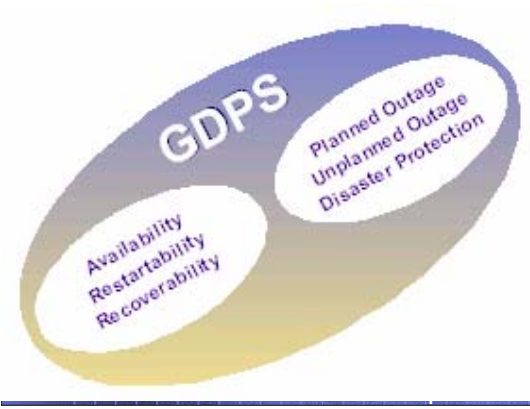

Continuous Availability of data within a single site

- ✓Capacity Backup (CBU) technology
- ✓HyperSwap technology
- ✓Unplanned and Planned disk reconfiguration with HyperSwap

Capacity BackUp Upgrade (CBU)

- **Who Needs It?**
  - Any business with a requirement for increased availability or Disaster Recovery
- **What Is It?**
  - Provides the ability to nondisruptively increment capacity temporarily, when capacity is lost elsewhere in the enterprise
  - Dual Microcode Loads
    - Provide two machine configurations in one box
  - Take advantage of "spare" PUs
  - Significant cost savings possible
    - Standby MIPS cost can be eliminated
    - IBM Software license charges on standby MIPS can be eliminated
  - Configure memory and channels to support production workload
- **How Can I Use It?**
  - Adjacent machines in the same location
  - Multiple images in the same Parallel Sysplex® cluster
  - Backup/Recovery site

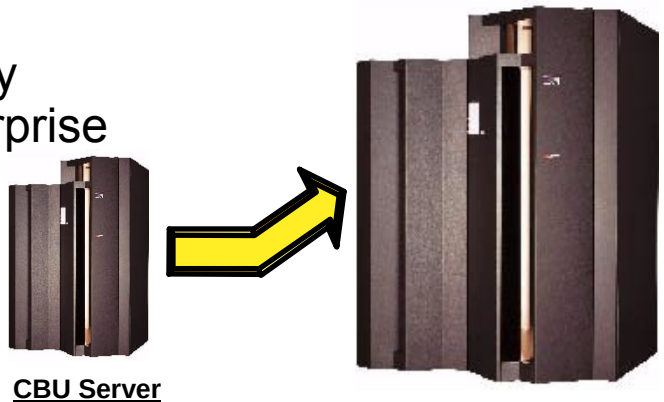

CBU Server

Production Server

I0140

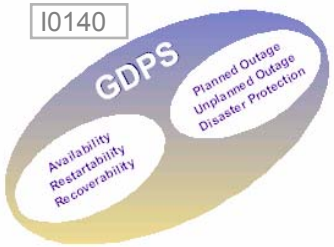

GDPS/PPRC HyperSwap – the Technology

Parallel Sysplex designed to provide Continuous Availability of systems, servers and applications within a single site
(degree of availability depends on exploitation of Parallel Sysplex)

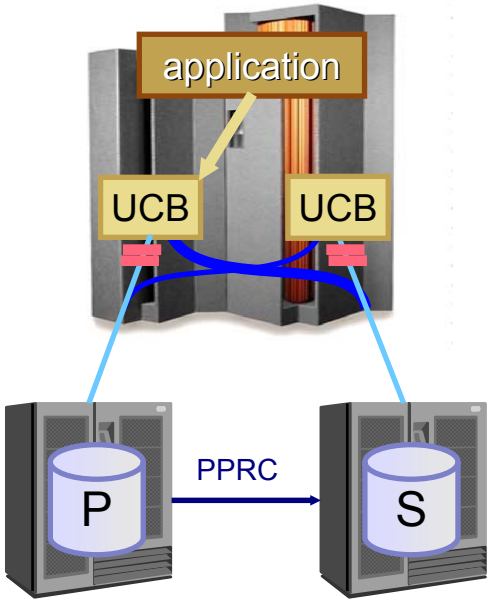

- Extends Parallel Sysplex availability to disk subsystems
- Masks primary disk subsystem failures by transparently switching to use secondary disks
- Provides ability to perform disk maintenance and planned site maintenance without requiring applications to be quiesced
- Substitutes PPRC secondary for 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 using same device addresses

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

GDPS/PPRC

Unplanned Disk Reconfiguration with HyperSwap

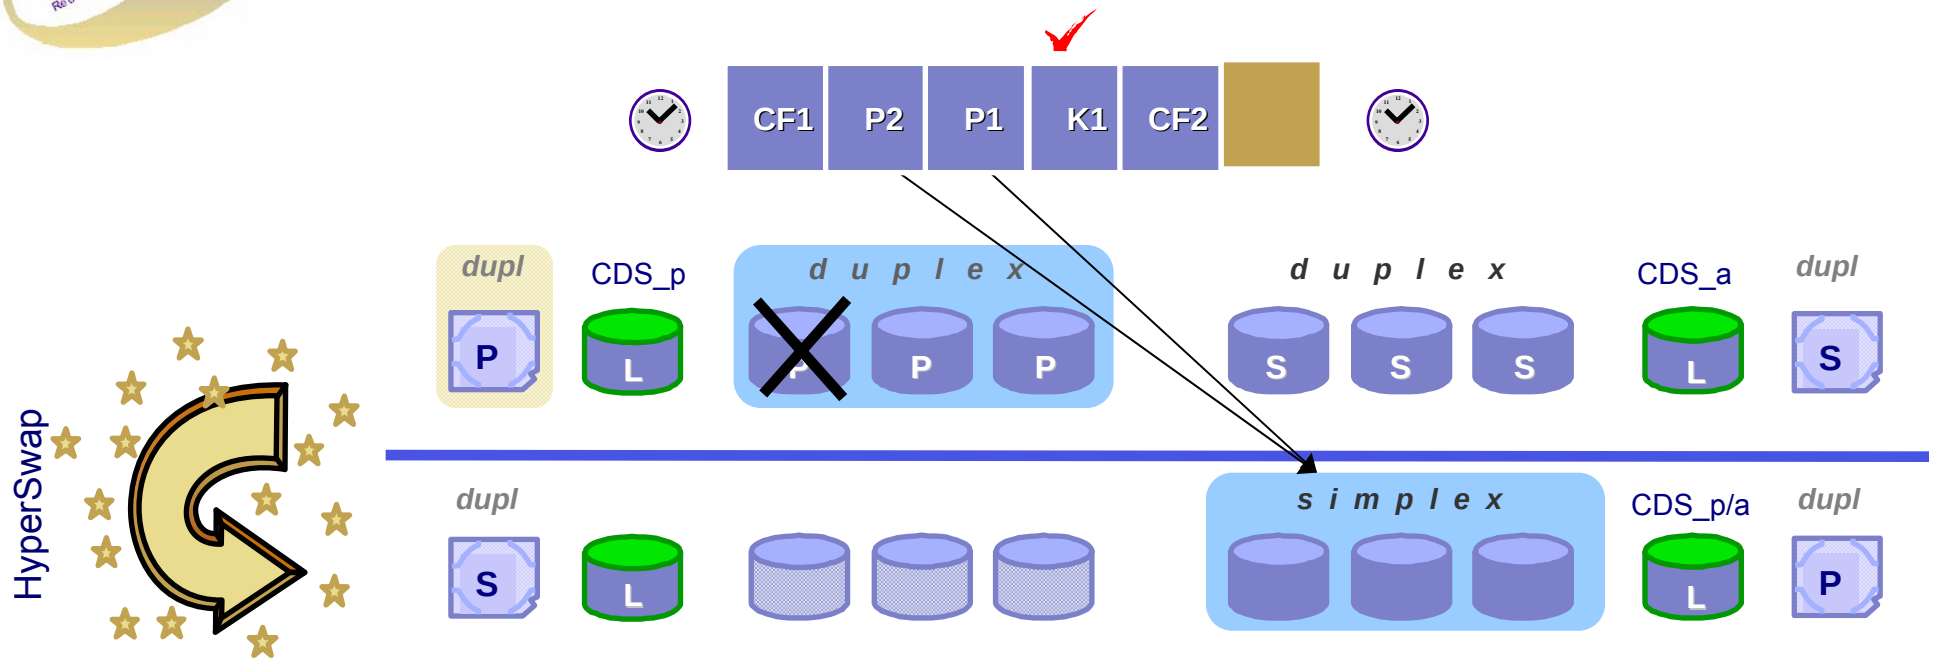

Without HyperSwap

FREEZE

TAKEOVER PROMPT

systems quiesced
(RESET)
if Freeze & Stop/Cond

Decision Time

TAKEOVER STARTED

remove systems from Sysplex
recover secondary PPRC volumes
restart systems

30 – 60 Minutes or more

With HyperSwap

25 Seconds!

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

P1, P2, remain active throughout the procedure

GDPS/PPRC

Planned Disk Configuration Maintenance with HyperSwap

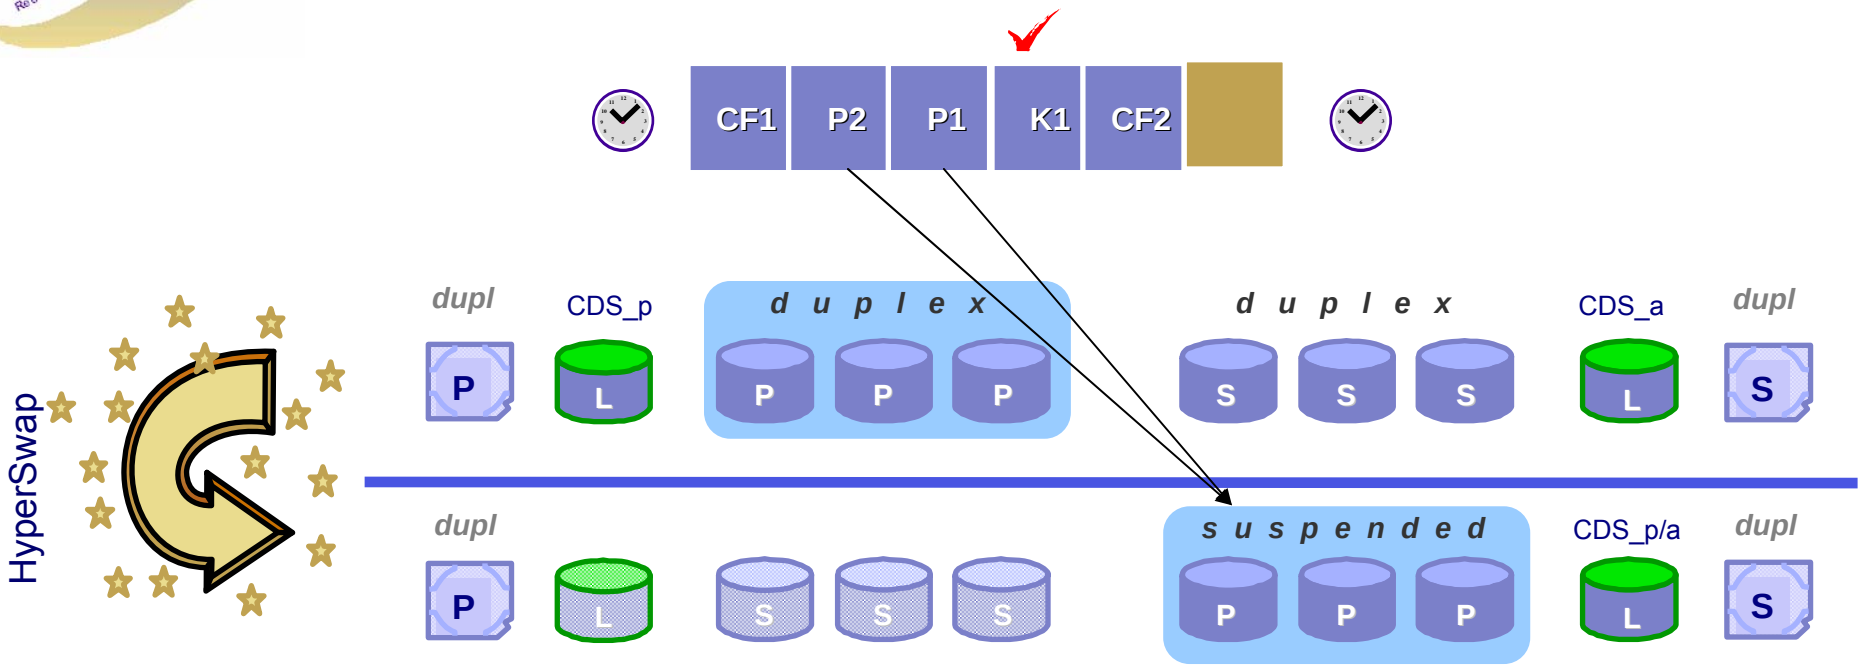

Without HyperSwap

PLANNED ACTION INITIATED

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

1-2 hrs (approx)

With HyperSwap

75 Seconds!

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

P1, P2 remain active throughout the procedure

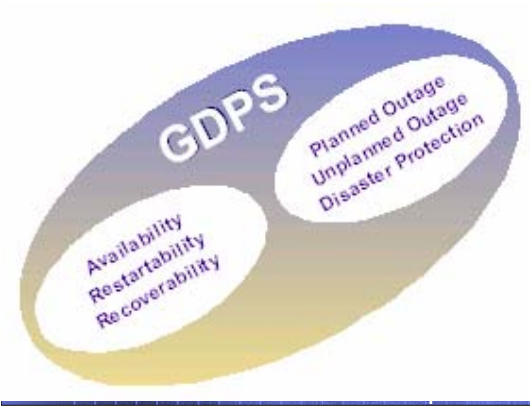

Metropolitan Distance Continuous Availability / Disaster Recovery Solution (2 sites)

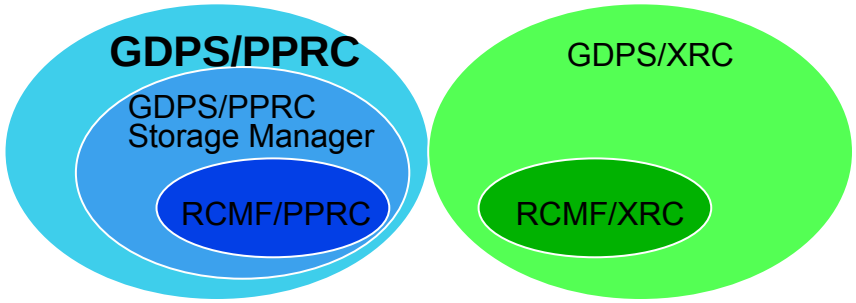

- ✓GDPS/PPRC
- ✓Configurations
- ✓Management of Open Systems LUNs

What is GDPS/PPRC? (Metro Mirror)

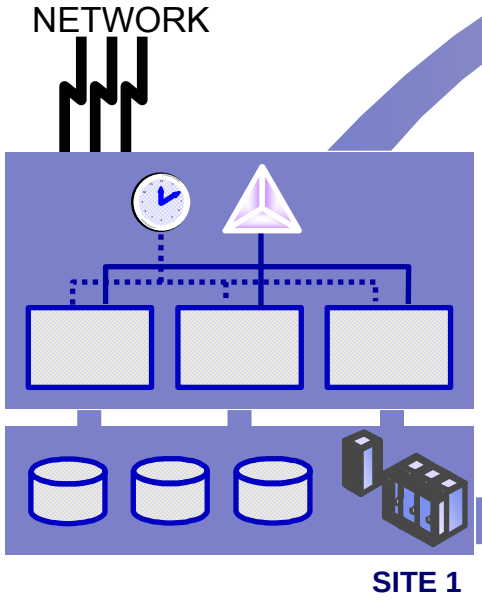

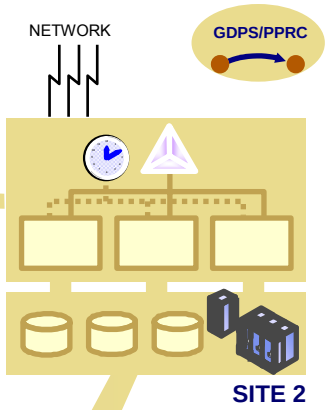

- **Multi-site base or Parallel Sysplex environment**
  - Sites separated by up to 100 km of fiber
- **Remote data mirroring using PPRC**
- **Manages unplanned reconfigurations**
  - z/OS, CF, disk, tape, site
  - Designed to maintain data consistency and integrity across all volumes
  - Supports fast, automated site failover
  - No or limited data loss - (customer business policies)
- **Single point of control for**
  - Standard actions
    - –Stop, Remove, IPL system(s)
  - Parallel Sysplex Configuration management
    - –Couple data set (CDS), Coupling Facility (CF) management
  - User defined script (e.g. Planned Site Switch)
  - PPRC Configuration management

Planned and Unplanned exception conditions

GDPS/PPRC Distance extended to up to 100 km

- ETR links (Sysplex Timer® to server) and ISC3 links (Peer mode) extended from 50 km to 100 km
  - Links between Sysplex Timers CANNOT exceed 40 km
  - DWDMs and/or amplifiers required
- GDPS supports ESS PPRC over FCP (fiber channel) links
  - Protocol used with fiber channel more efficient than ESCON
  - **Lower Total Cost of Ownership (TCO)**
    - Only 2 cross site FCP links / ESS required for most workloads
  - **Better performance**
    - Able to increase distance between sites while maintaining acceptable application performance
    - One protocol exchange vs 2-3 with ESCON

Single Site Workload - Cross-site Sysplex
Near Continuous Availability Configuration

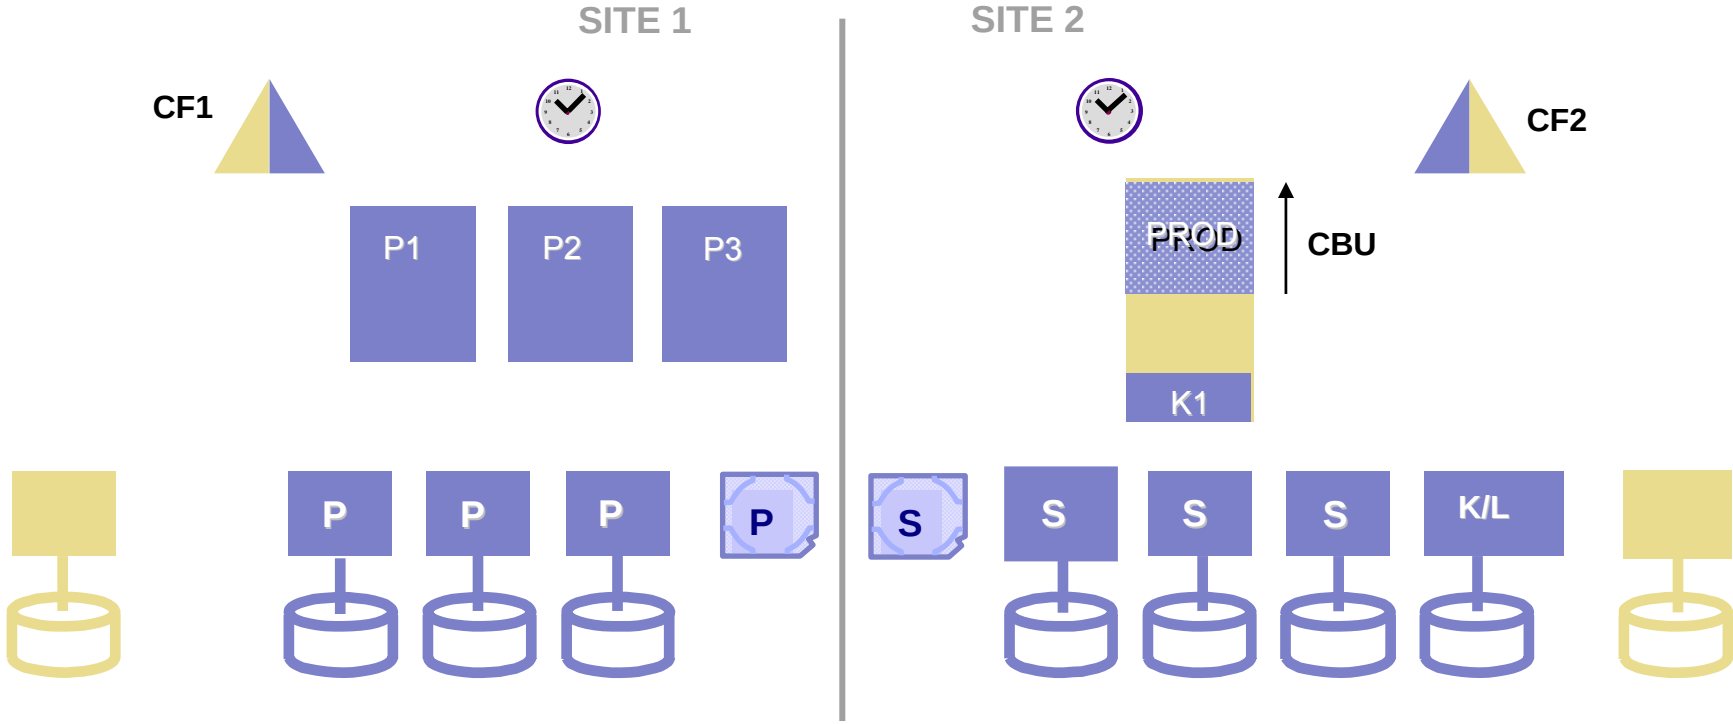

Site recovery by restarting failed system images in Site 2
Designed to support continuous access to data from site 1
Unplanned and planned disk reconfiguration with HyperSwap

Multiple Site Workload - Cross-site Sysplex Continuous Availability Configuration

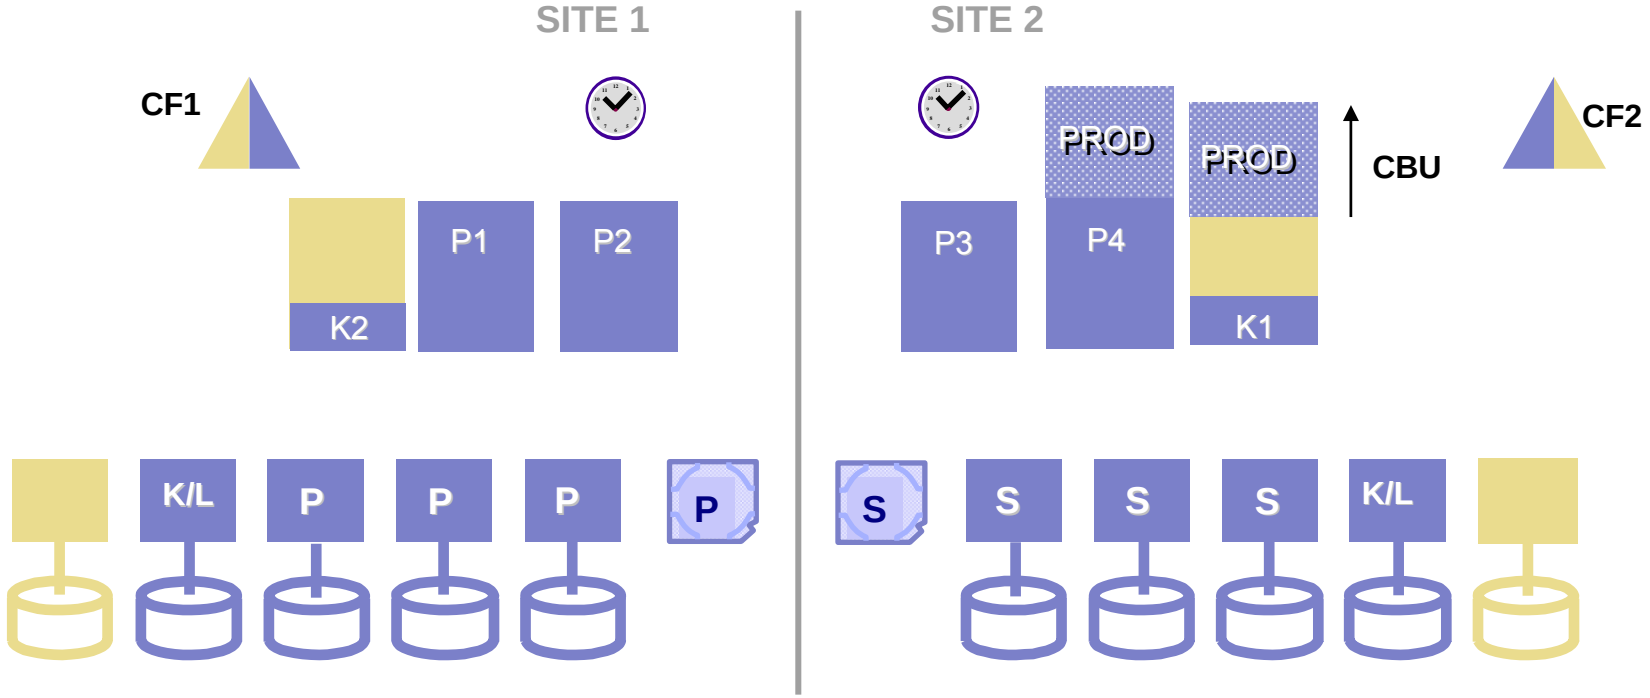

Supports Planned and Unplanned Site and Disk reconfiguration via HyperSwap
Operating systems remain active, applications need to be recycled
Designed to provide continuous access to data from either site

GDPS/PPRC

Sysplex in a single site (aka BRS configuration) PPRC across sites

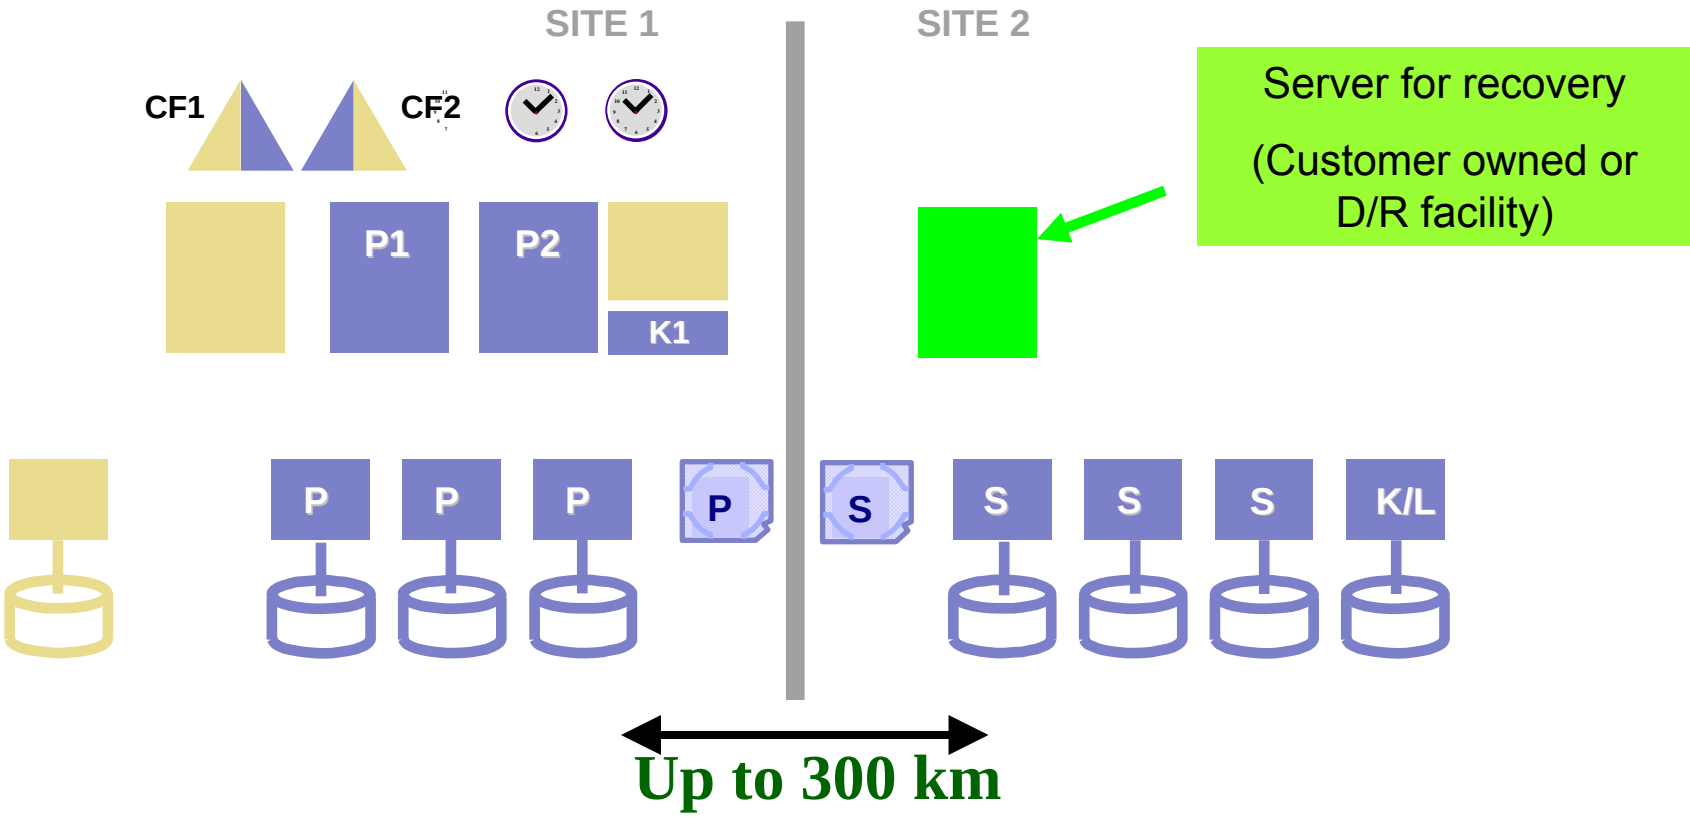

IPL K system in Site 2
GDPS automation restarts Production systems and applications
No Data Loss Full data consistency

New Functions with HyperSwap technology

- **Unplanned reconfigurations**
  - **Production systems can remain active during disk failover**
  - **Site 1 failover**
    - **Applications cloned and exploiting data sharing across 2 sites**
    - **Production systems can remain active**
    - **Workload needs to be restarted**
- **Planned reconfigurations**
  - **Disk Maintenance**
    - **Swap Primary/Secondary disks is nondisruptive to applications**
  - **Site Maintenance without stopping applications**
    - **Applications cloned and exploiting data sharing across 2 sites**
  - **Data Migration**
    - **HyperSwap and Terminate**

Recovery Scenario with PPRC only

DR Team On-Site

- Remote Location: Operations Center/Home
- Remote Connectivity to SITE2 Required

System Recovery (Manual)

- Shutdown Discretionary Workload
- POR CEC / Load Production IODF
- Reconfigure Tape Library
- Reset CFs
- Network

System/Application/Data Validation

- IPL Production zOS LPARs
- Data Consistency
- Application Access/Functioning
- Network Connectivity

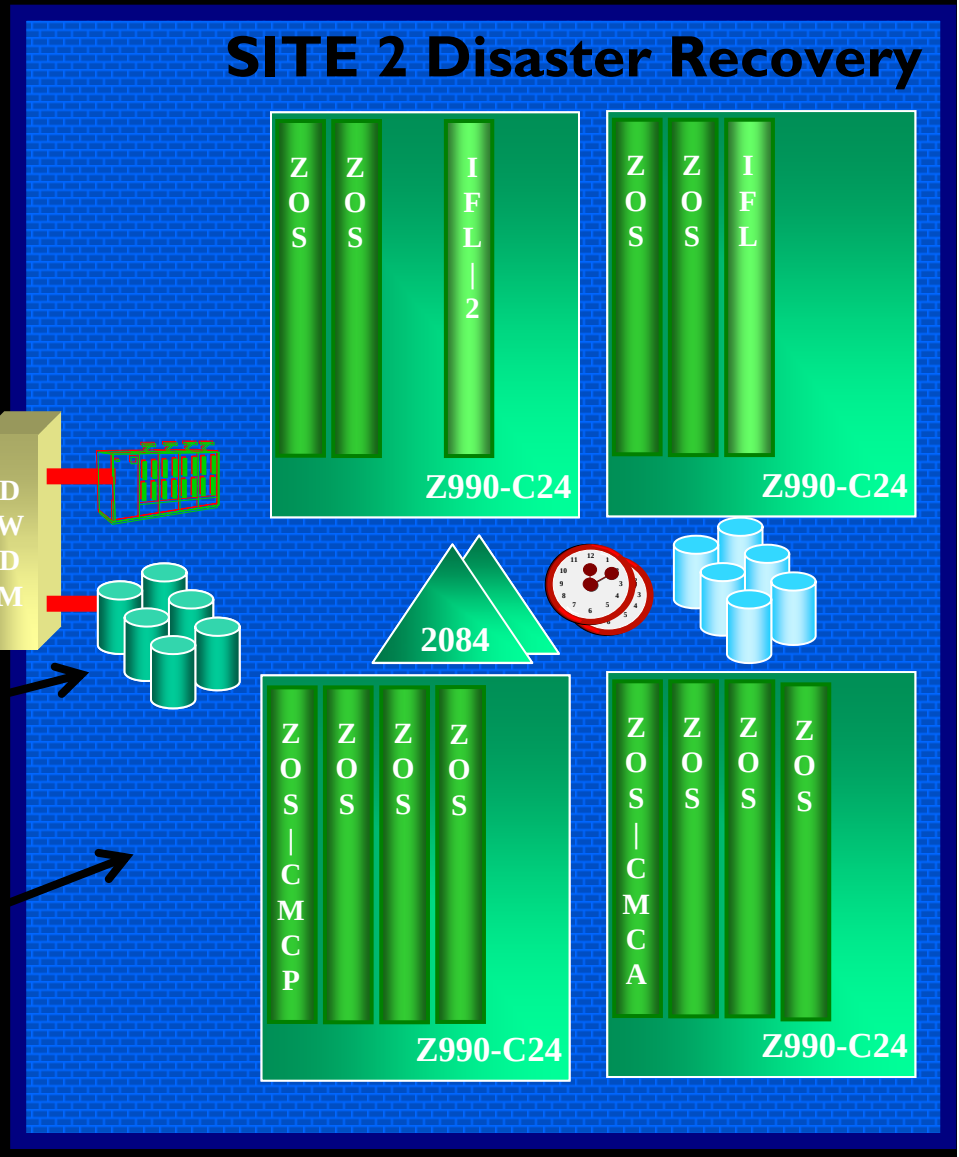

Declare Disaster | Team On Site | Restore Systems | Application/Data Validations

2 Hours

Time → Money

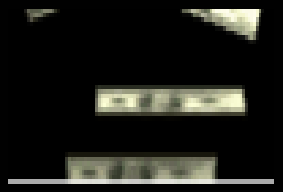

Recovery Scenario with PPRC only

Data Base
Recovery/Restore

DR /
Recovery Team

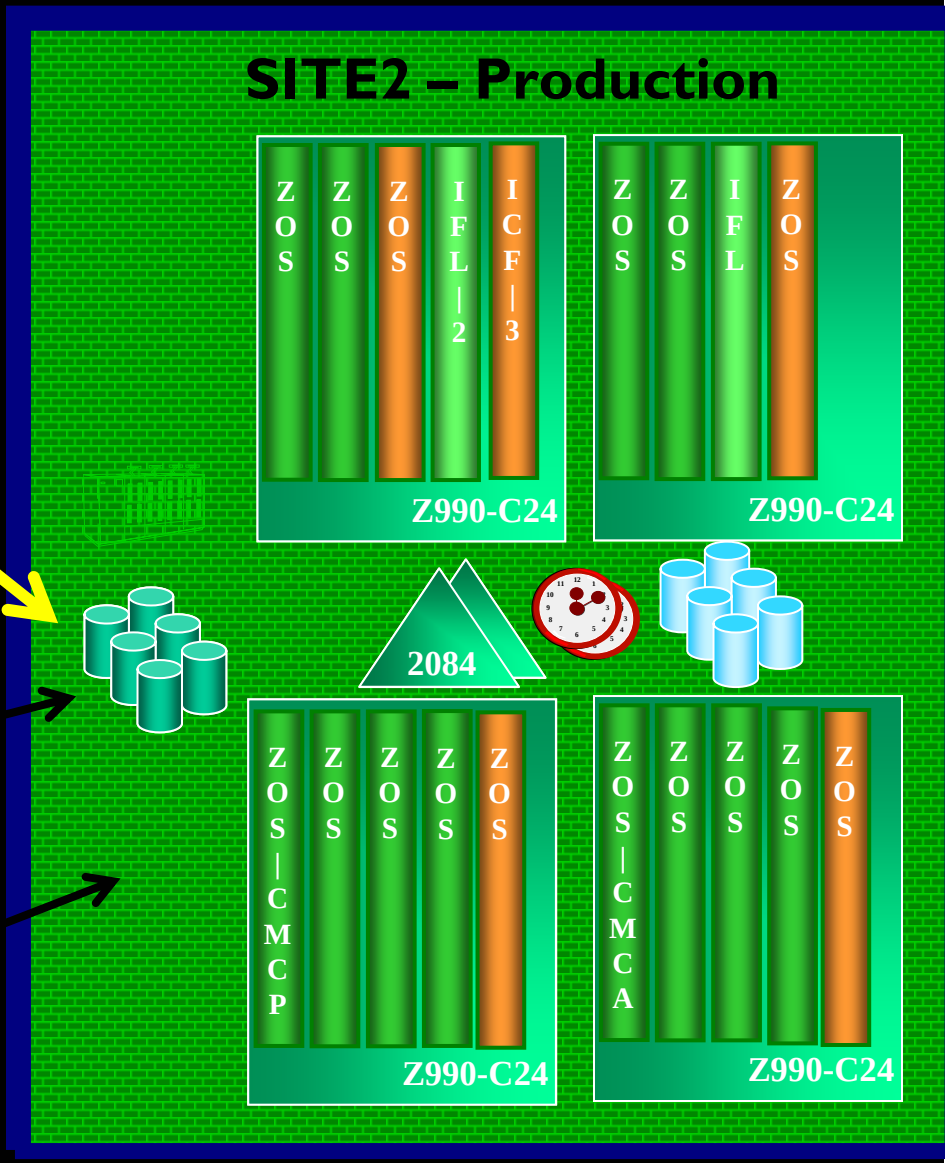

Declare Disaster

Team onSite

Restore Systems

Application/Data Validations

DB Recovery

6 Hours

2 Hours

Time

Money

Enhanced Recovery with GDPS/PPRC

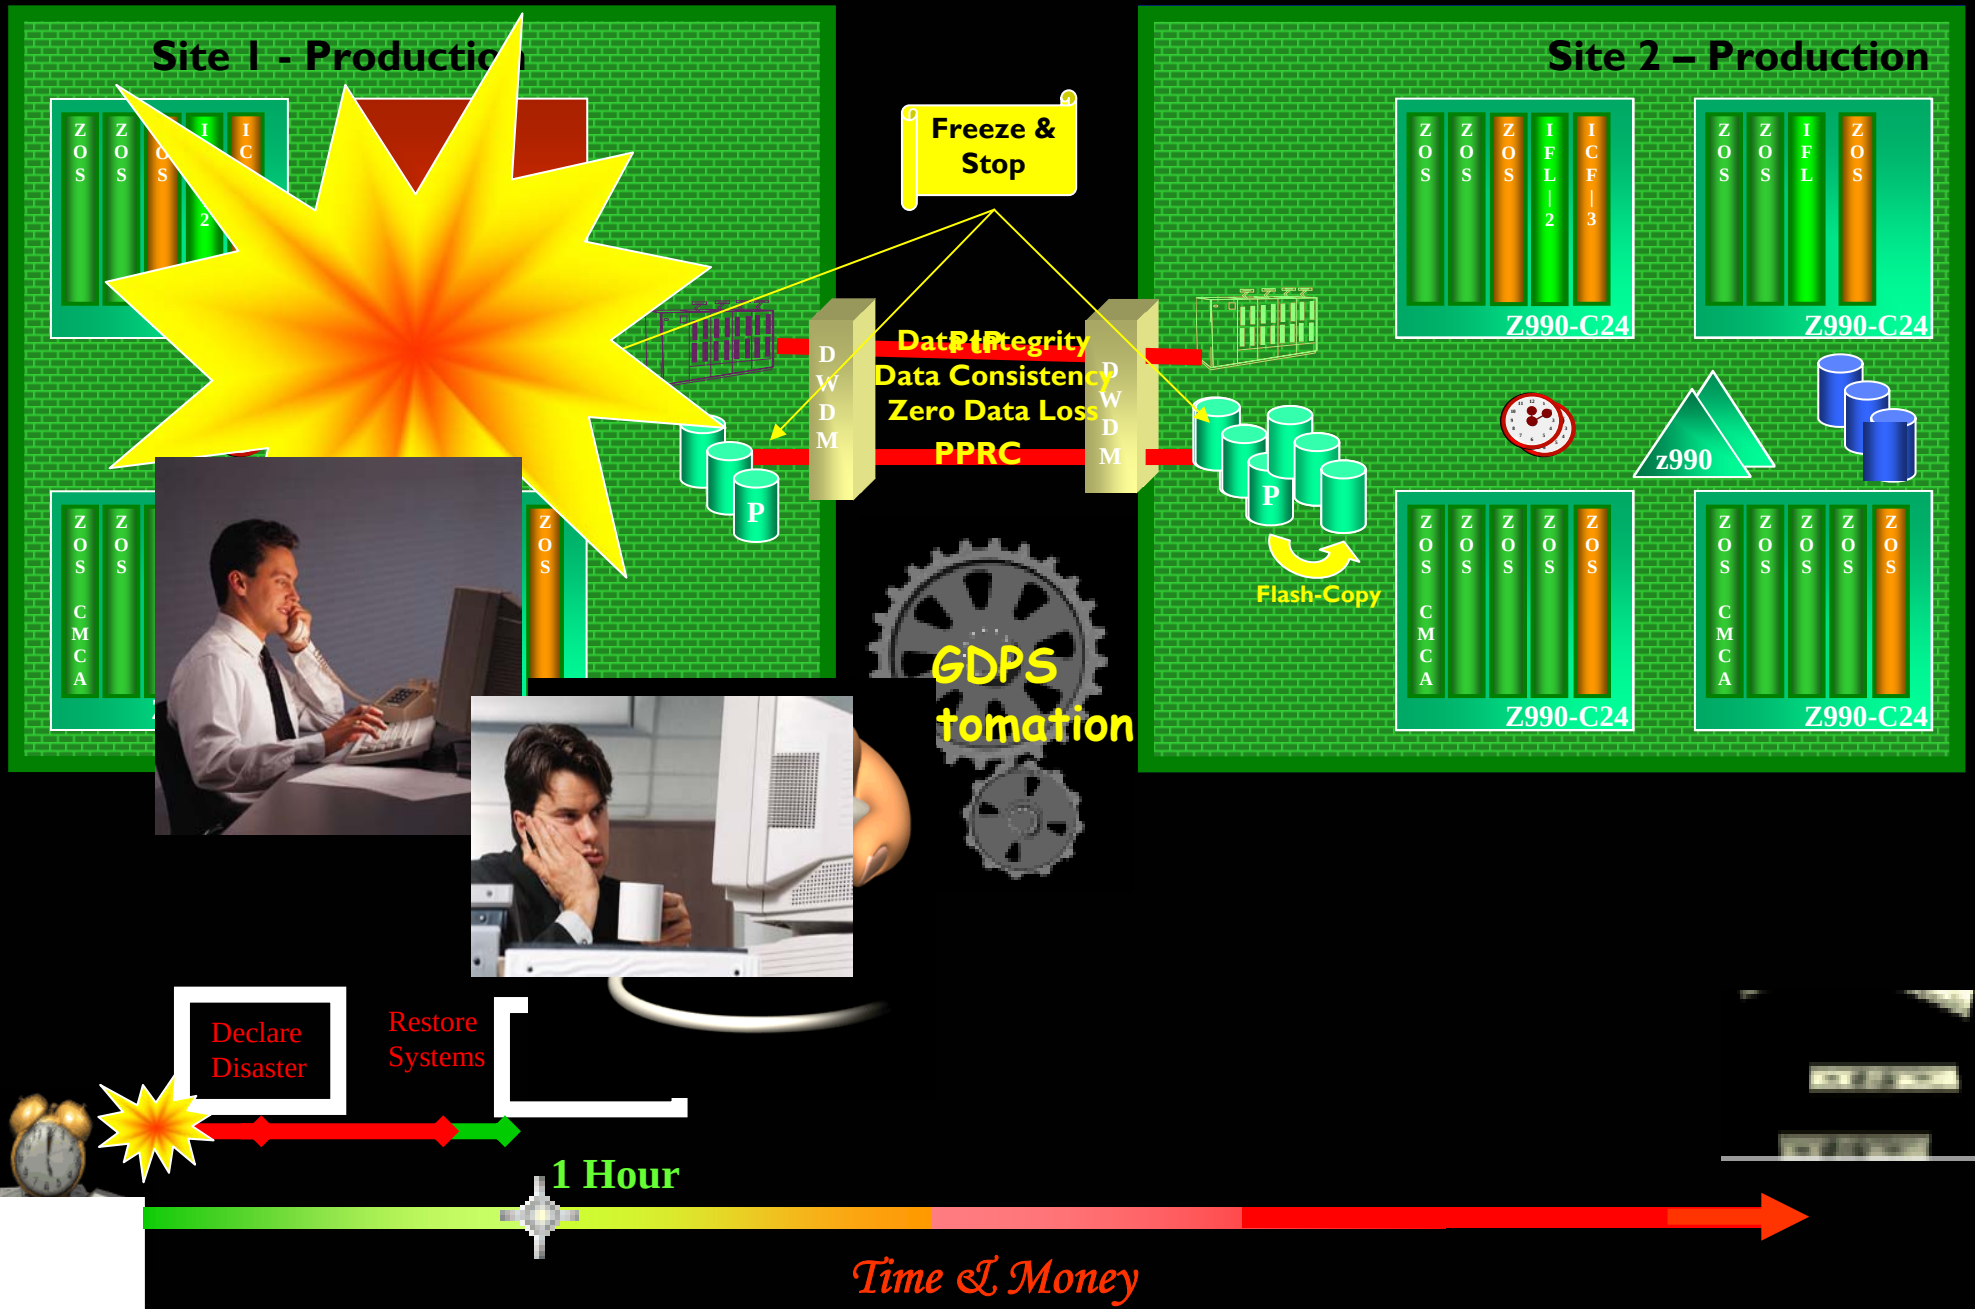

GDPS/PPRC

GDPS/PPRC management of Open Systems LUNs

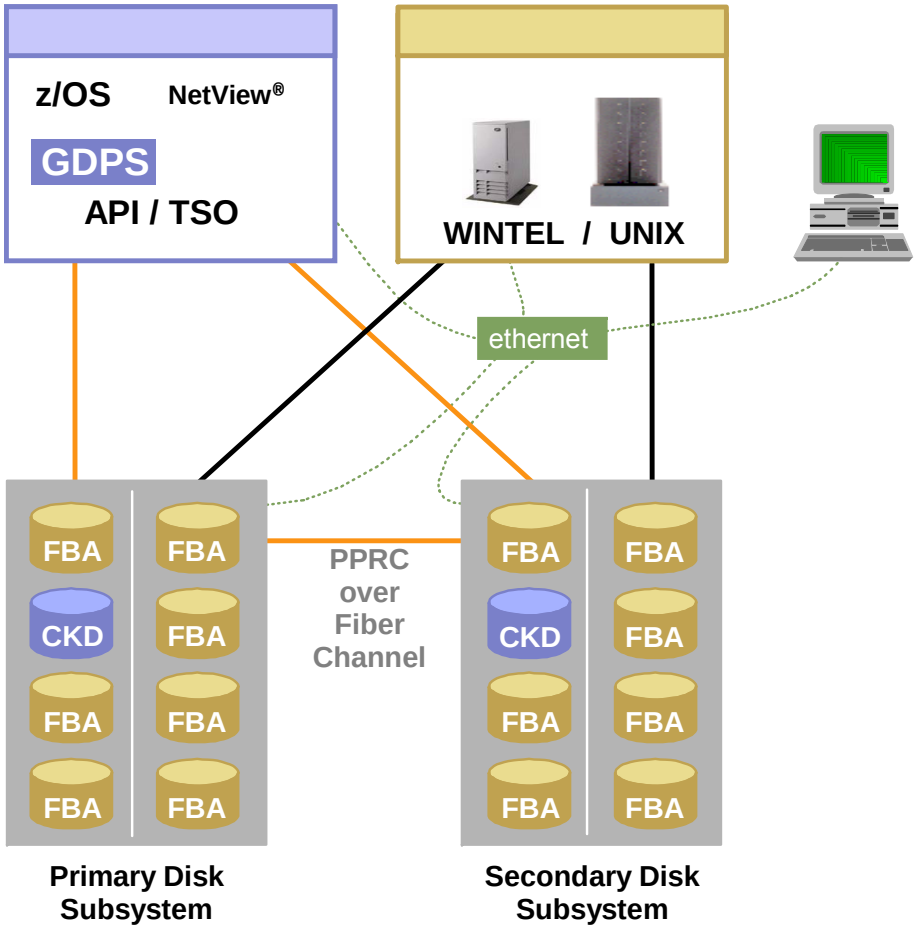

- Extends GDPS/PPRC technology to manage distributed applications across multiple platforms
  - z/OS and open systems data (Unix, NT, Linux)
- GDPS/PPRC running in a z/OS system manages the PPRC status of devices that belong to the other platforms
- Provides data consistency across both z/OS and/or open systems data when failures occur
- Requires
  - Some CKD capacity in disk subsystem
  - PPRC level 4
- Support details
  - Supports x-platform or platform level Freeze
  - FlashCopy not supported for Open
  - No GDPS Code running on Open Systems host - suspend reported through SNMP alert
  - Manual restart of Open systems required

Helps provide enterprise-wide Disaster Recovery with data consistancy!

GDPS/PPRC Multi Platform Resiliency for zSeries

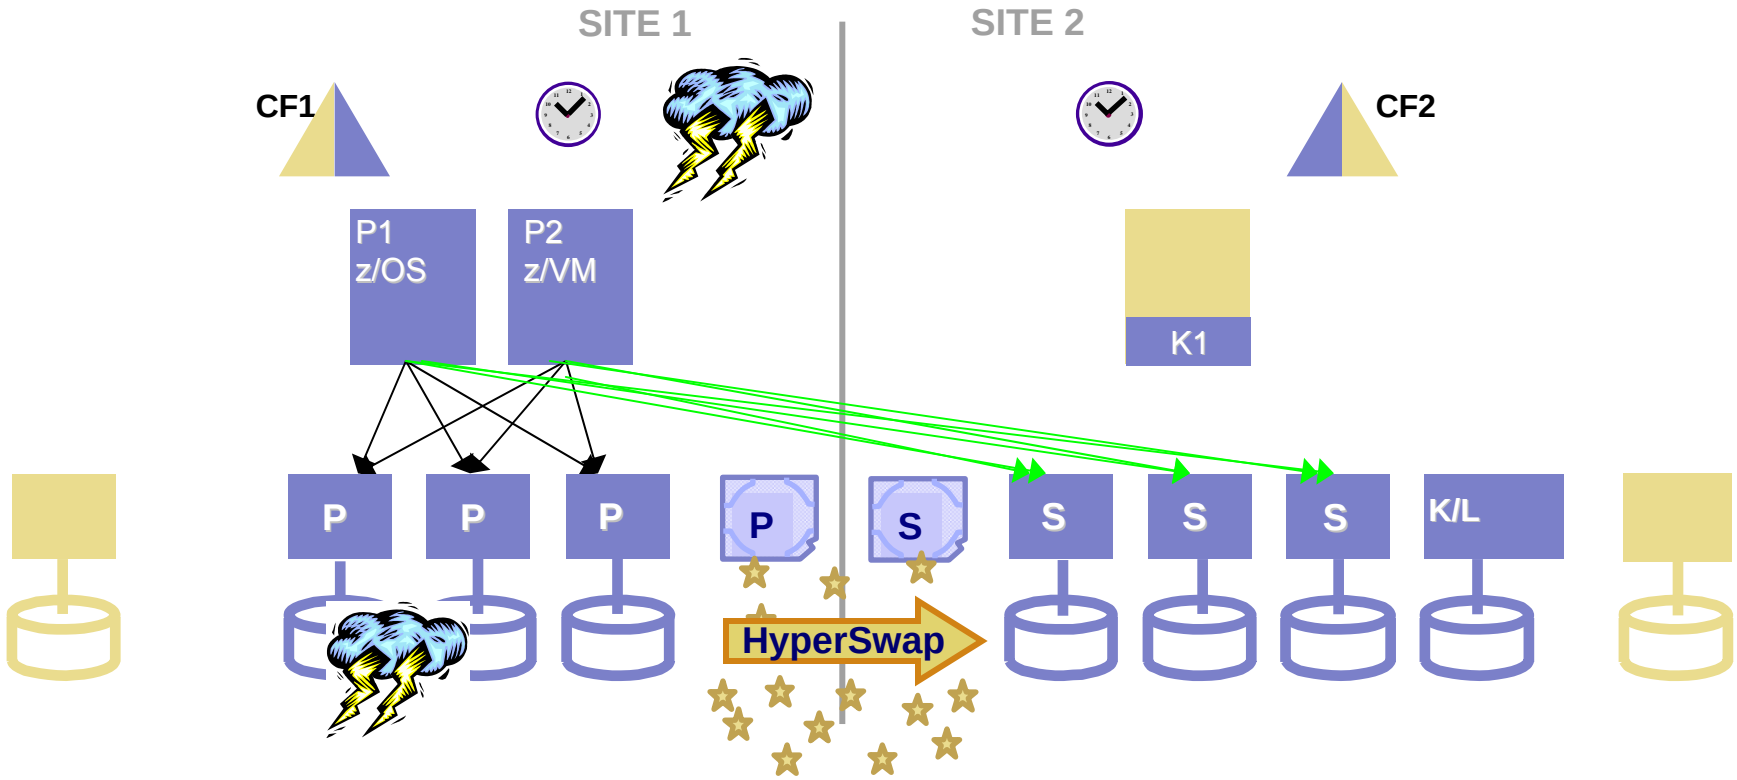

✓**Coordinated near-continuous availability and DR solution for z/OS and Linux guests running under z/VM**
- ***Valuable for customers with distributed applications***
- *SAP application server running on Linux for zSeries*
- *SAP DB sever running on z/OS*

✓**Planned and Unplanned Reconfigurations**

Metropolitan Distance Disaster Recovery (2 sites)

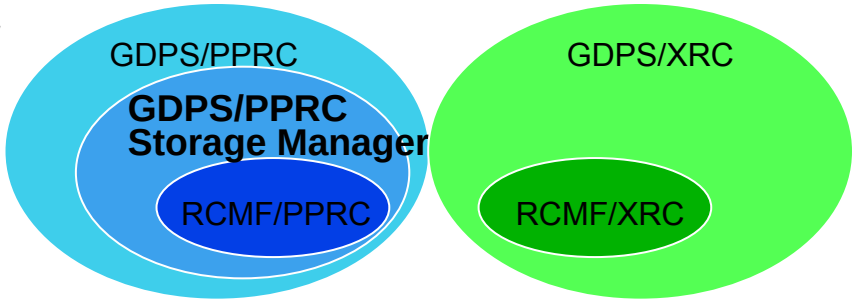

- ✓GDPS/PPRC Storage Manager

GDPS/PPRC Storage Manager

- Benefits
  - **Combines the features of Remote Copy Management with the automation of GDPS**
  - **Helps ensure data consistency with minimal data loss**
  - **Simplifies management of your Remote Copy configuration**
  - **Helps to reduce storage management costs**
  - **Can reduce time required for remote copy implementation**
- Investment protection
  - **Positioned to upgrade to full GDPS/PPRC**

Storage Manager Functional Overview

- Automation shell that manages data availability in and across sites
- Provides Single Point of Control for PPRC configuration management
- Helps ensure data consistency in the event of failures or disaster
  - **Freeze function assures secondary data consistency allowing systems and applications to be restarted in site 2**
  - **Restart has to be invoked by customer**
- FlashCopy support
  - **Auto initiated by GDPS prior to resynchronization**
  - **User initiated**
- User interface through panels
  - **Status and planned actions**
  - **Facilitates Primary/Secondary disk swaps for Planned Disk/Site Maintenance**
- Upgradeable to "full" GDPS/PPRC

Lower cost offering
Single point of control

Site 1 Failure

Site 1 Failure

- Freeze secondary disks (GDPS automation)
- IPL K1 in site 2 (if configuration is “Sysplex in a single site”)
  - Customer task
- Recover secondary disks (GDPS panels)
- Invoke CBU, reIPL production systems, restart applications
  - Customer task

Unlimited Distance Disaster Recovery (2 sites)

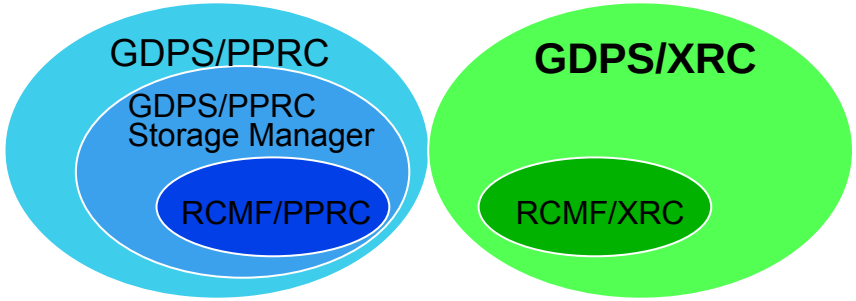

✓GDPS/XRC

What is GDPS/XRC?
(Global Mirror for zSeries)

- **Productivity tool that integrates management of XRC and FlashCopy**
  - Full-screen interface
  - Invoke scripted procedures from panels or through exit
- **GDPS/XRC runs in the SDM location and interacts with SDM(s)**
  - Manages availability of SDM Sysplex
  - Performs fully automated site failover
- **Single point of control for multiple / coupled Data Movers**

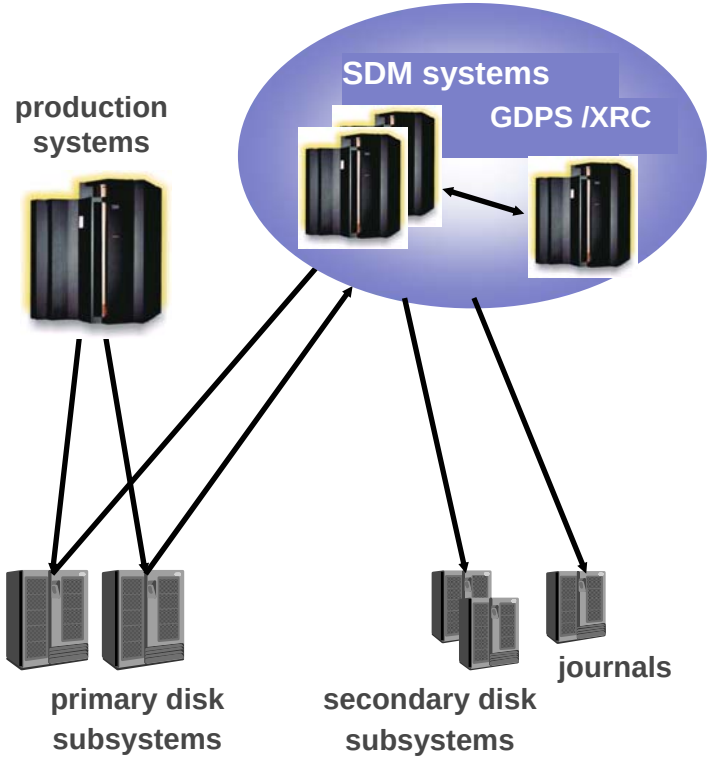

GDPS / XRC Configuration

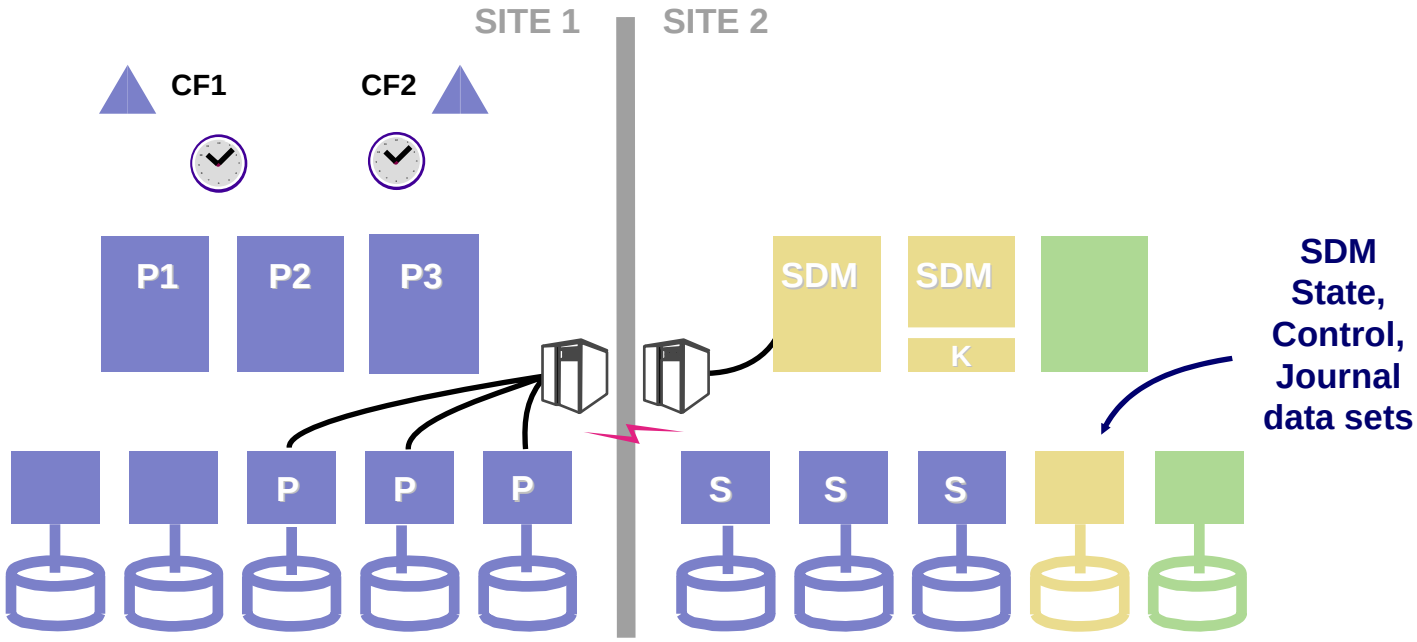

- Production system can be no, Base or Parallel Sysplex environment
- Common time reference required in site 1
- Primary data in site 1, secondary data in site 2
- System Data Mover(s) and Controlling System in site 2
- System Data Mover(s) must run in Base or Parallel Sysplex
- Expendable workload in site 2 and/or CBU capable processors

GDPS/XRC - Primary Site Failure

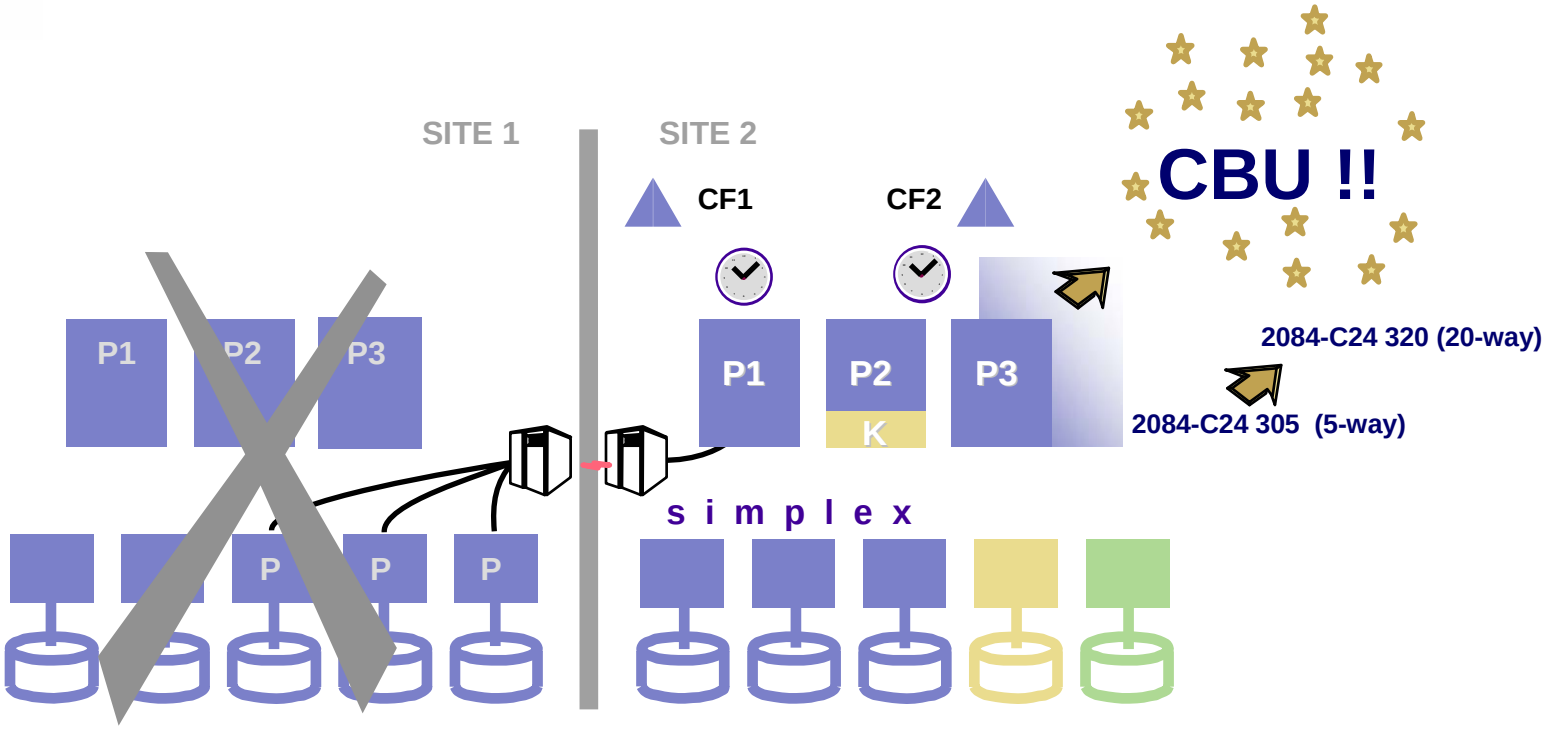

Automates recovery of production environment
Automates invocation of CBU

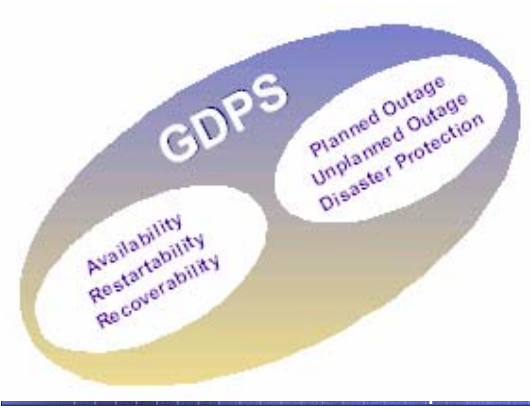

Continuous Availability and Disaster Recovery Solutions (3 site)

- ✓Continuous Availability – Metro distance
- ✓Disaster Recovery at unlimited distances

zSeries Solution designed to provide Continuous Availability and Disaster Recovery at unlimited distance

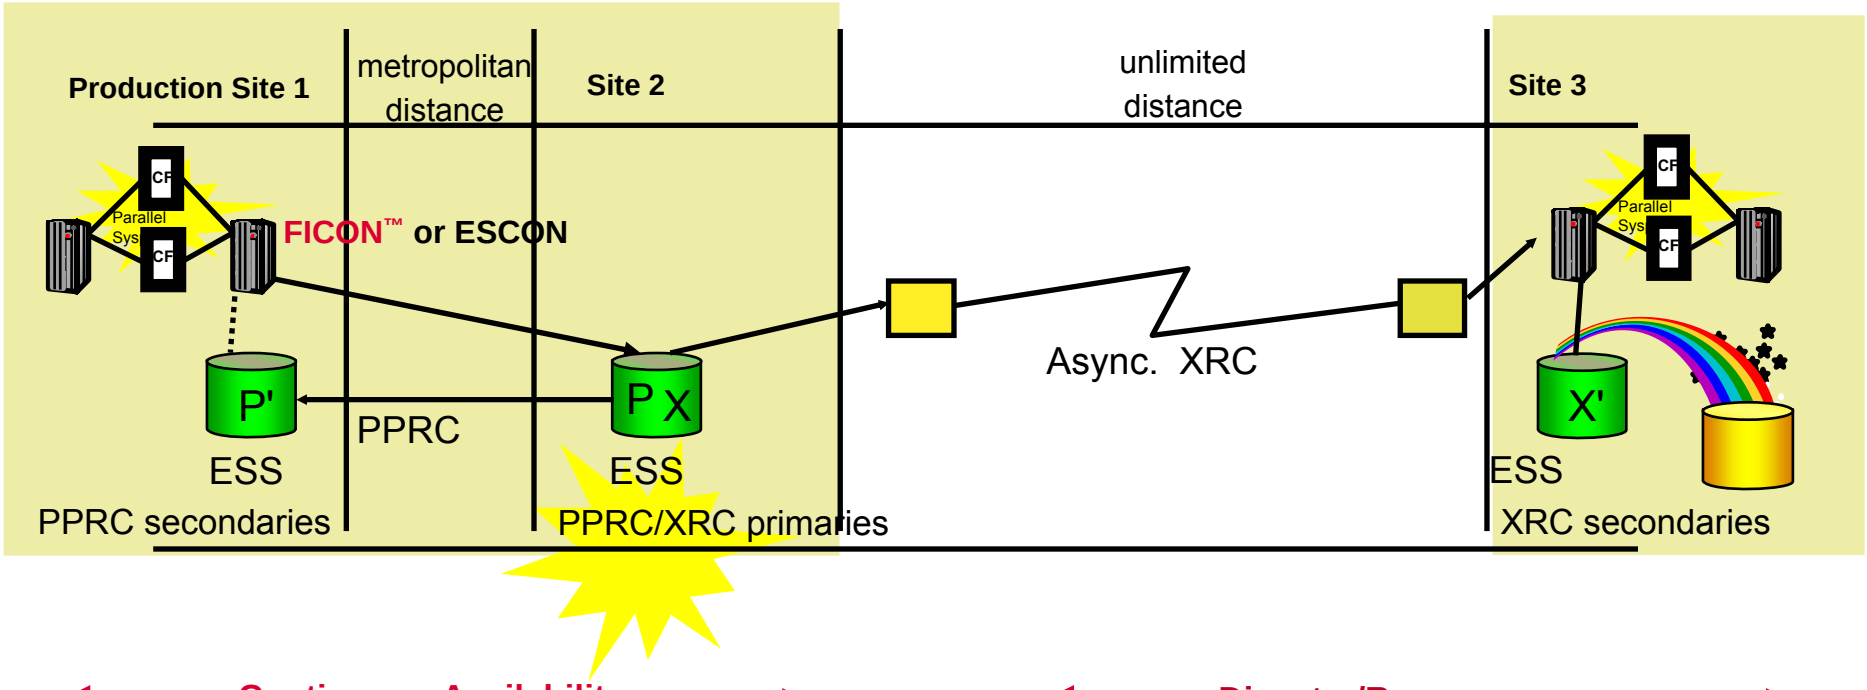

- Designed to provide continuous availability and no data loss between sites 1 and 2
- Sites 1 and 2 can be same building or campus distance to minimize performance impact

- Production site 1 failure
  - Site 3 can recover with no data loss in most instances
- Site 2 failure
  - Production can continue with site 1 data (P')
- Site 1 and 2 failure
  - SIte 3 can recover with minimal loss of data

Common Functions

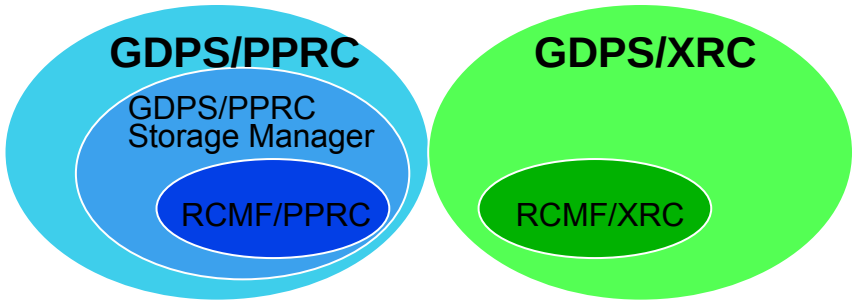

- ✓User Initiated Flash Copy
- ✓PtP Virtual Tape Server
- ✓Managed Linux for zSeries

GDPS/PPRC or GDPS/XRC or User Initiated FlashCopy

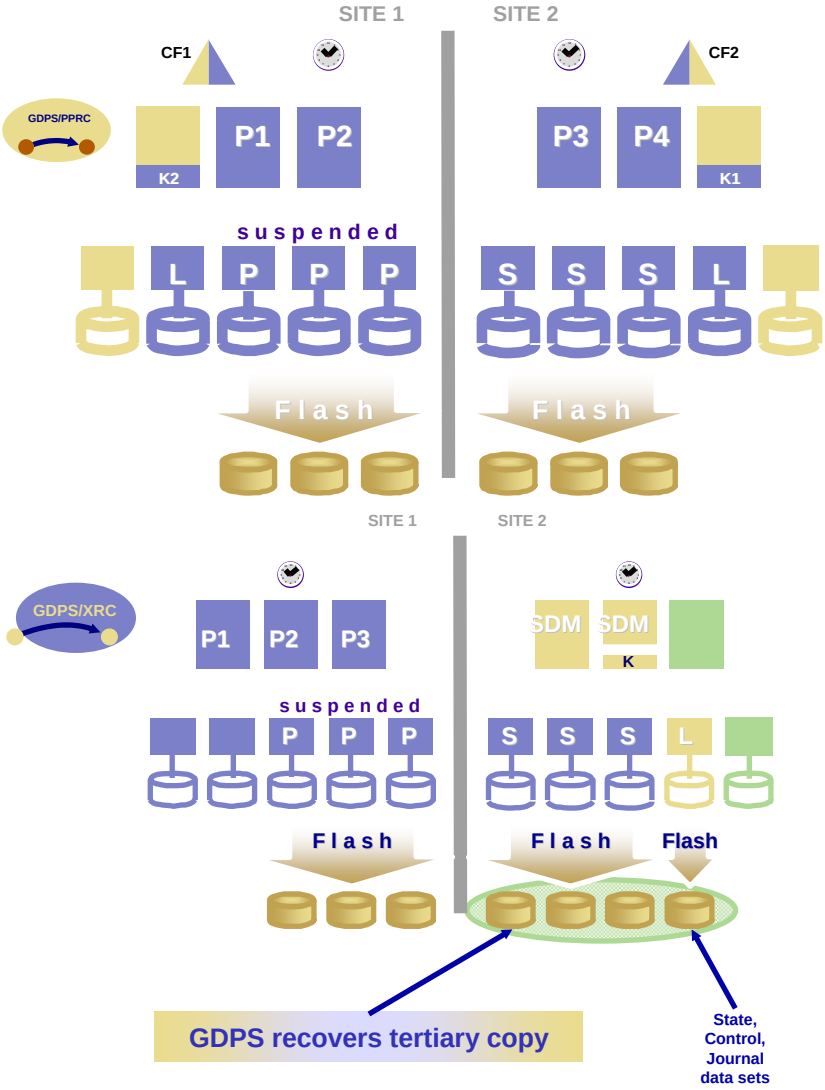

- **GDPS-initiated FlashCopy**
  - Before resynchronization
  - Maintain D/R readiness while resynchronizing
- **User-initiated FlashCopy**
  - Requested from GDPS panel or through exit
  - Enable parallel processing against tertiary copy: (point in time) backup, ...
- **FlashCopy secondaries and optionally primaries**
  - GDPS can also FlashCopy State, Control and Journal
  - GDPS can recover tertiary copy

GDPS/PPRC & GDPS/XRC Supports IBM TotalStorage Peer-to-Peer Virtual Tape Server (PtP VTS)

- Existing workload modes: Balanced, Preferred
- New workload mode to support GDPS: Primary
  - Primary VTS performs all host I/O
  - Secondary receives copies
- GDPS sets copy mode to 'immediate'
- Tape control data sets to be placed on PPRC volumes
- Planned and unplanned site switch
  - Coordinated disk / tape failover in single GDPS script
  - GDPS lists volumes 'in flight' to facilitate manual adjustment of tape control data sets after unplanned failover

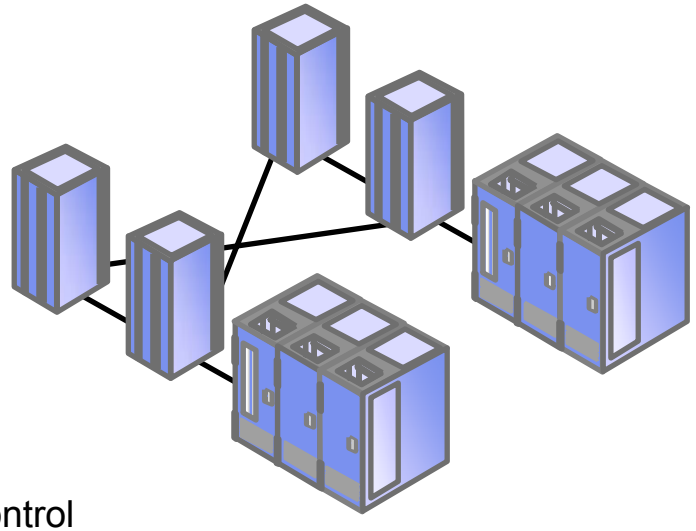

Secondary data consistent across disk and tape
Integrated disk / tape site failover
Vastly reduced complexity in site failover

GDPS/PPRC and GDPS/XRC manages Linux** for zSeries

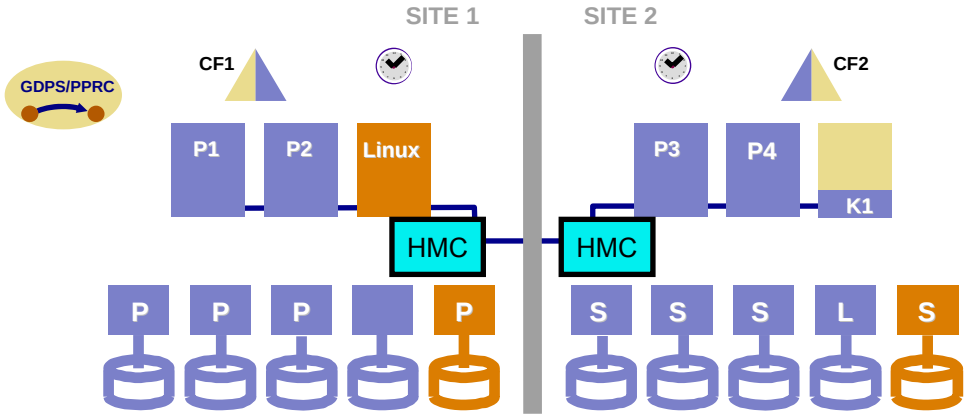

- Planned and Unplanned reconfigurations
- Unplanned Site Reconfiguration driven by z/OS
  - Controlling System recovers secondary disks including Linux
  - Expendable workload stopped
  - CBU invoked, if applicable
  - Site 1 production systems restarted including Linux
- GDPS manages PPRC Linux volumes

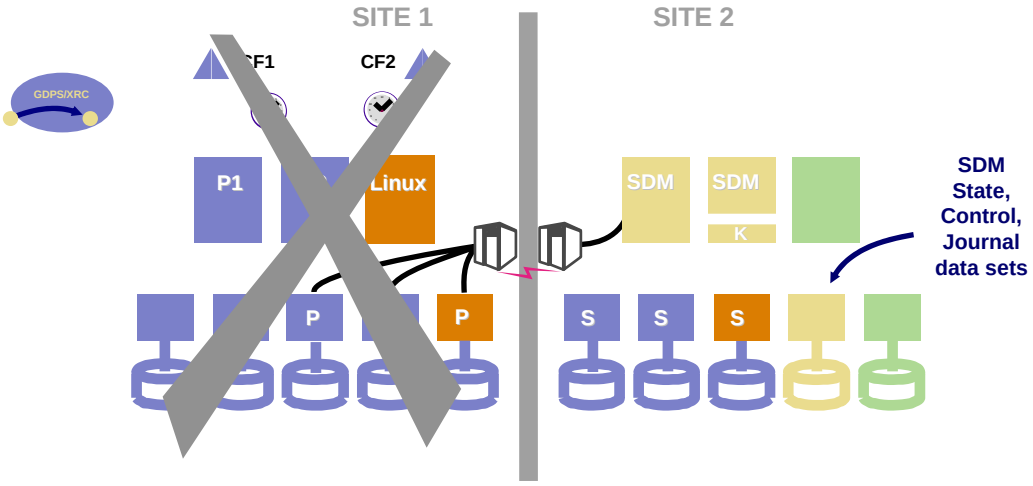

- Unplanned Site Reconfiguration manually initiated by Customer
  - Controlling System recovers secondary disks including Linux
  - Expendable workload (SDM) stopped
  - CBU invoked, if applicable
  - Site 1 production systems restarted including Linux
- GDPS manages XRC Linux volumes
- Linux for zSeries extended to time stamp data

Another step towards Enterprise-wide D/R management

**** GDPS manages all zSeries Operating Systems (e.g. z/VM, VSE/ESA)**

To Summarize

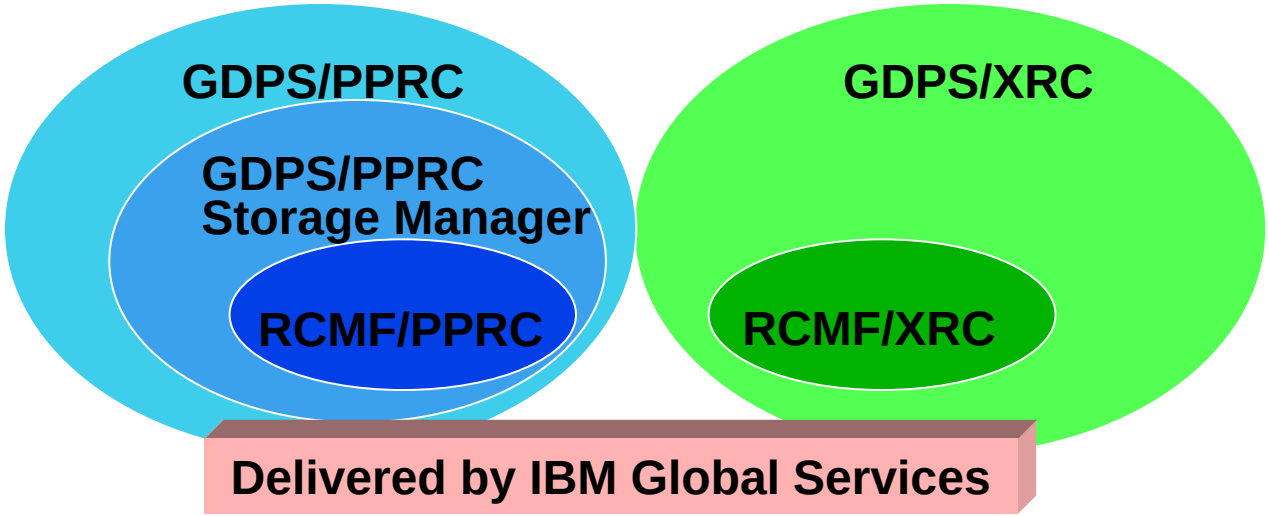

- ✓Five Implementation Options
- ✓References
- ✓Planned Enhancements
- ✓Additional Information

Summary

- In case of disaster
  - Designed to ensure data consistency and integrity
    - No data loss (GDPS/PPRC) or
    - Minimal data loss (GDPS/XRC)
  - Offers prompt, responsive disaster recovery through end-to-end automation
- GDPS/PPRC addresses both Disaster Recovery and Continuous Availability
  - Based on uninterrupted data availability through HyperSwap
- Simplifies routine management of systems, disk subsystems and data mirroring
  - Single point of control
  - Covering z/OS, z/VM, VSE/ESA, Linux and other Open Systems platforms
  - Removes stress from software, hardware or site facilities maintenance
- Solution is application independent

GDPS Experience HOUSEHOLD USA

Financial services

First GDPS/PPRC in USA
Multi Site Workload,
Freeze and Go

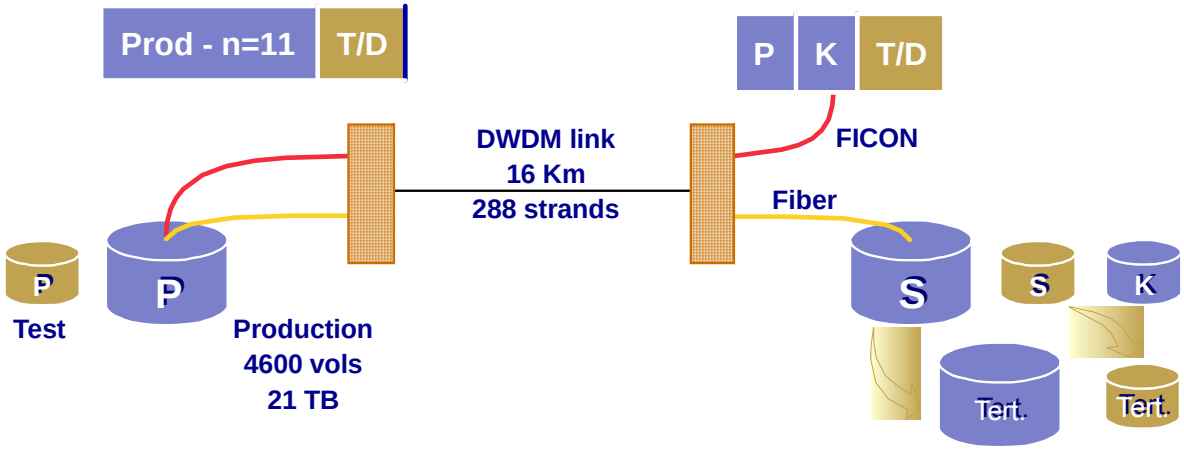

Project started March 2001
In production since Sept. 2001
Testplex considered critical success factor
GDPS implementation
- Standard actions
- Sysplex management
- Control and takeover scripts

- Business Continuity assured through GDPS/PPRC automation
- D/R test preparation time reduced from 72-96 hours to 2 hours
- Full failover test planned for 2003

GDPS/PPRC Experience

AUSTRIA

Business Objectives

- no loss of committed data
- < 2 hours recovery time (today)
- no more than five minutes disruption in the event of catastrophic systems or data center failure
- support site maintenance without application outage

- recovery window reduced from 48 hours to less than two hours
- planned site switch completed in the two hour target (without HyperSwap)
- significant reduction of on-site manpower and skill level required to manage planned and unplanned reconfigurations
- Planned HS: Dynamic switching of disk subsystems in the smallest GDPS is between 32-36 seconds and in the two larger GDPS it is about 71-75 seconds

Spardat / EB GDPS
(n=5 PS, CICS/Natural, Adabas, DB2)

Bank Austria Creditanstalt GDPS
(n=10 PS (CICS/DB2, IMS/DB, VSAM/RLS))

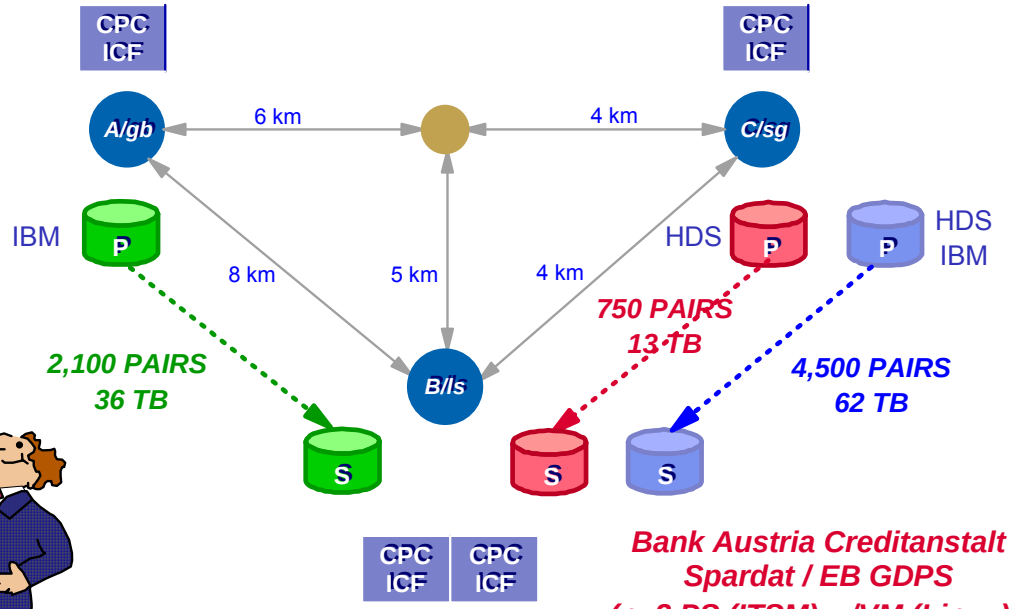

Bank Austria Creditanstalt
Spardat / EB GDPS
(n=3 PS (ITSM), z/VM (Linux))

Refer to Application Brief GM13-0142-00

Pershing , a BNY Securities Group Company

GDPS/XRC

- **Vision**
  - Leading investment banking and securities firm needed to meet Federal Reserve Bank requirement to dramatically strengthen disaster recovery capabilities
- **Challenge**
  - Reduce the Disaster Recovery Time Objective (RTO) to 2 hours, reduce the Disaster Recovery Point Objective (RPO) to 5 minutes, enable the addition of dynamic capacity and provide continuous availability
- **Solution**
  - IBM Global Services enabled an IBM GDPS/XRC configuration spanning 2 data centers and using a 2-way IBM eServer zSeries Parallel Sysplex cluster, 8 Enterprise Storage Servers, and XRC to mirror 7,350 volumes to the recovery data center
- **Value**
  - IBM eServer zSeries enabled Capacity Backup (CBU), one of the features of the autonomic computing initiative, to provide dynamic expansion of server capacity and recovery times of less than 1 hour

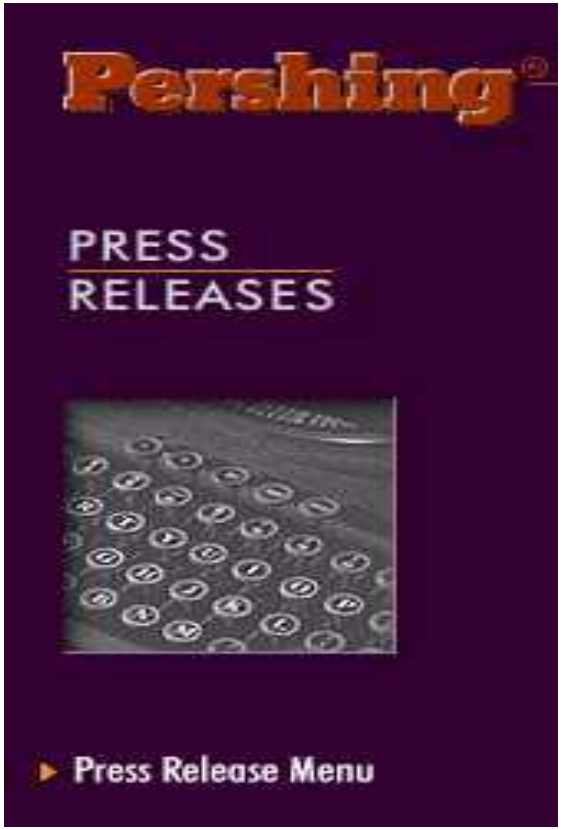

May 12, 2004
IBM And Pershing Set New Precedent For Data Recovery

Planned Enhancements

- GDPS/PPRC HyperSwap extensions
- GDPS/PPRC Multi Platform Resiliency for zSeries
- Unlimited distance solution for zSeries and Open data
  - Support for Global Mirror (2 sites)
- Unlimited distance solution for zSeries and Open data
  - Continuous Availability/Disaster Recovery (3 sites)

Business Continuity Services Offerings

- GDPS Technical Consulting Workshop (TCW)
  - **Designed to ensure the GDPS Availability & Recovery solution will meet the Client's business requirements as they relate to continuous availability and recovery. The workshop will look at the site-to-site connectivity necessary to implement GDPS and identify the high level tasks that will be needed to implement.**
- Business Continuity Solution Workshop
  - **This program is designed to introduce the elements of IBM's products and services that form a Business Continuity Solution. Your time will be divided between interactive presentations tailored to your specific requirements and "hands on labs" that allow you to actually experience the capabilities of each element. Over the course of three days at our Washington System Center you will explore topics such as: Disk and Tape Copy Services, Network Options, Server Considerations, System Performance Planning, and Implementation Services.**
- BCRS Business Continuity Health Check
  - **The Health Check is an independent review that creates an action plan addressing continuity issues such as existing capabilities, costs, future technology, and resource requirements.**
- I/O Bandwidth Analysis
  - **IBM will use trace data collected from the customer environment to determine the requirements to configure and implement Remote Copy. IBM will create a written report of the I/O Sizing and Bandwidth Analysis of your existing environment. The report will include an analysis of your full mainframe DASD environment, as well as an analysis of a subset of that environment representing the minimum DASD required to support Remote Copy.**

Additional Information

- **Detailed GDPS Presentation and Information e-mail:**
  - gdps@us.ibm.com
- **White Papers:**
  - *Business Continuity Considerations and the IBM eServer zSeries*
  - *GDPS - The Ultimate e-business Availability Solution*
- **Publications:**
  - *TotalStorage Disaster Recovery Solutions Redbook* – SG24-6547-01
  - *z/OS Advanced Copy Services* – SC35-0428
  - ESS Copy Services on zSeries Redpiece - SG24-5680
  - ESS Copy Services on Open Redpiece – SG24-5757
- **TotalStorage Business Continuance Sales Kit**
- **GDPS Services Offerings**
  - GDPS Announcement
  - GDPS/XRC Announcement
- **Business Continuity and Recovery Services**

Appendix

- ✓Metro Mirror (PPRC Synchronous), Global Mirror for zSeries (XRC) and Global Mirror (PPRC Asynchronous) Overview
- ✓Need for Time Consistency
- ✓GDPS/PPRC across 100 Km
- ✓RCMF

PPRC and XRC Overview

- **PPRC**
- 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

- **XRC**
- Asynchronous remote data mirroring
  - Application receives "I/O complete" as soon as primary disk is updated
- Unlimited distance support
- Performance impact negligible
- System Data Mover (SDM) provides
  - Data consistency of secondary data
  - Central point of control

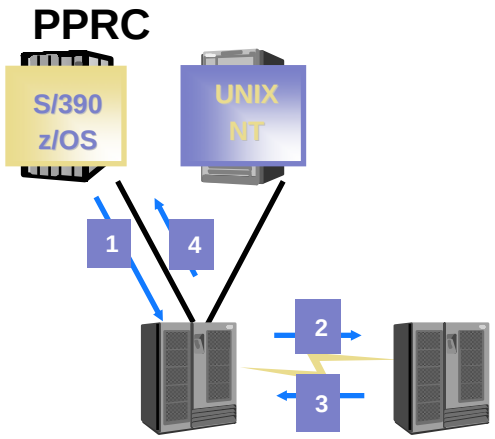

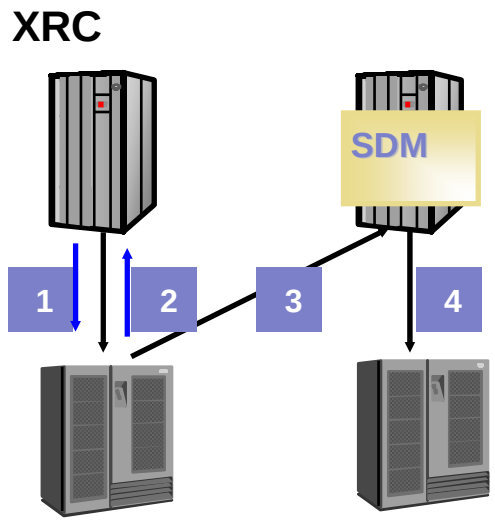

Storage Consolidation and Disaster Recovery with *Global Mirror for ESS*

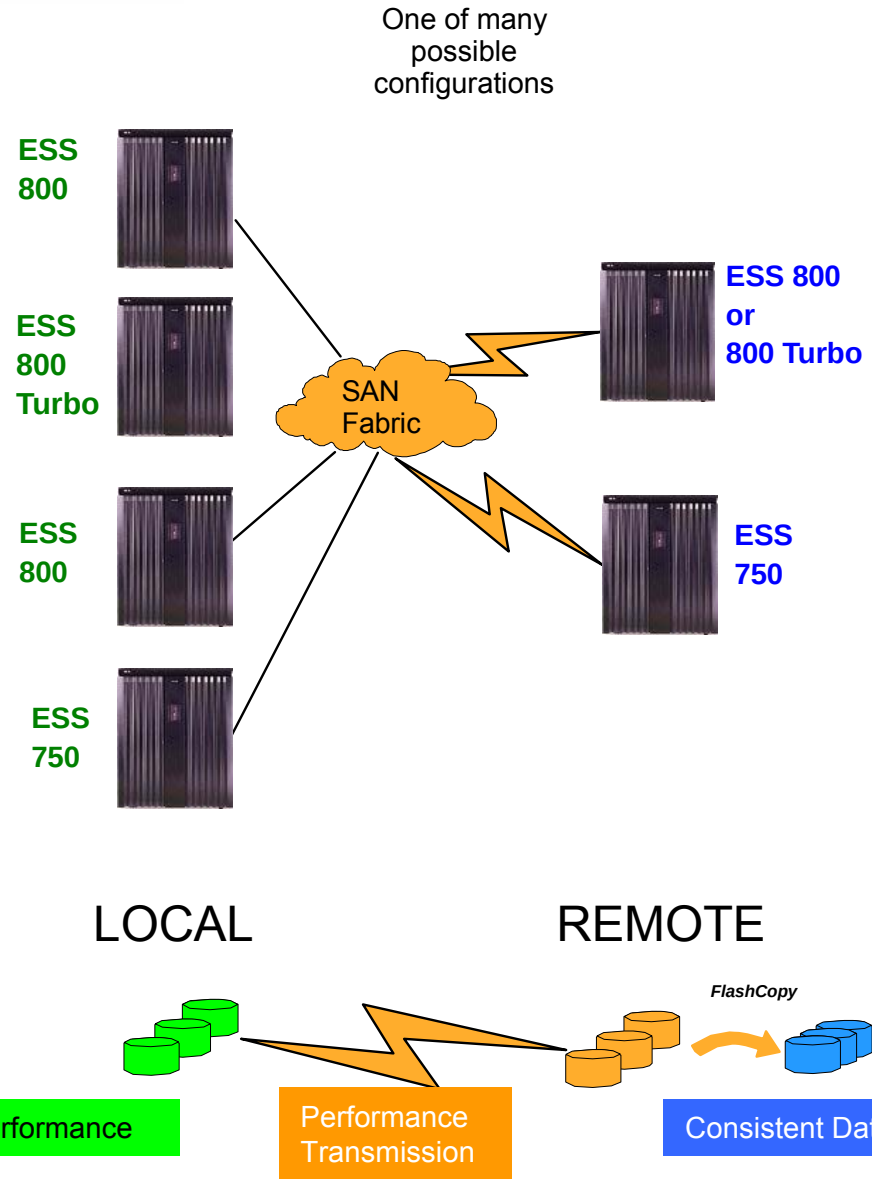

- Designed to Provide:
  - **Global Distance:** Two-site, unlimited distance, data consistent asynchronous disk mirroring
  - **Scalability:** Consistency Group supported across up to 8 total ESSs in Global Mirror session
  - **Flexibility:** Many possible configurations
  - **Heterogeneous:** Data can span zSeries® and open systems data, and can contain a mix of zSeries and open systems data
  - **Application Performance:** Native
  - **Mirroring Performance:** Two ESS Fibre Channel disk mirroring links per ESS sufficient for almost all workloads
- Intended Benefits
  - **Autonomic:** No active external controlling software required to form consistency groups
  - **Saves cost:** No server cycles required to manage consistency groups, fewer
  - **Lowers TCO:** designed to provide improved performance, global distances; and lower costs

Global Mirror for ESS Configuration Requirements

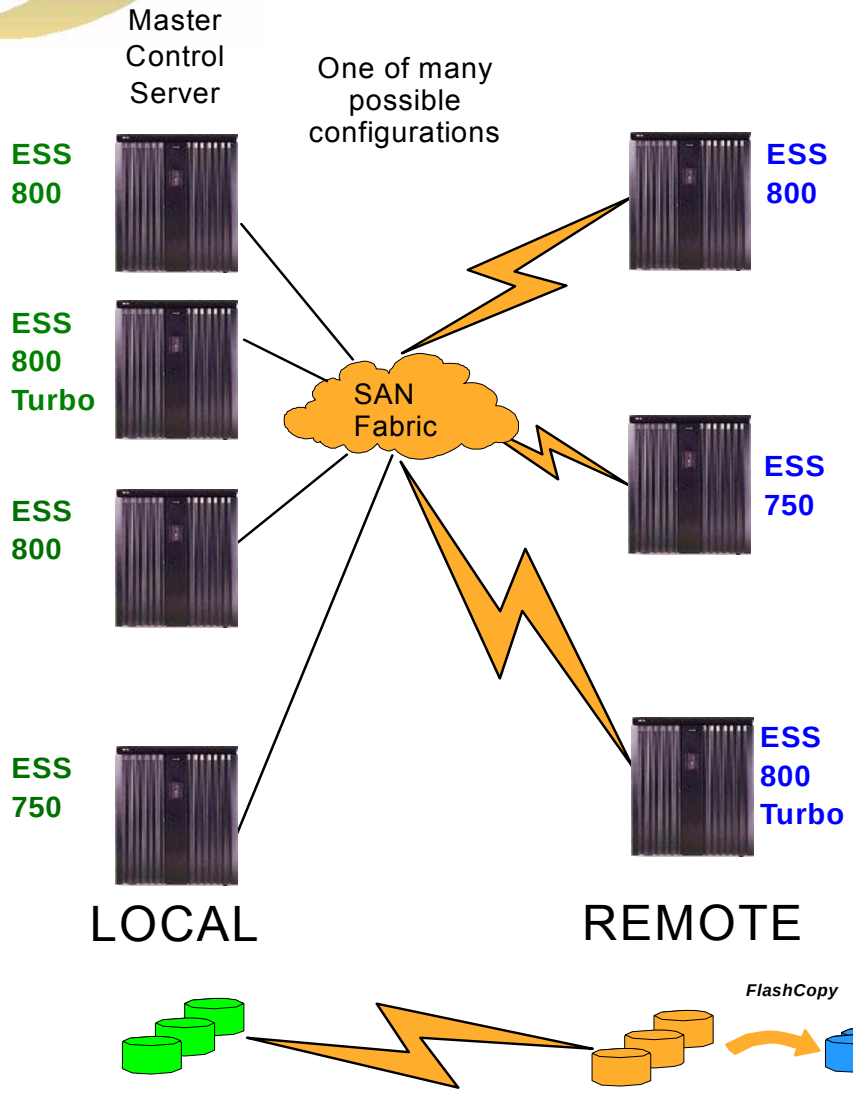

- Configuration Requirements
  - **Copy consistency:** managed autonomically by Master Control Server in master ESS
  - **Requires Fiber Channel paths:** between all ESS's in the Global Mirror session, for control commands and data
  - **Exploits Fibre Channel SAN infrastructure:** uses standard Fibre Channel SAN infrastructure, distance extension, switch blades, and wide area network topologies
  - **Maximum configuration:** Up to 8 ESSs in a PPRC Global Mirror session
- Prerequisites
  - ESS Model 800, Model 800 Turbo, or Model 750 with ESS 2.4 microcode
  - Fibre Channel connections between all ESSs in session
  - PPRC V2 at local site
  - PPRC V2 and FlashCopy V2 at remote site

Need for Time Consistency

Recovery	Restart
Process measured in hours or days Restore last set of Image Copy tapes Apply log changes to bring database up to point of failure	Process measured in minutes To start a DB application following an outage without having to restore the database

Database restart required for today's high availability requirements
(Can your business tolerate a lengthy database recovery ?)

- Many examples where the start of one write is time dependent on the completion of a previous write
  - Database & log
  - Catalogs, Volume Table of Content
  - Index & data components
- Time sequence could be exposed during Rolling Disaster
- GDPS Freeze function helps assure time consistency of data
- XRC manages secondary consistency
  - Across any number of primary subsystems
  - All writes time-stamped and sorted before committed to secondary devices

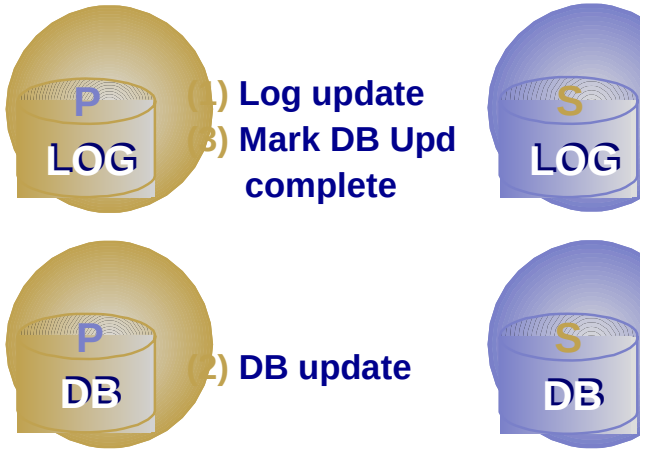

GDPS/PPRC

GDPS/PPRC across 100 km (example)

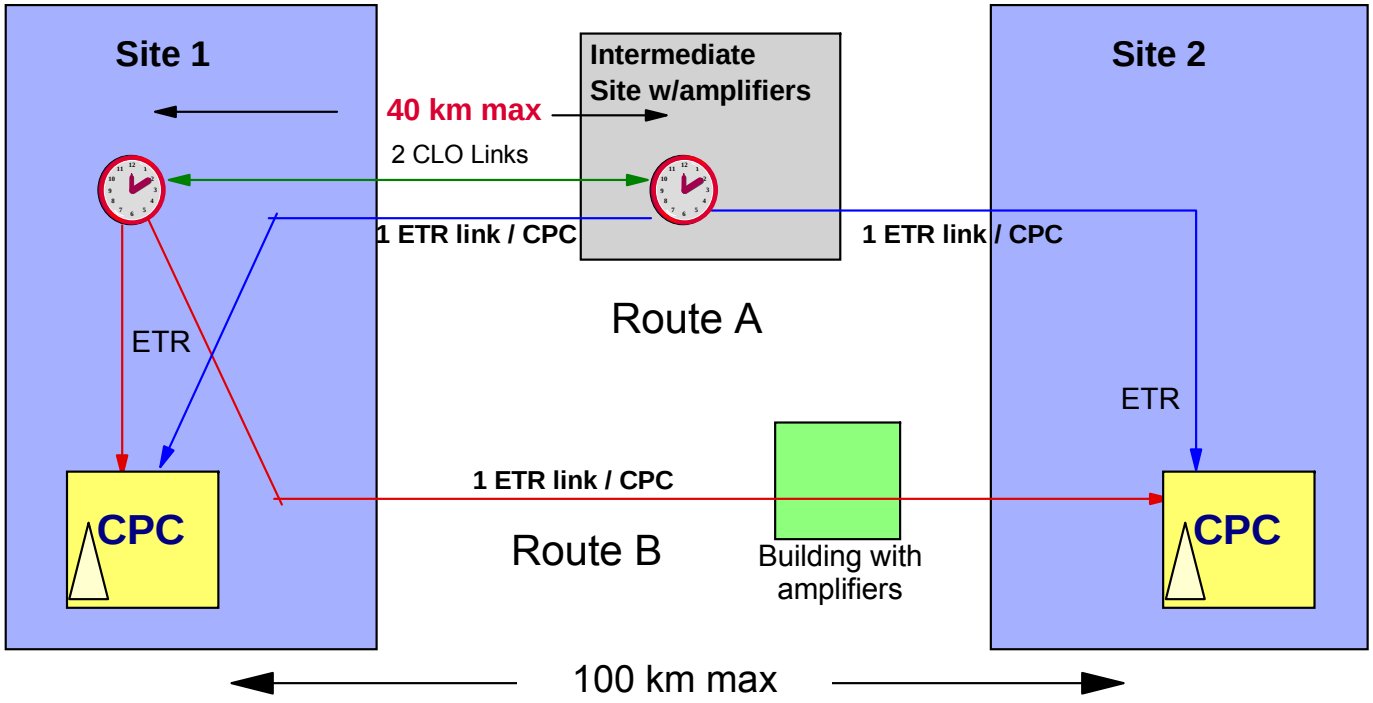

Redundant ISC3 cross site links from Site 1 to Site 2 also extended to up to 100 km max

XRC Intersite Connectivity Options

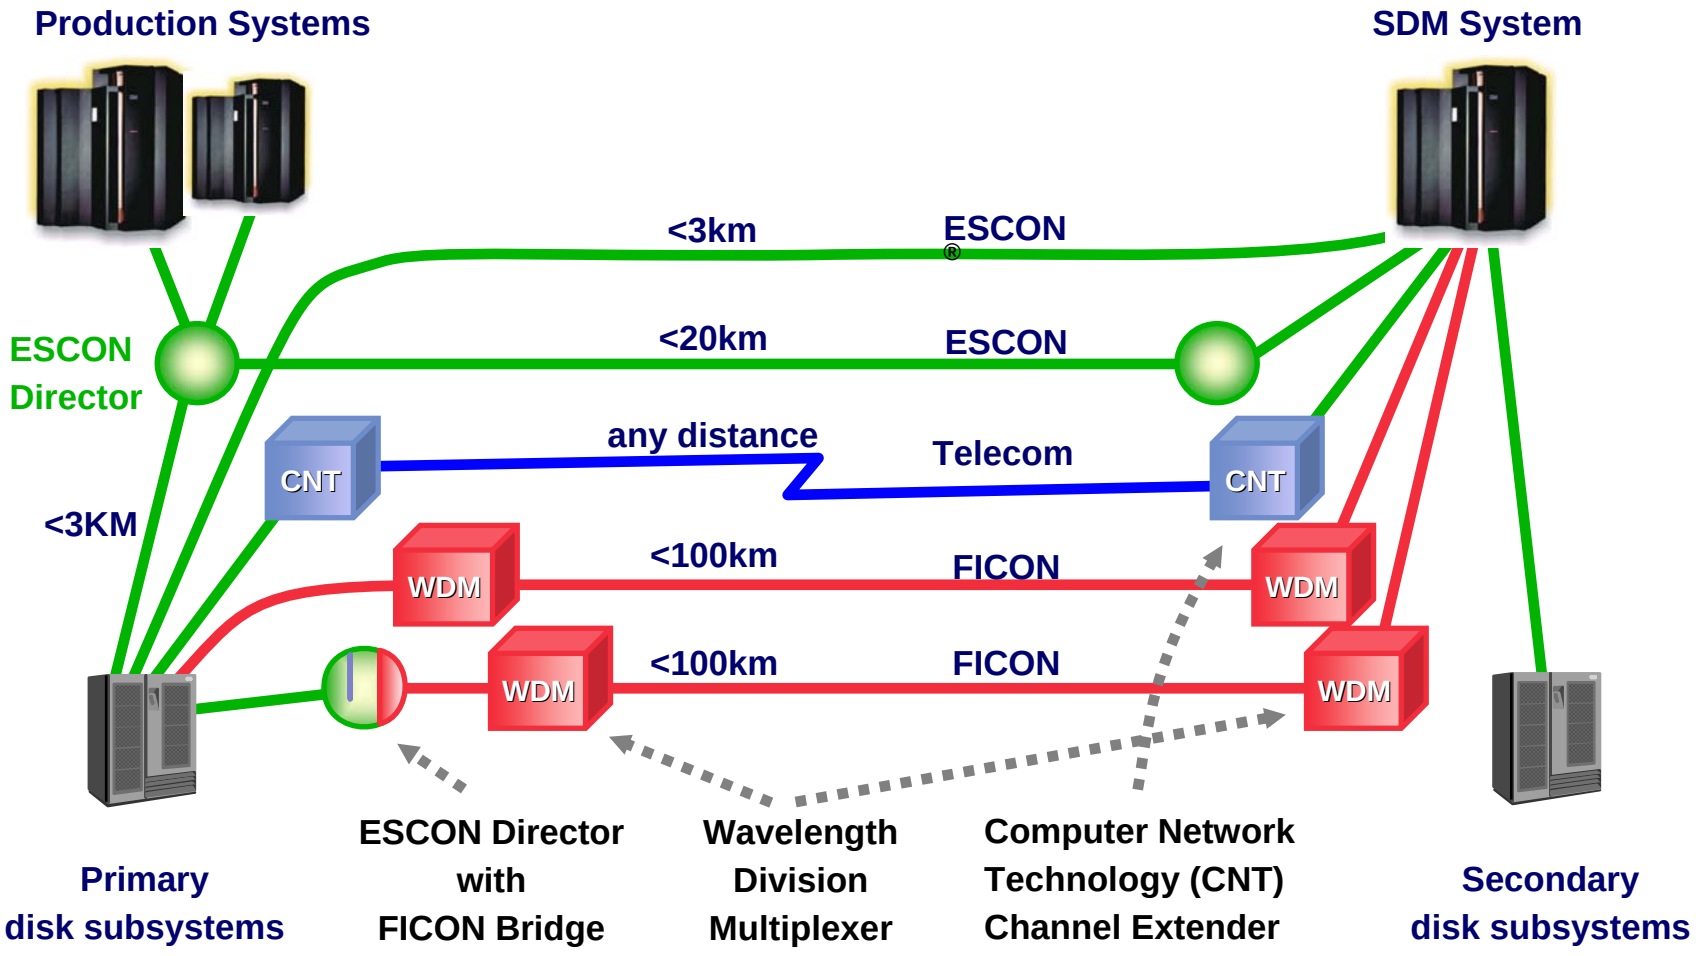

RCMF

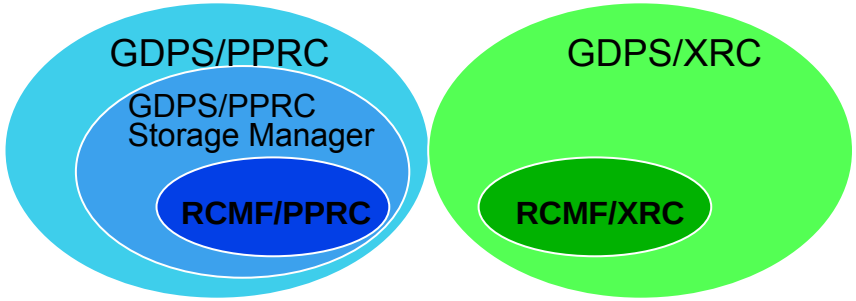

- ✓RCMF/PPRC
- ✓RCMF/XRC

Remote Copy Management Facility / PPRC

- Central point of control - full screen
  - Global PPRC configuration awareness
  - Functional, tree-structured interface
  - TSO commands not needed
- Initialize & maintain Remote Copy configuration
  - Single keystroke function invocation
  - Initiate functions per pair, subsystem or all
  - Automatically establish target configuration at system startup
  - Supports adding, moving, removing pairs, subsystems, links
  - (RCMF does not manage secondary consistency)
- User-initiated status & exception reporting
- Runs as a NetView application - SA for z/OS not required

Manage remote copy configurations vs. remote copy pairs

Remote Copy Management Facility / XRC

- Central point of control - full screen
  - Functional, tree-structured interface
  - TSO commands not needed
- Initialize & maintain Remote Copy configuration
  - Single keystroke function invocation
  - Initiate functions per pair, subsystem or all
  - Automatically establish target configuration at system startup
  - Supports adding, moving, removing pairs, subsystems, links
- User-initiated status & exception reporting
- Runs in System Data Mover system and controls single SDM
  - Single SDM can handle 1000-2000 3390-3 pairs
- z/OS executing in 31-bit mode)
- Runs as a NetView application - SA for z/OS not required

Manage remote copy configurations vs. remote copy pairs
XRC (with RCMF) provides a valid D/R solution