IBM Systems and Technology Group University 2005

IBM Systems and Technology Group University 2005

What is new with GDPS in 2005?

Course #: ZT32

Noshir Dhondy
Senior Engineer
dhondy@us.ibm.com

Agenda

- GDPS solutions
- Near-Continuous Availability (NCA) of Data
  - HyperSwap Manager
  - Disk Maintenance and Disk Failures with HyperSwap
- Metropolitan Distance NCA/Disaster Recovery (D/R) Solution (2 sites)
  - Configuration Options
  - Planned and Unplanned Site Reconfiguration with HyperSwap
  - Open LUN Management
  - Multi Platform Resiliency with zSeries
- Unlimited Distance D/R Solution (2 sites)
  - z/OS® data only
- NCA/DR Solution (3 sites)
  - z/OS® data only
- Planned Enhancements

Learning Objectives

- At the end of this training, you should be able to
  - Understand how you can extend Parallel Sysplex Availability to disk subsystems with a new IGS offering, GDPS/PPRC HyperSwap Manager
  - Describe how GDPS provides a coordinated Near-Continuous Availability and Disaster Recovery solution for both z/OS and Linux for zSeries
    - Multi Platform Resiliency for zSeries
  - Understand GDPS planned enhancements for 2005 and the GDPS Roadmap beyond 2005

IBM Systems and Technology Group University 2005

GDPS Solutions

Business Resiliency/Security Roadmap
What is GDPS?

Extending zSeries Business Resiliency Strengths

Creating a "Secure Vault[1]" for Heterogeneous Environments

Today's Capabilities

- Up to 99.999% availability[2] across all zSeries resources to avoid planned and unplanned outages
- Minimizes interruption through intelligent management of system resources based on business priorities
- Enables integrity of zSeries data with capability to secure at granular user level

Initial Focus*

- Enable asset and resource associations to be mapped by business function.
- **Extend GDPS to heterogeneous environments and single site environments.**
- Extend zSeries ability to provide seamless integration of Security through utilization of open technologies.

Extension of on demand value*

- Provide capability to monitor business functions and identify appropriate recovery actions.
- Extend zSeries resiliency into heterogeneous environments through utilization of common interfaces.
- Position zSeries to lead in managing heterogeneous assets and resources.

Future Vision*

Become the "secure vault" across the enterprise.

- Become the "world's most resilient enterprise" through zSeries autonomic management based on business policy support for heterogeneous assets and resources.
- Become the enterprise trust authority through integration of zSeries Security leadership with the ODOE.

Note 2: Based on Parallel Sysplex implementation

Note 1: Secure Vault Represents the role and ability of zSeries to provide a highly resilient and security-rich enterprise-wide environment for enterprise data and transactions

Business Resiliency / Security Roadmap

"Secure Vault" and Enterprise Trust Authority*

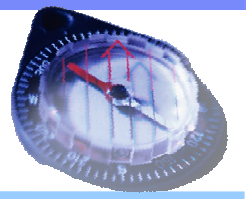

Business Value

Technology

Data Protection

- Automated data protection helps reduce business risk and frees up zSeries resources
- Integrated zSeries data protection and recovery

Open Industry Standards

- Enhance customer investments in availability by leveraging technologies such as GDPS and Data Mirroring
  - Enable zSeries componentry to be a key participant in an end to end environment
  - Start integration of silo solutions to achieve real business resilience
  - Provide availability impact analysis of resources and their associated business processes

Heterogeneous End-to-end support

- Leverage the resilience of the zSeries platform to manage heterogeneous enterprise availability
  - Align availability to business goals
  - Enable the granular recovery of business components
  - Enable integration with vendor availability solutions
  - Enable businesses to focus on IT innovation not IT management

Autonomics

- e2e resources managed to overall business objectives
  - BR fully integrated with overall systems management fabric
  - Enables businesses to focus on business priorities, not IT management

Future Vision

Tomorrow

- Integration of other Availability solutions into OD Framework
- Industry Specific Solutions
- Coordinated integration with service providers
- Complete and integrated mapping of business goals to infrastructure deployment

- Enable z/OS as business resilience management server
- Integrate with customer automation.
- Enable business modeling tools to generate availability policy

Today

- BR Autonomic Manager
- **Integrate GDPS into BR Framework**
- End to end topology of bus. apps.
- Standard and reuse-able model of application and infra resources
- Common Information Model (2005)

- GDPS Multi Platform Resiliency for zSeries support for Linux on zSeries running as a z/VM guest (Sept. 2004).
- **GDPS HyperSwap Manager (1Q2005).**
- **GDPS/Global Mirror (1Q/2Q05)**

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

It is our intention to:

Innovation

Provide leadership in innovation to enhance the use of IBM eServer zSeries to support increasingly integrated and flexible business processes for the on demand business. *

On demand capability Roadmaps:

Business Integration

- z/OS 1.6 and zAAP exploitation
- Communication Controller on Linux (2005*)

Business resiliency & security

- **GDPS Multi Platform Resiliency for zSeries supporting Linux for zSeries running as z/VM guest (2004*)**
- **GDPS Hyperswap Manager (2005*)**

Intelligent Business Director:

- eWLM for zOS (2004*)
- eWLM for Linux (2005*)
- Common Information Model (2005*)
- IBM Director Multiplatform (2005*)
- TBSM and TEC for managing business service levels

Value

Enhance the value proposition and lower the cost of computing of zSeries solutions in a way that is compelling, clear, and consistent.*

Flexible & Responsive Pricing Models:

- Subcapacity pricing for key platform Software
  - Recently adopted by BMC and CA
- Broad portfolio of Capacity on Demand offerings

zSeries ServerProven Rebate Offering***

Offerings for Integration and Simplification

- zAAPs and IFLs
- Application transformation and integration services
- Scorpion Studies

Offerings for Resiliency and Security

- **GDPS®/PPRC Implementation Services (now supports GDPS Multi Platform Resiliency for zSeries)**
- **GDPS® HyperSwap Manager – significant savings on select function SA and Tivoli Netview suite**
- **GDPS/Global Mirror**
- Base zSeries and z/OS Security workshops

Offerings for System Mgmt and Optimization

- End to End Systems Management Services
- IT Optimization Solution Offering
- 12-Step Strategic Virtualization Assessment

Community

Support programs designed to foster vitality in the zSeries community, helping to promote a strong application portfolio and world-class support services.*

Broadening ecosystem to enable customer with support and skills

- Greater support to enable participation
- Over 50 new ISVs with 150 new applications in 2003
- Already 39 new ISVs and 131 new applications through 1H04 and growing

FSS Reference Architectures

- Live banking demo in MOP

Exponential increase in Scholars Program

- 70 Universities enrolled
- Targeting 20,000 new zSeries trained people in market by 2010.

** Excerpted from the Mainframe Charter – August 2003

*** US and Canada Only

What is GDPS?

- **Automated solution 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**

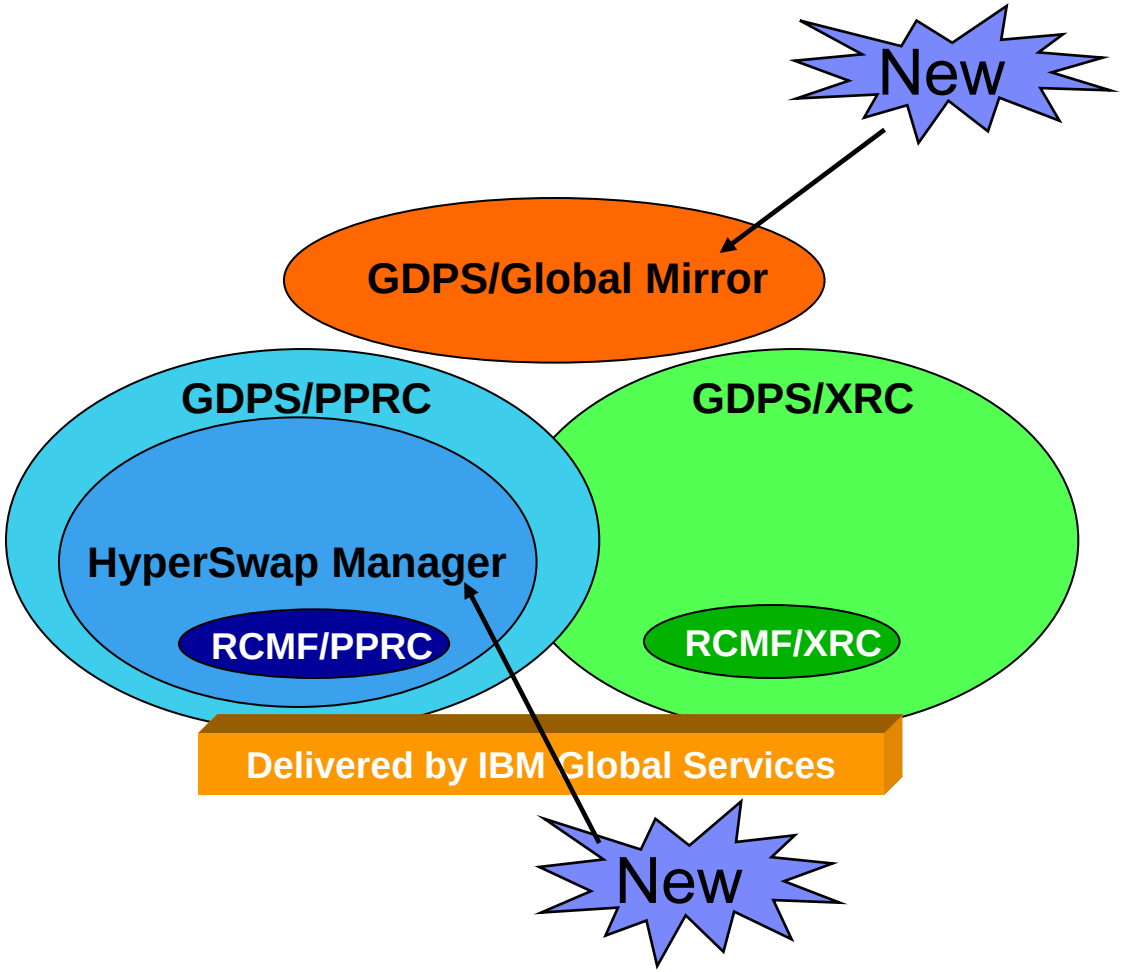

Designed for Near-Continuous Application & Data Availability
Single point of control
Delivered through IBM Services

GDPS Solutions

Near-Continuous Availability (NCA) of Data

Solution	Target Customer	Value
GDPS/PPRC HyperSwap Manager **(1)**	Parallel Sysplex	NCA of zSeries and Open Systems data

Metropolitan Distance NCA/DR 2 sites

Solution	Target Customer	Value
RCMF/PPRC	Disk mirroring	PPRC management Ease of use
GDPS/PPRC HyperSwap Manager **(1)**	Entry Level Disaster Recovery / zSeries and Open data	NCA of data Site failure protection; RPO<1 min. RTO depends on cust automation
GDPS/PPRC **(4)** Sysplex/PPRC across 2 sites - Single site or Multi site Workload	DR for zSeries and Open data - Continuous zSeries data availability	Unplanned/Planned reconfig RPO=0; RTO < 1hr
GDPS/PPRC BRS configuration Sysplex in one site PPRC across sites	Disaster Recovery - zSeries and Open data	Unplanned/Planned reconfig RPO=0; RTO < 4 hrs

Unlimited Distance D/R 2 sites

Solution	Target Customer	Value
RCMF/XRC	Disk mirroring	XRC management Ease of use
GDPS/XRC	zSeries Disaster Recovery	Site failover RPO=< 2 min; RTO 1-2 hrs.
GDPS/Global Mirror **(2)**	zSeries + Open Disaster Recovery	Site failover RPO=minutes RTO <2 hrs

NCA/DR 3-sites (Metro + Unlimited dist)

Solution	Target Customer	Value
GDPS/PPRC+ GDPS/XRC (zSeries data only)	Economically essential businesses, Ultimate in BR	Metro distance CA for zSeries data + Unlimited distance DR
GDPS Metro/Global cascading (zSeries+Open data) **(3)**	zSeries + Open Economically essential businesses, Ultimate in BR	Metro distance CA for zSeries data + Unlimited distance DR

(1) Not announced - Feb 2005 target

(2) Not announced - z/OS data Feb 2005 target
- z/OS & Open data 2Q05 target

(3) Not announced - 4Q05 target

(4) Multi Platform Resiliency for zSeries (zLinux guest operating under z/VM) - GA Sep 2004

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

Near-Continuous Availability of data

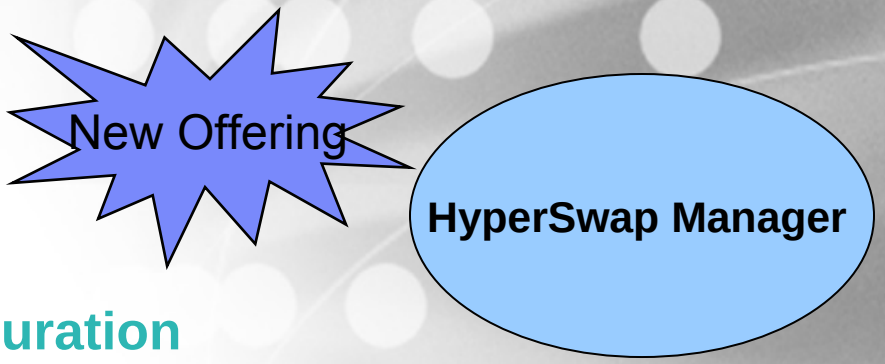

Unplanned disk reconfiguration
Planned disk reconfiguration
Specially priced Tivoli NetView and System Automation

ON DEMAND BUSINESS™

GDPS/PPRC HyperSwap – the Technology

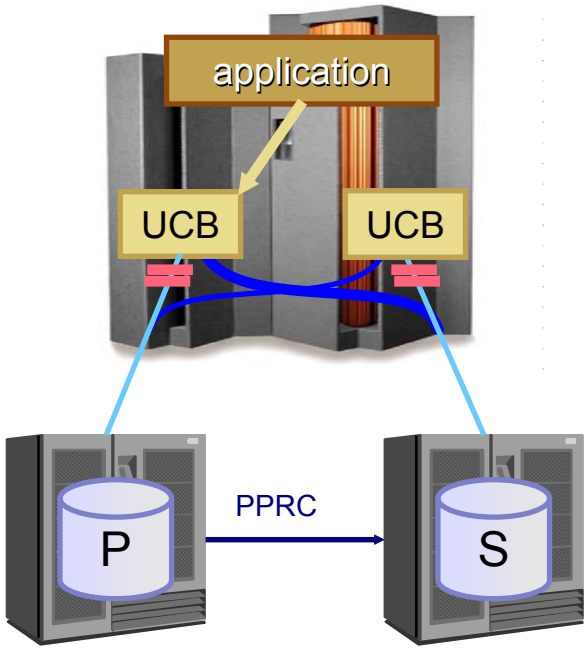

- Substitutes PPRC secondary for primary device
  - No operator interaction - GDPS-managed
  - Can swap large number of devices - fast
  - Includes volumes with Sysres, page DS, catalogs
  - Non-disruptive - applications keep using same device addresses

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

GDPS/PPRC HyperSwap Manager (HM) Benefits

- Extends Parallel Sysplex Availability to disk subsystems
- Effective entry level offering for customers that require high levels of availability
- Specially priced Tivoli NetView and System Automation products
- Combines the features of Remote Copy Management with the automation of GDPS
- Simplifies management of Remote Copy configuration, reducing storage management costs
- Reduces time required for remote copy implementation
- Investment protection
  - Positioned to upgrade to full GDPS

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

GDPS/PPRC HM Functional Overview

- Single point of control to manage the remote copy configuration
  - zSeries and Open data
    - Cannot HyperSwap Open Data; Open data will be “frozen” to maintain data consistency
- Masks primary disk subsystem failures by transparently switching to use secondary disks
- Provides ability to perform disk maintenance without requiring applications to be quiesced
- Enables data consistency in the event of failures or disaster
- 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

PPRC Failover / Failback

- **The new primary volumes (at the remote site) records changes while in failover mode. The original mode of the volumes at the local site is preserved as it was when the failover was initiated.**
- **Only need to resynchronize from time of failover, not entire data set**

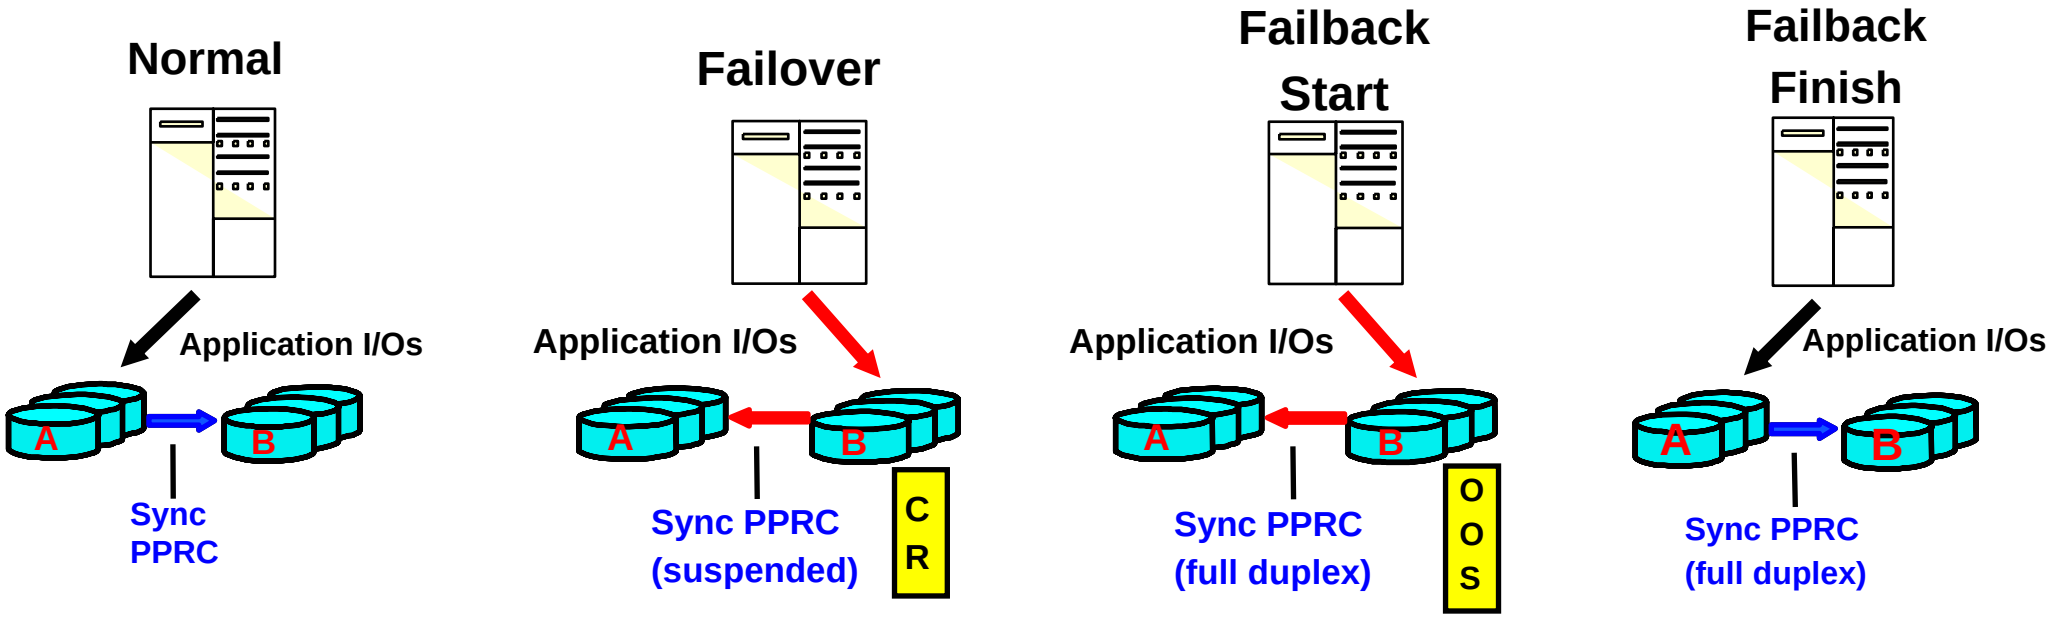

- Faster Resynchronization
- Less resource consuming

Unplanned Disk Reconfiguration with HyperSwap

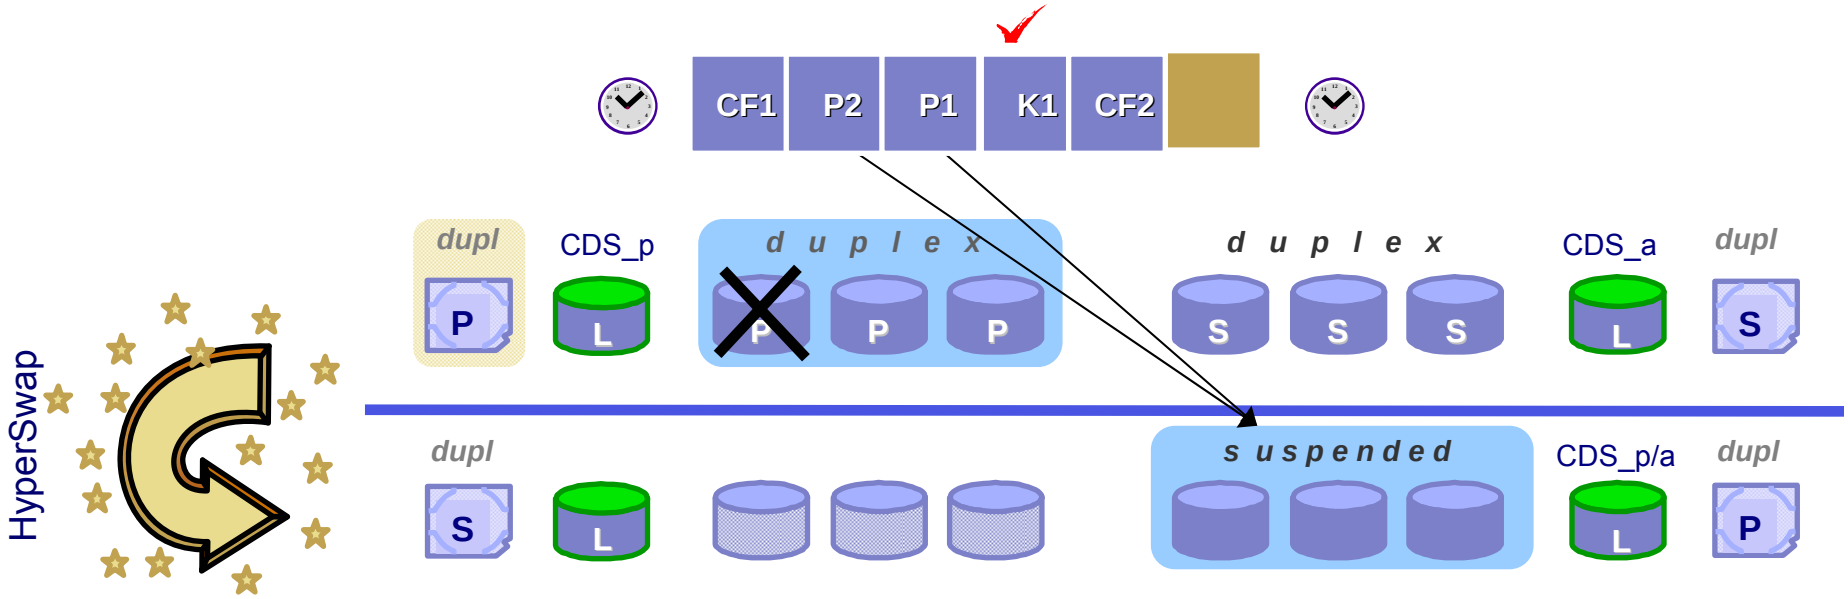

- Parallel Sysplex – P1, P2, K1 (GDPS controlling system)
- **Disk Failure detected**
- **GDPS automation invokes**
  - HyperSwap disk configuration
    - Swap primary / secondary disks
    - **Failover invoked (secondary disks in suspended state)**
  - After primary disk failure fixed
    - **Failback invoked (updates to data copied)**
    - Execute HyperSwap again to return to original configuration

PPRC failover/failback exploitation eliminates need for full copy when going back

P1, P2, remain active throughout the procedure

GDPS/PPRC

Benchmark Measurements – Unplanned Disk Reconfiguration

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 Without FO/FB

18 Seconds! (2900 vol pairs 4.6 TB)
Initial copy of all volumes required to restore to original configuration

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

With HyperSwap and FO/FB

15 Seconds! (2900 vol pairs 4.6 TB)
Only changed data needs to be copied to restore to original configuration

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

New Function

GDPS/PPRC

Planned Disk Configuration with HyperSwap

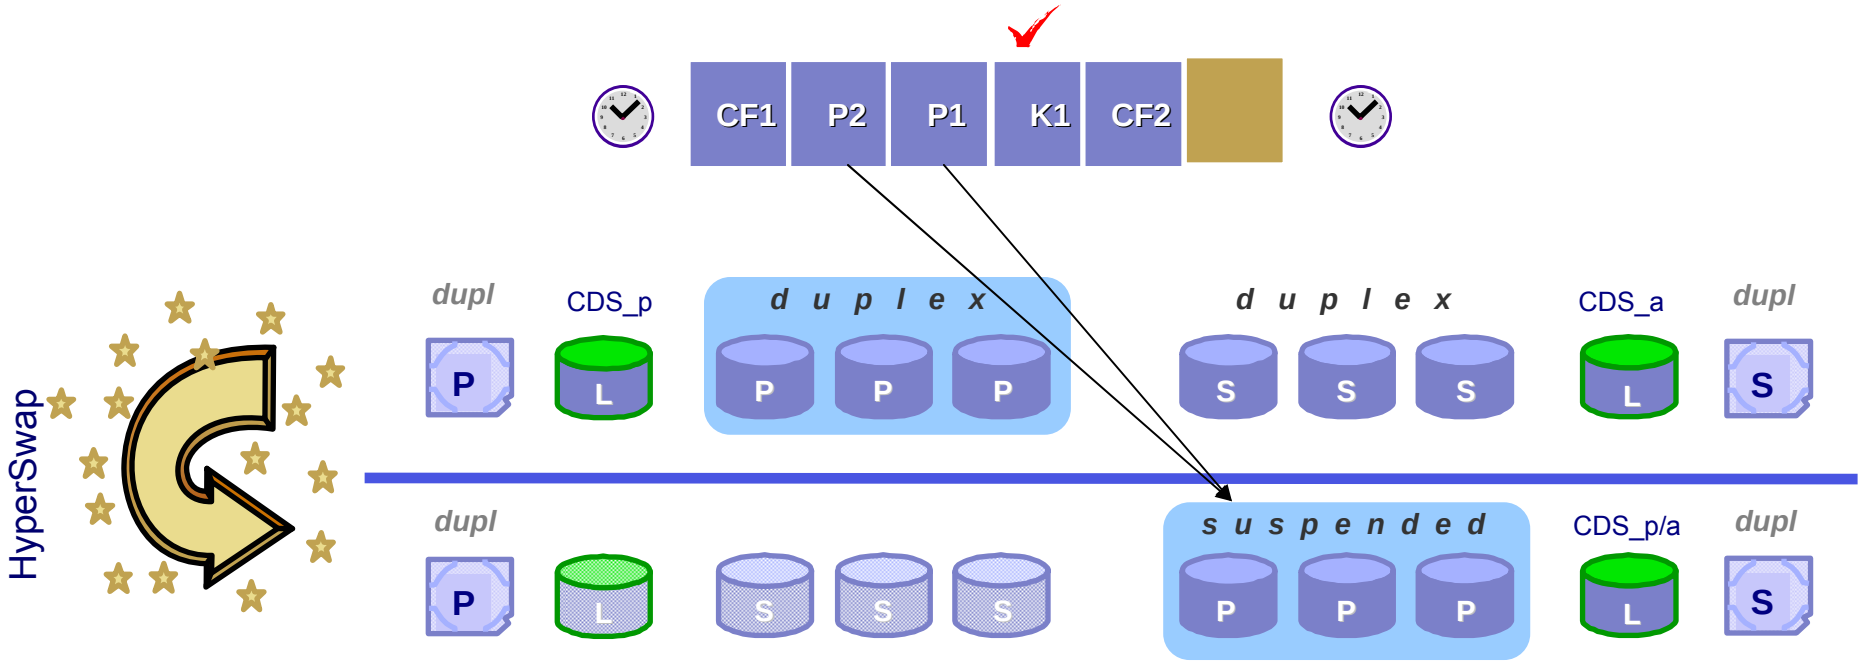

- **GDPS command**
- HyperSwap disk configuration
  - Swap primary / secondary disks
  - **Failover invoked (secondary disks in suspended state)**
- After maintenance is completed
  - **Failback invoked (updates to data copied)**
  - Execute HyperSwap again to return to original configuration

P1, P2 remain active throughout the procedure

PPRC failover/failback exploitation expected to reduce planned swap times by 50%

Benchmark Measurements – Planned Disk Reconfiguration

GDPS/PPRC

Without HyperSwap

PLANNED ACTION INITIATED

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

1-2 hrs (approx)

With HyperSwap

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

Reference Customer	Configuration	Switch Time (without FO/FB)	Switch Time (with FO/FB)
ARZ (Austria)	2300 vol pairs (14 TB) (Note 1)	82-84 secs	
Postbank (Germany)	1800 vol pairs (32 TB)	80-84 secs	
iT Austria (Austria)	4200 vol pairs (24 TB)	75 secs	
iT Austria (Austria)	4500 vol pairs (76 TB)	75 secs	
IBM test facility (MOP)	2900 vol pairs (4.6 TB)	93 secs	18 secs
	Note 1:TB depends on 3390-3 or 3390-9 type vols		

New Function

IBM Systems and Technology Group University 2005

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

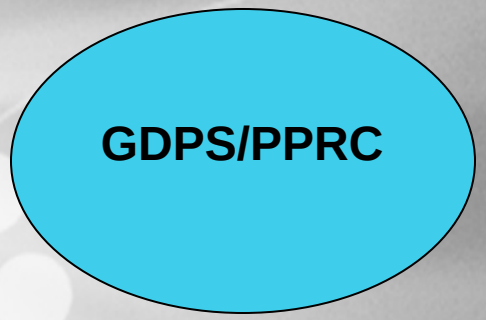

GDPS/PPRC
Configurations
Management of Open Systems LUNs
Multiplatform Resiliency for zSeries

ON DEMAND BUSINESS™

What is GDPS/PPRC? (Metro Mirror)

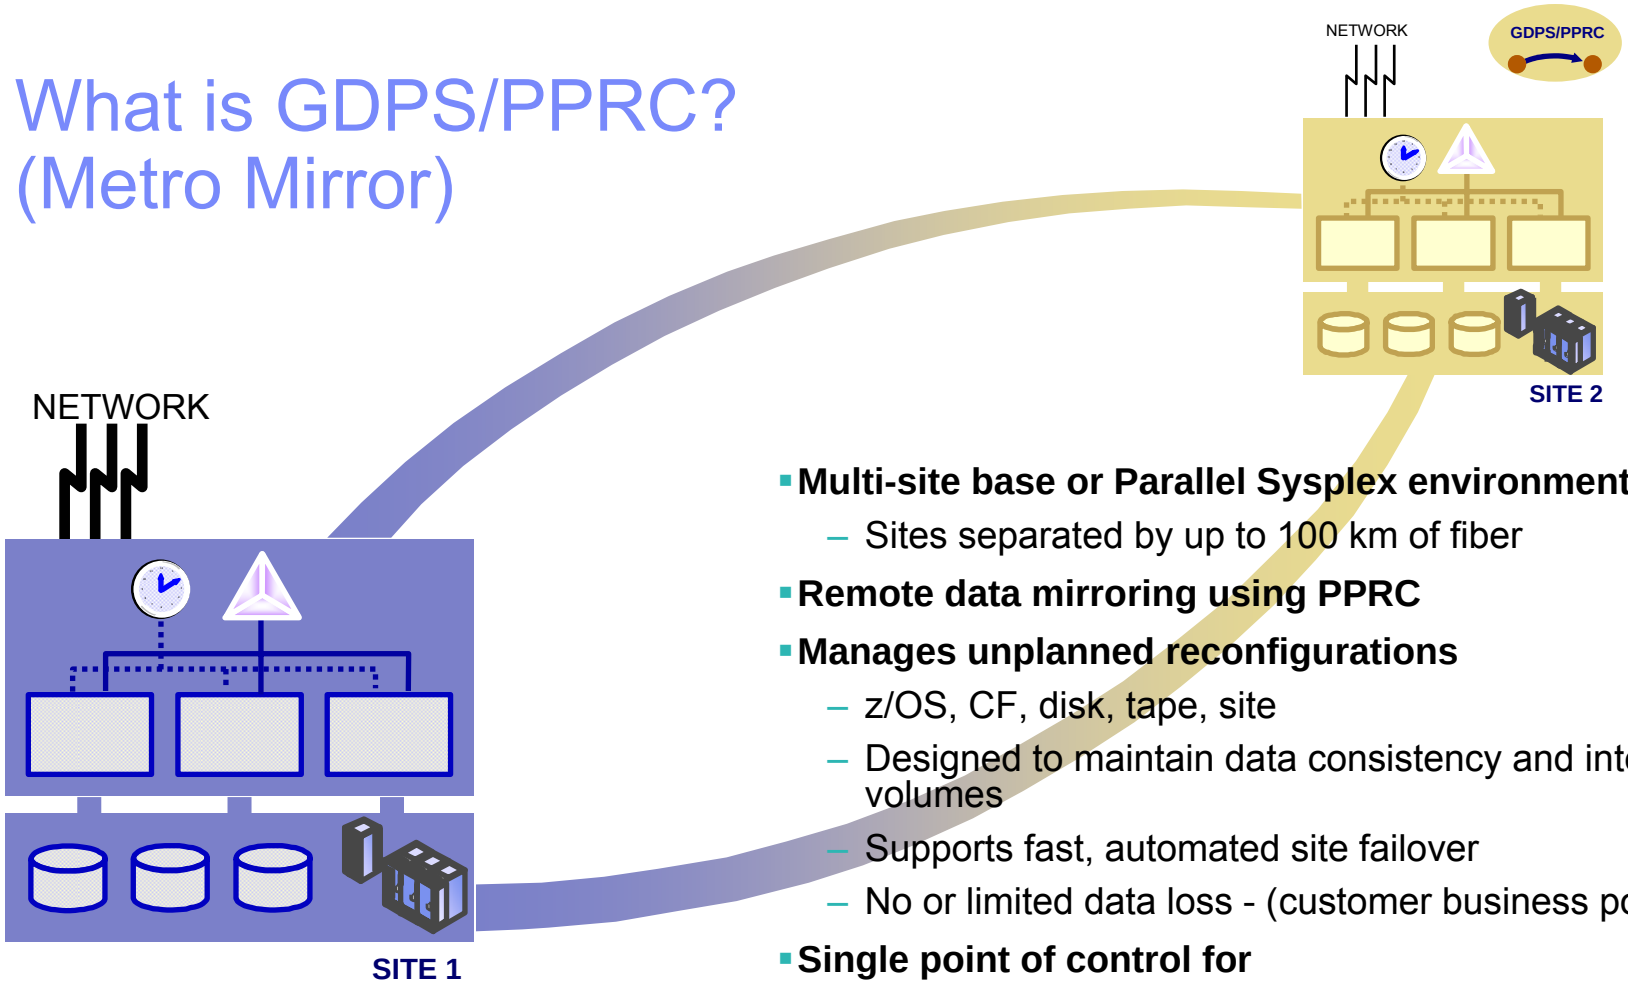

- **Multi-site base or Parallel Sysplex environment**
  - Sites separated by up to 100 km of fiber
- **Remote data mirroring using PPRC**
- **Manages unplanned reconfigurations**
  - z/OS, CF, disk, tape, site
  - Designed to maintain data consistency and integrity across all volumes
  - Supports fast, 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

Unplanned Cross-Site Reconfiguration with HyperSwap

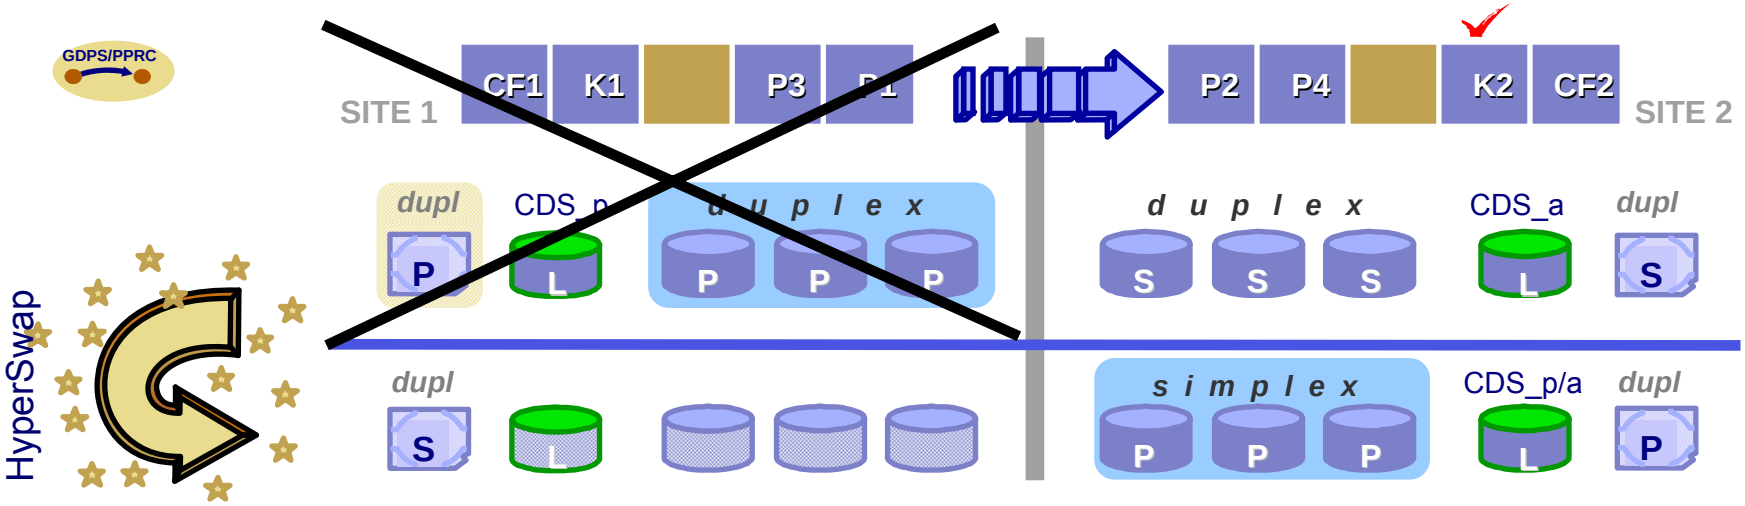

Autonomic processing

- HyperSwap disk configuration (swap prim/sec PPRC volume UCBs)

Execute SWAPSITE1 - runs automatically

- SYSPLEX = 'CDS SITE2' - move Couple Data Sets to Site2
- IPLTYPE = 'P1 MODE=SITE2' (repeat statements for P2, P3)
- TAPE = 'SWITCH SITE2' - switch PtPVTS configuration to have Site2 as the primary location
- TAPE = 'STOP SECONDARY' - suspend the PtPVTS processing

Execute takeover script due to P1 and P3 failure detection - initiated by customer

- SYSPLEX = 'CFRECOVER UNCOND' - Switch to site 2 CFRM policy
- CBU='ACTIVATE CPC=CPC2A' - activate CBU capacity on the CPC in Site2
- Restart workload running on P2 and P4
- Restart failed workload running on P1 and P3 to release resources
- IPLTYPE='P1 ABNORMAL' – Point P1 to its abnormal location (repeat for P3)
- SYSPLEX='LOAD P1' - IPL P1 in its abnormal location using Site2 DASD. (repeat for P3)

P2 and P4 remain active throughout procedure

Planned Site Shutdown with HyperSwap

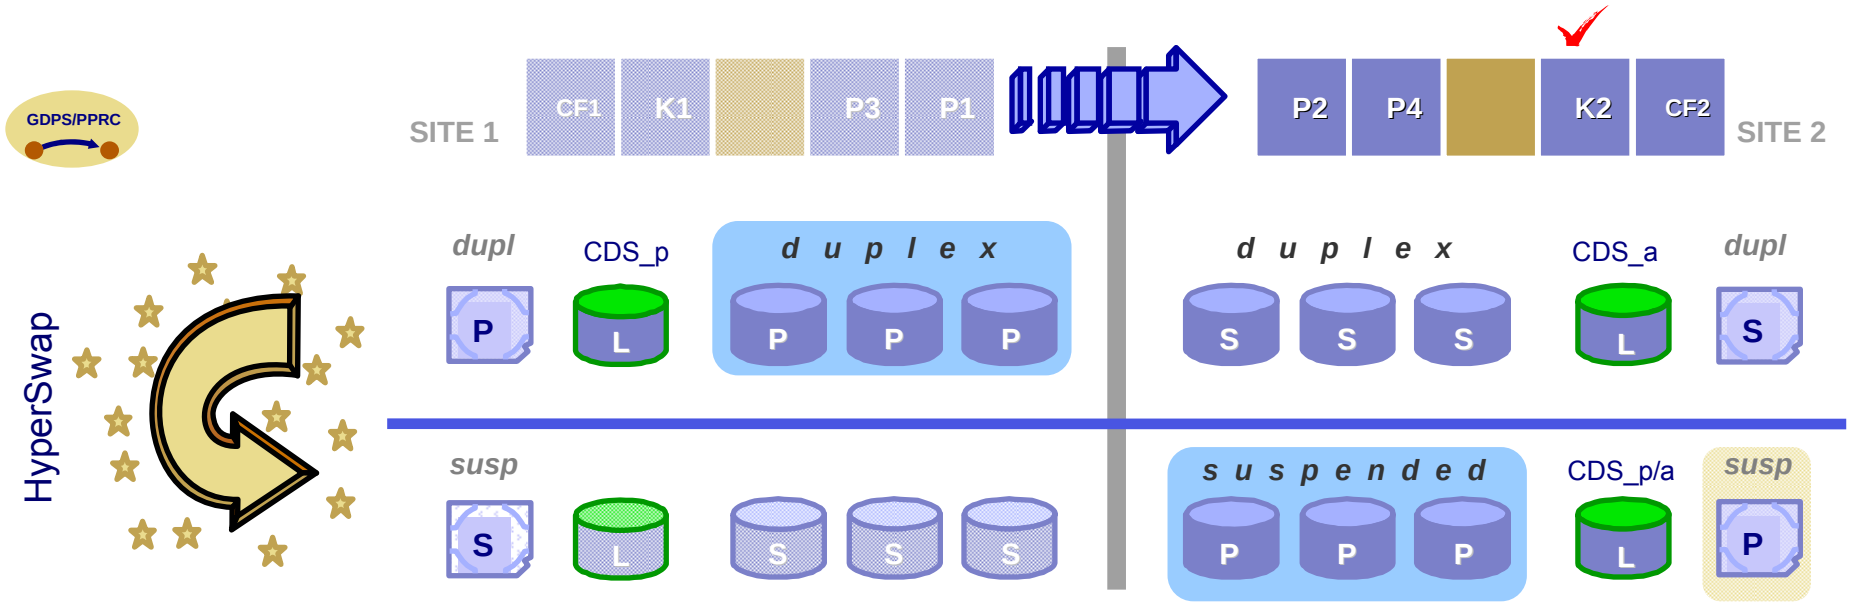

Site1 Maintenance Control Script

- SYSPLEX = 'STOP SITE1' - *orderly shutdown of all Site1 systems*
- SYSPLEX = 'CDS SITE2' - *move Couple Data Sets to Site2*
- SYSPLEX = 'CF SITE2' - *Move all structures to Site2 CFs*
- DASD = 'SWITCH HYPERSWAP SUSPEND' - *swap PPRC primary and secondary volumes; suspend PPRC*
- TAPE = 'SWITCH SITE2' - *switch PtPVTS configuration to have Site2 as the primary location*
- TAPE = 'STOP SECONDARY' - *suspend the PtPVTS processing*
- IPLTYPE = 'P1 ABNORMAL' - *point P1 to its abnormal LPAR for subsequent IPLs (repeat statement for P3)*
- SYSPLEX = 'ACTIVATE P1 LPAR' - *activate the backup LPAR for P1 in Site2 (repeat statement for P3)*
- IPLTYPE = 'P1 MODE=SITE2' *Point P1 to Site 2 disks for subsequent IPL (repeat statements for P2, P3, P4)*
- SYSPLEX = 'LOAD P1' - *load P1 into its backup LPAR in Site2 (repeat statement for P3)*

  Note: Since P2 and P4 were never stopped, they do not need to be re-IPLed.

GDPS/PPRC management of Open Systems LUNs

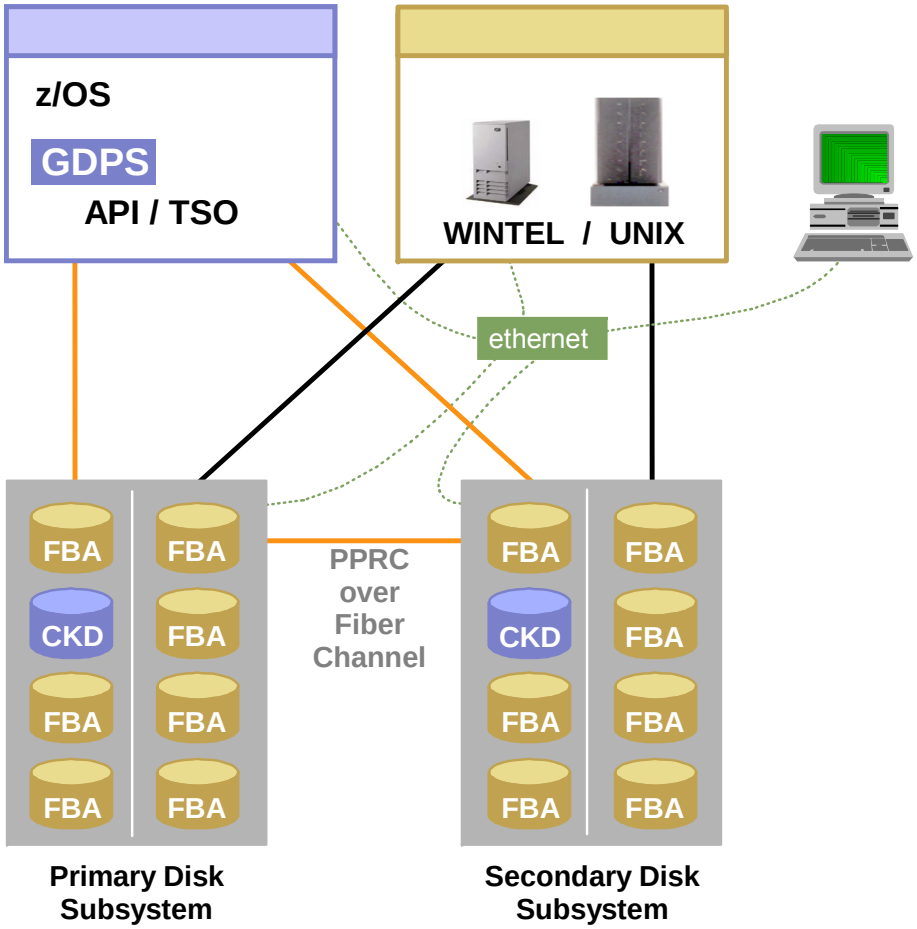

- Extends GDPS/PPRC technology to manage distributed applications across multiple platforms
  - z/OS and open systems data (Unix, NT, Linux)
- GDPS/PPRC running in a z/OS system manages the PPRC status of devices that belong to the other platforms
- Provides data consistency across both z/OS and/or open 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

GDPS/PPRC Multi Platform Resiliency for zSeries

New

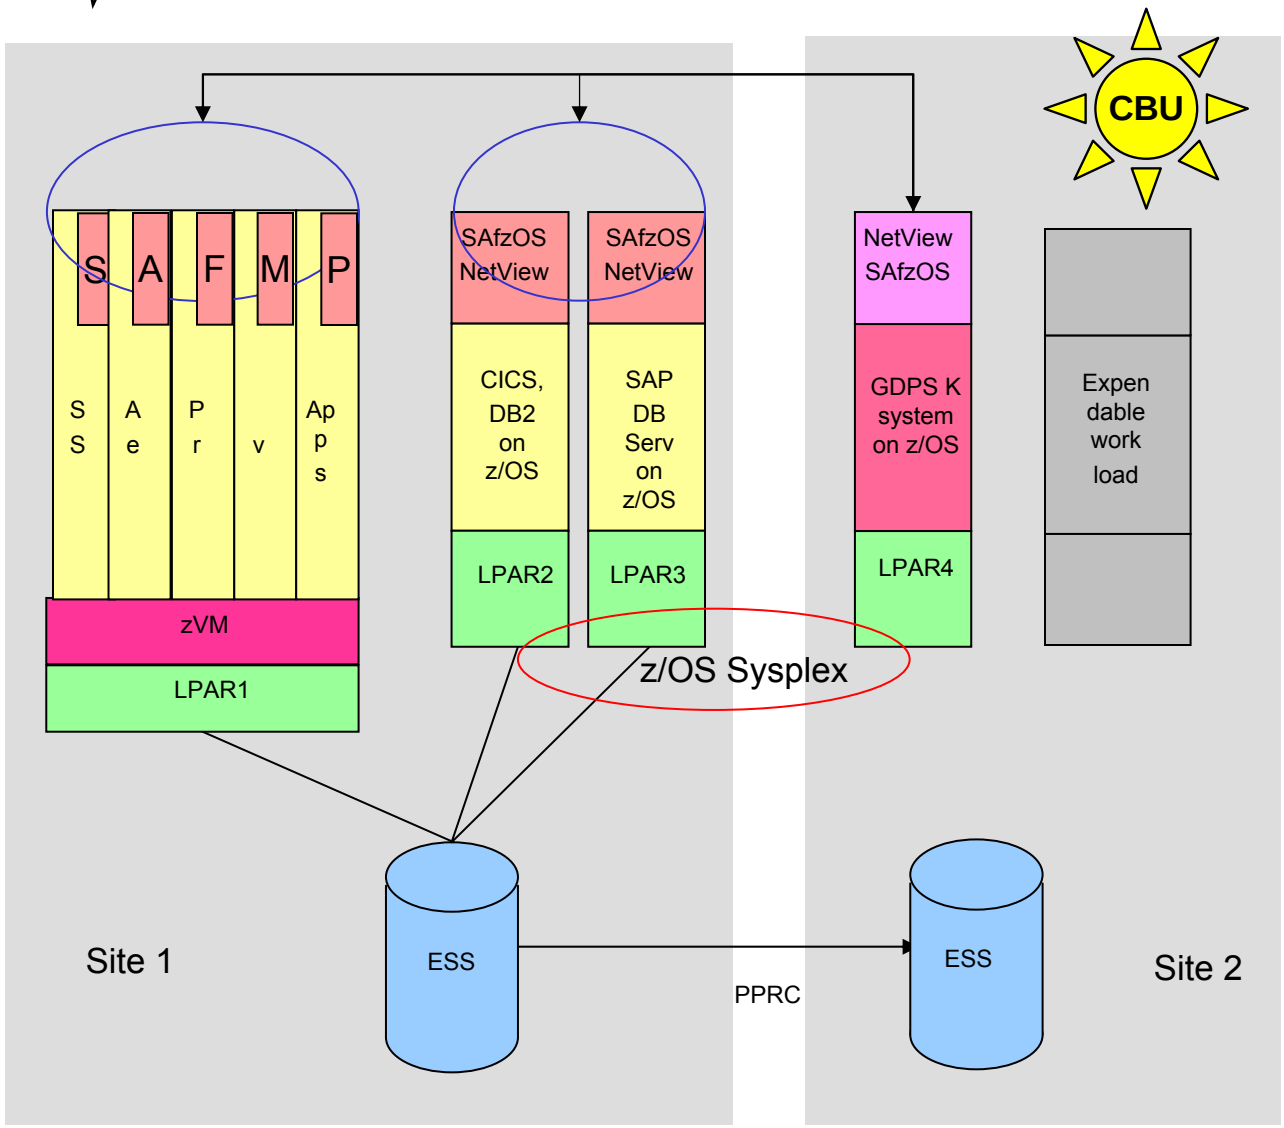

- **Valuable to customers with distributed applications**
  - SAP application server running on Linux on zSeries;
  - SAP DB server running on z/OS
  - etc.
- **Coordinated near-continuous availability and DR solution for z/OS and Linux guests running under z/VM**
- **GDPS exploits z/VM HyperSwap function to switch to secondary disks mirrored by PPRC**

GDPS/PPRC Multi Platform Resiliency for zSeries – Site 1 Failure

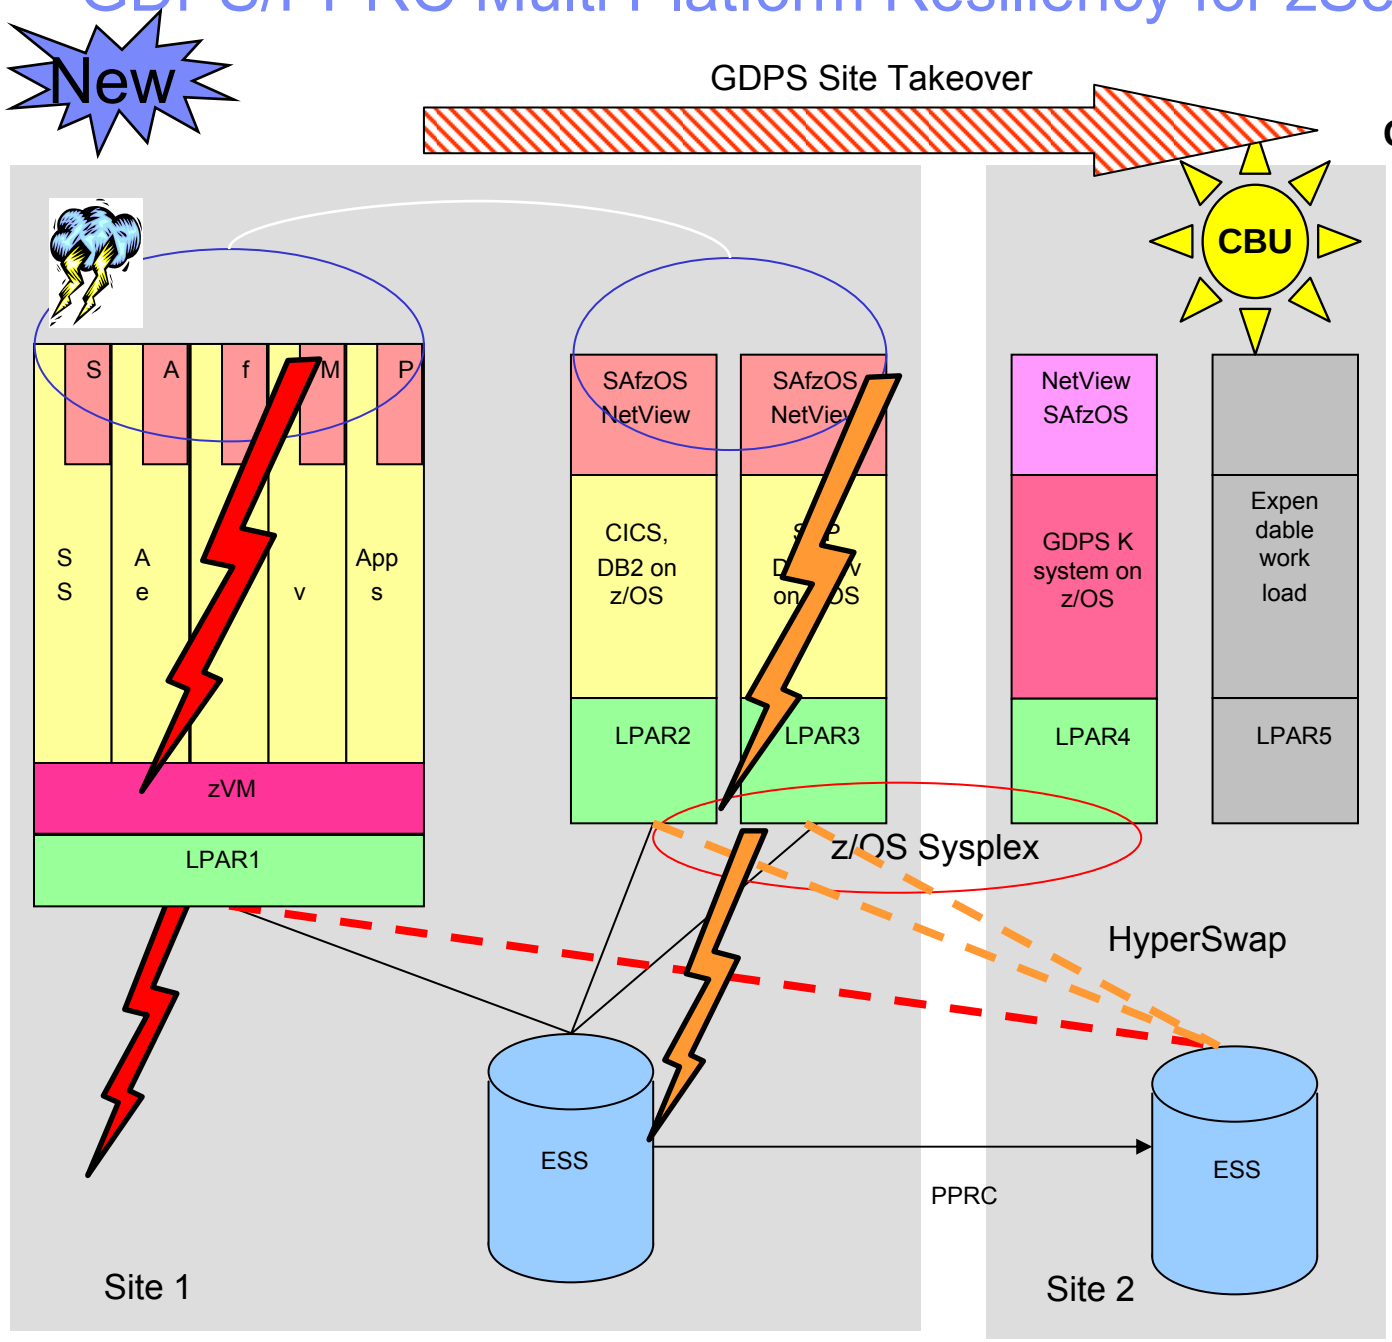

GDPS/PPRC provides planned and unplanned reconfiguration capabilities for Linux on zSeries and z/OS

- Unplanned site takeover triggered by Linux on zSeries or z/OS
  - Coordinated HyperSwap across both z/OS and z/VM disks
  - Stop expendable work in site 2 and/or invoke CBU
  - Restart Site 1 production systems in site 2
- Planned coordinated HyperSwap or site takeover also supported (control script)

IBM Systems and Technology Group University 2005

Unlimited Distance Disaster Recovery (2 sites)

What is GDPS/XRC

GDPS/XRC configuration – Site 1 Failure

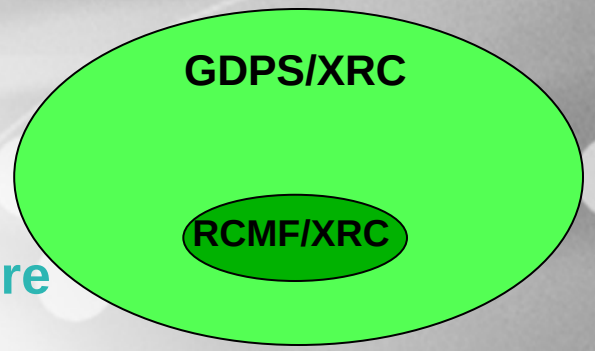

ON DEMAND BUSINESS™

What is GDPS/XRC ?
z/OS Global Mirror

- 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 System Data Movers

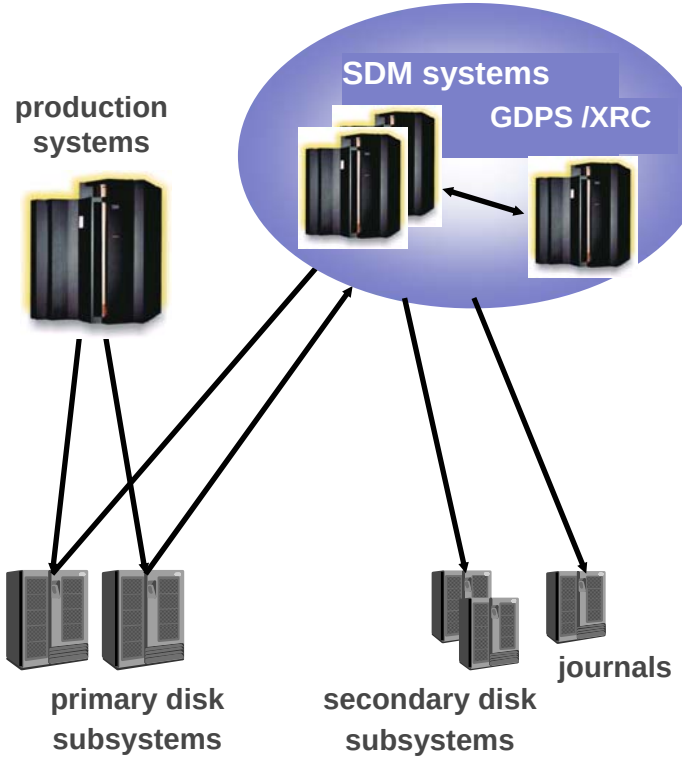

XRC manages secondary consistency
Across any number of primary subsystems
All writes time-stamped and sorted before committed to secondary devices

GDPS/XRC - Primary Site Failure

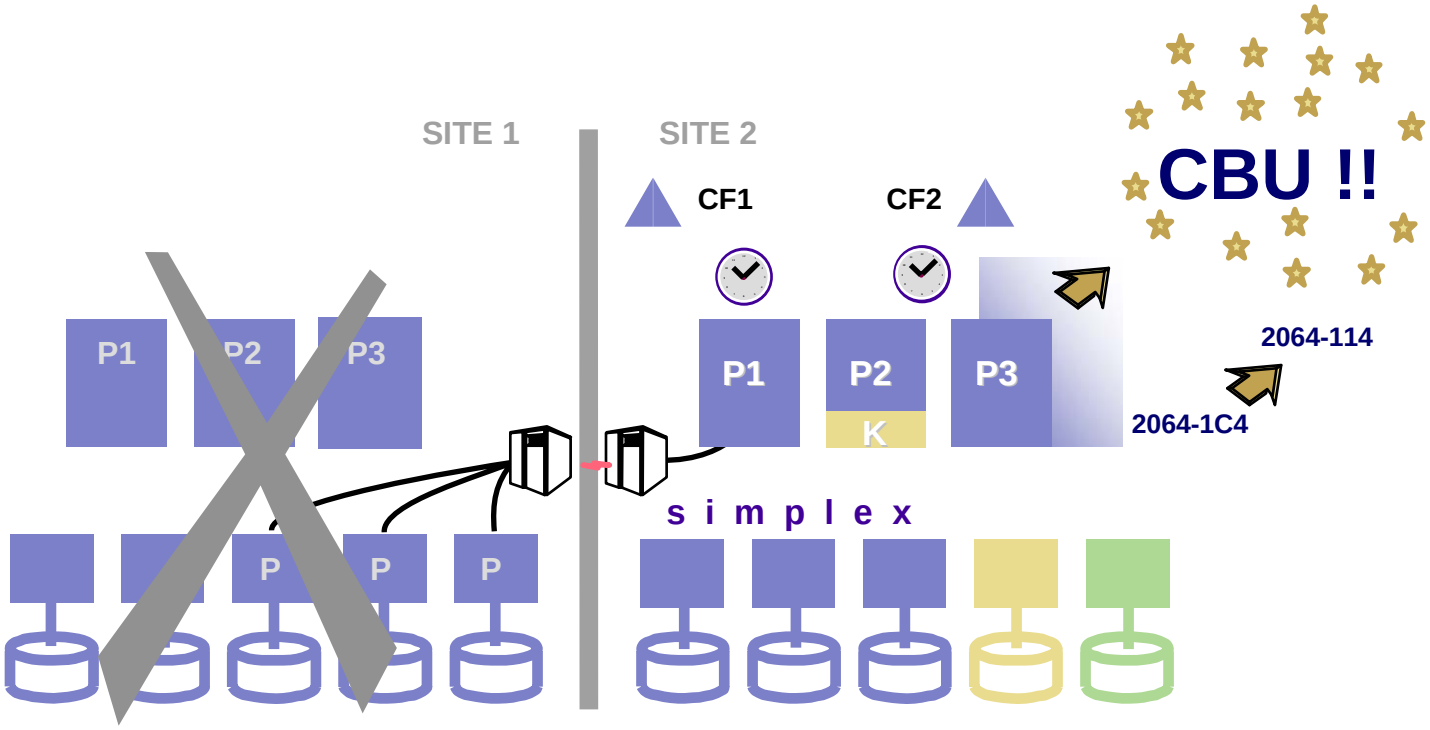

- Controlling System recovers secondary disks
- Stops expendable workload (SDM)
- Performs Capacity Backup, if applicable
- Restarts site 1 production systems in site 2

IBM Systems and Technology Group University 2005

Near-Continuous Availability and Disaster Recovery Solution (3 site)

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

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

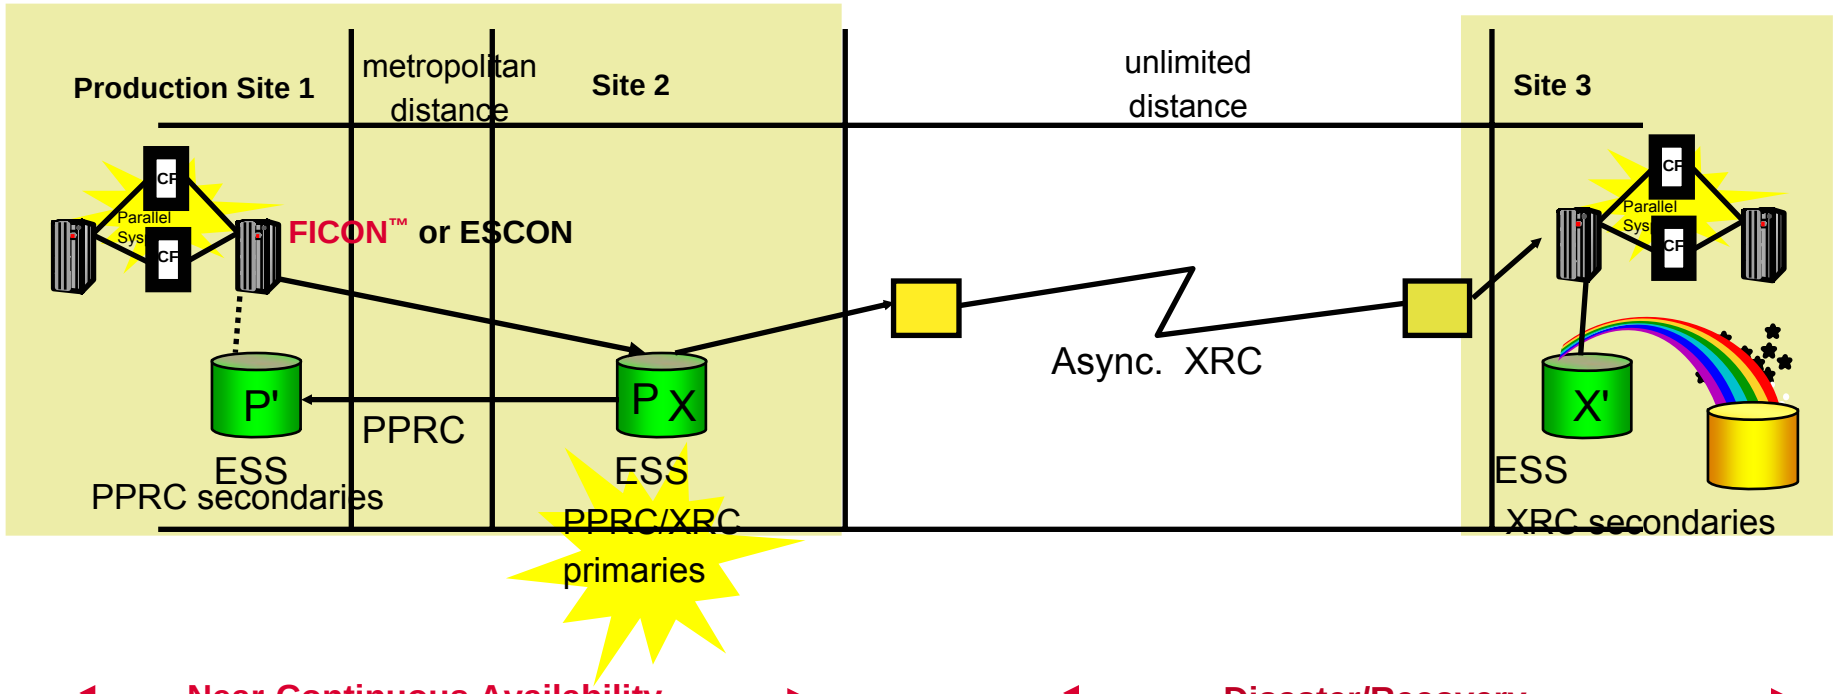

Near-Continuous Availability

- 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

Disaster/Recovery

- 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

Planned Enhancements

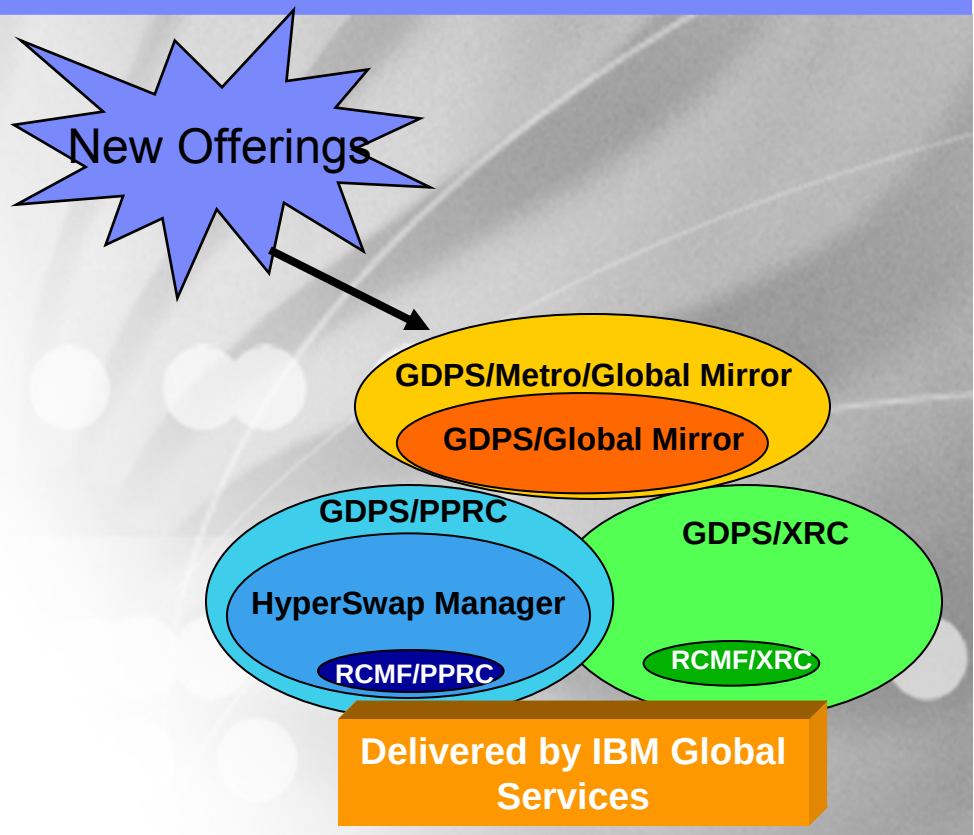

ON DEMAND BUSINESS™

Planned Enhancements

- GDPS/PPRC Multi Platform Resiliency for zSeries
  - Support for Linux for zSeries in native mode
- Unlimited distance Disaster Recovery for zSeries and Open data (2 sites)
  - Support for Global Mirror
- Unlimited distance CA/DR solution for zSeries and Open data (3 sites)
  - Support for Metro/Global Mirror
- Metropolitan Distance CA solution for site failures
  - CF Hint (exploitation of System Managed CF structure duplexing)

Global Mirror for ESS Configuration Requirements

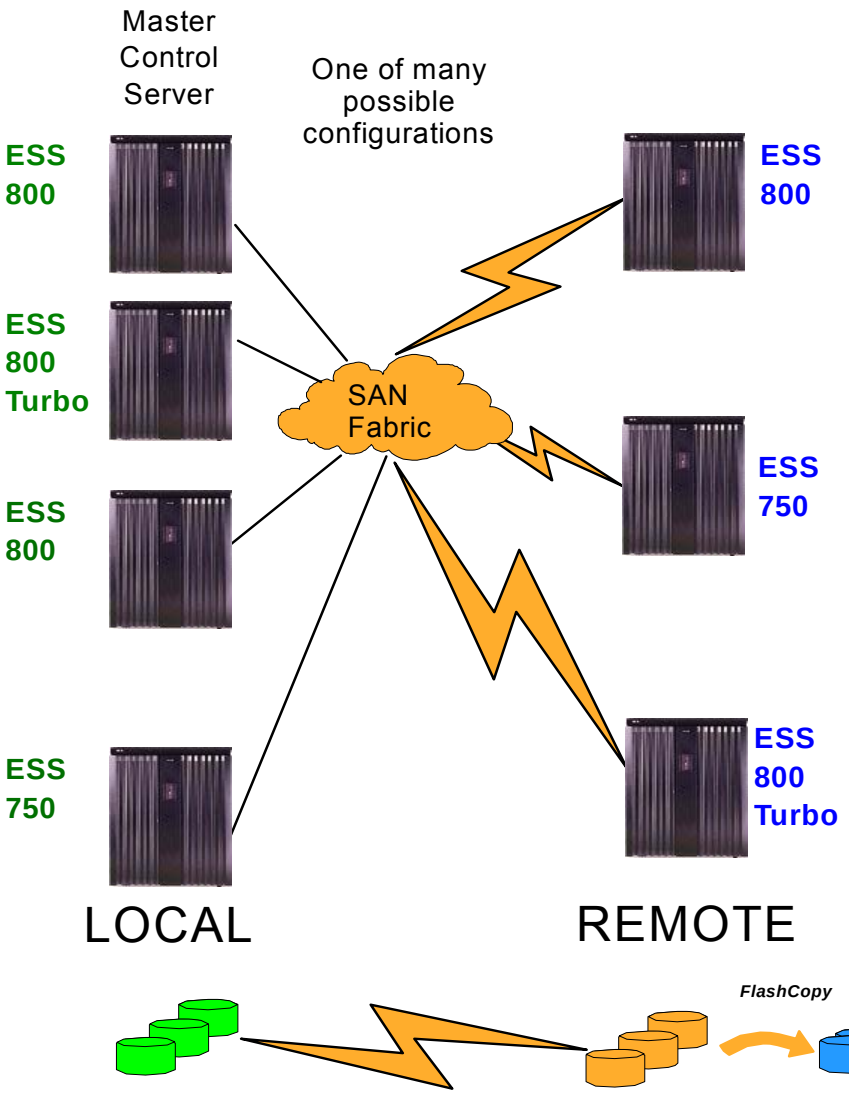

- Configuration Requirements
  - **Copy consistency:** managed autonomically by Master Control Server in master ESS
  - **Requires Fiber Channel paths:** between all ESS's in the Global Mirror session, for control commands and data
  - **Exploits Fibre Channel SAN infrastructure:** uses standard Fibre 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

GDPS/Global Mirror
Disaster recovery at virtually unlimited distance for zSeries and Open

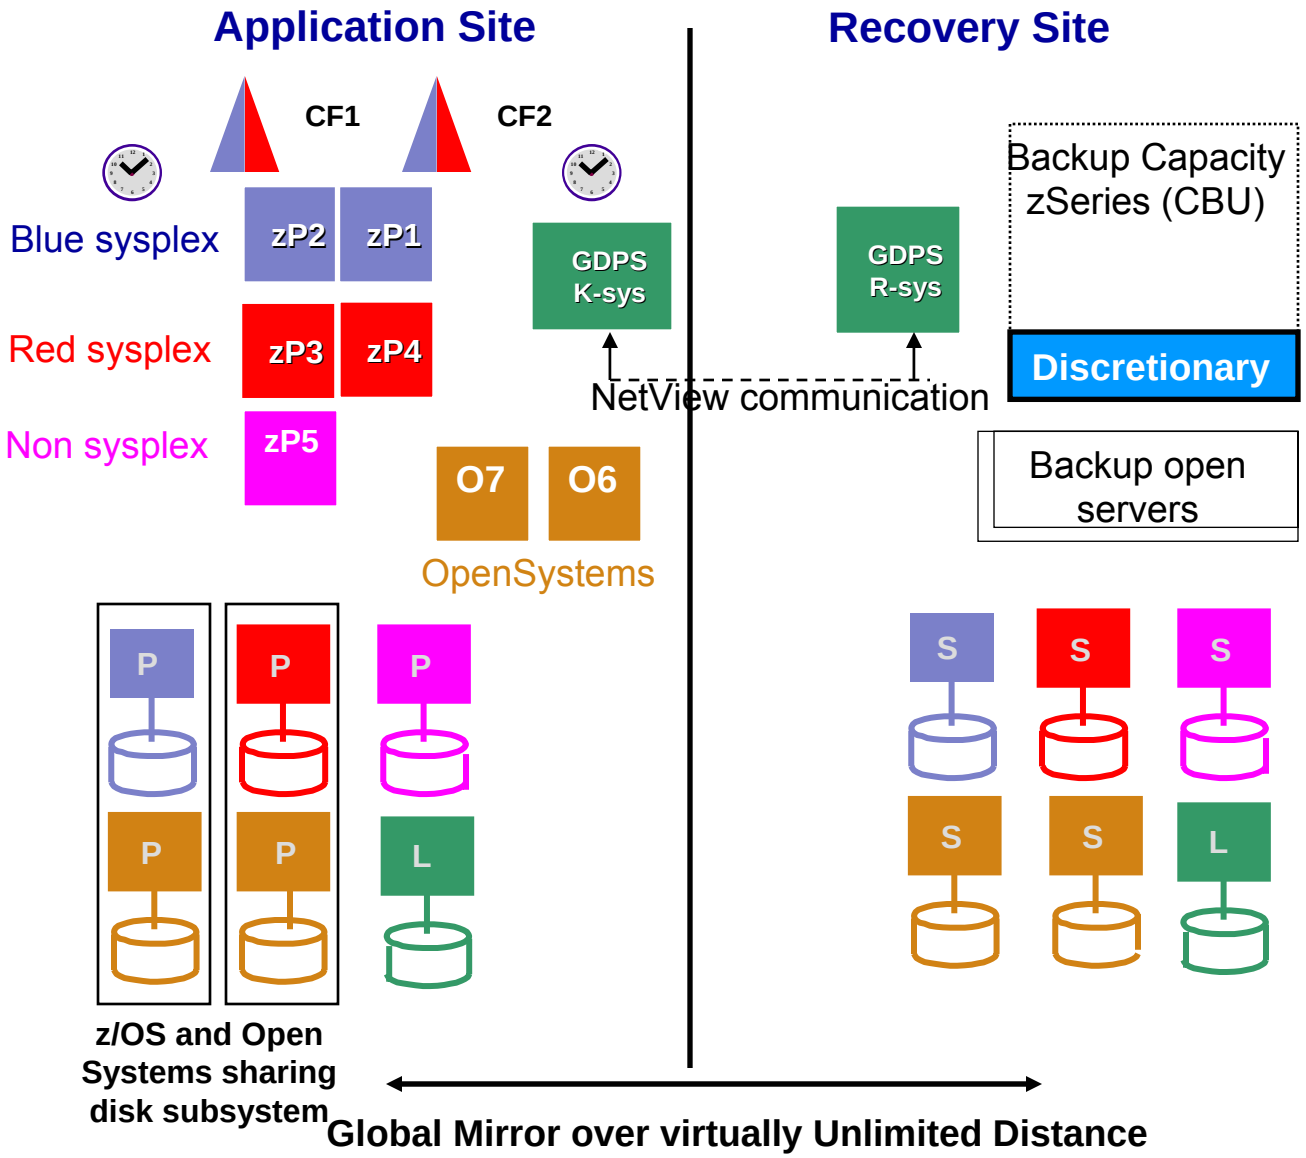

- Application site can have single z/OS Systems, Open Systems, Systems in a Sysplex
- z/OS and Open Systems data can be mirrored using Global Mirror
- GDPS manages multiple Global Mirror sessions
- Phase1 z/OS data
- Phase2 Open data

GDPS/Global Mirror - Site 1 Failure

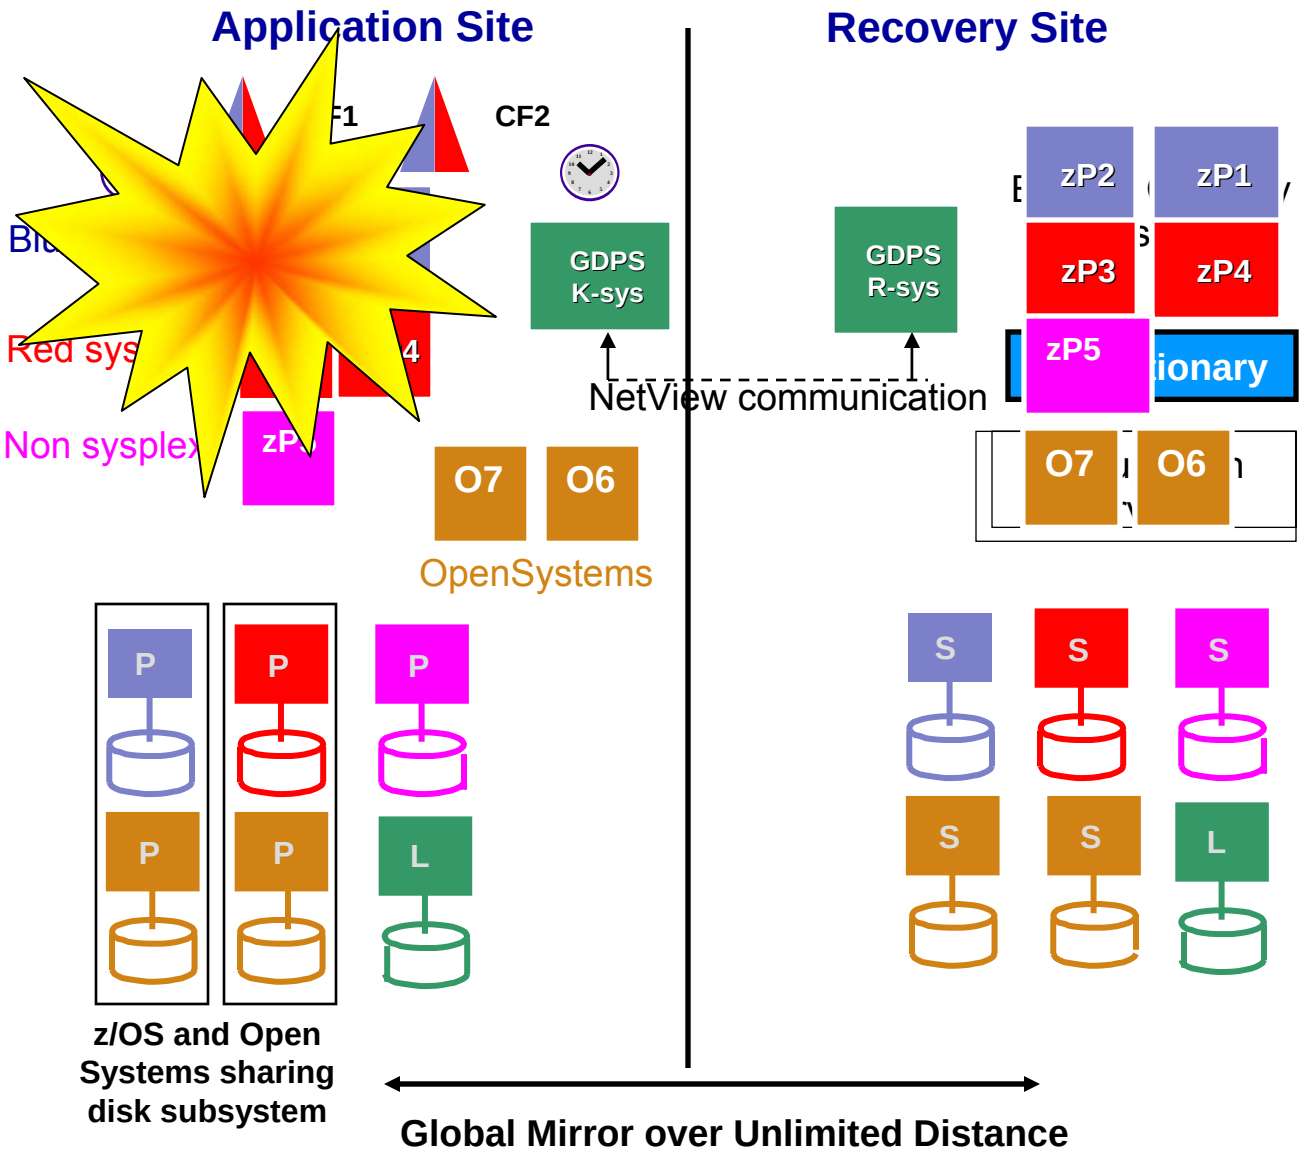

R-sys activities:
Secondary disk recovery, CBU activation, activate backup LPARs, IPLs systems.

- Target RTO < 2 hours
- RPO < 1 min
  - (depends on bandwidth)

GDPS Metro/Global Mirror - Solution for zSeries and Open Designed to provide Continuous Availability and Disaster Recovery at virtually unlimited distance

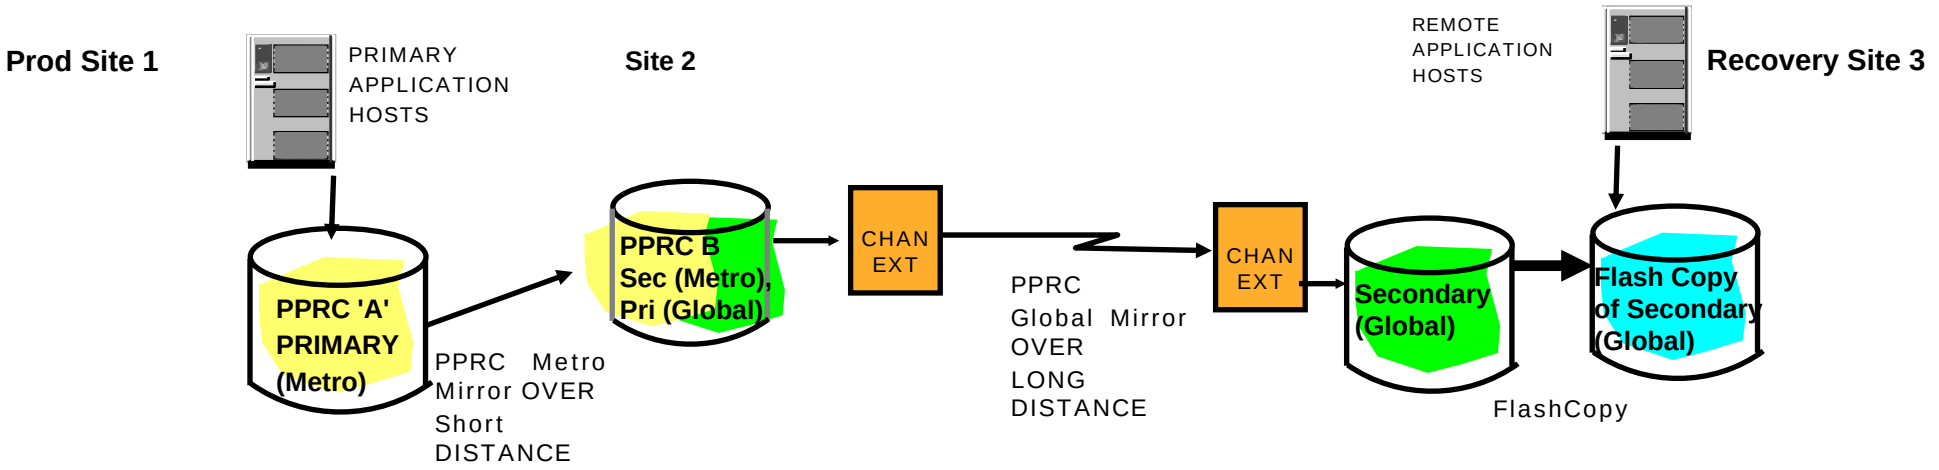

◄ **Near-Continuous Availability** ► ◄ **Disaster/Recovery** ►

- 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 (A)
- Site 1 and 2 failure
  - Site 3 can recover with minimal loss of data

GDPS Managed coordinated solution for zSeries and open systems

GDPS/PPRC 'CF Hint'

Should we keep Site1 or Site2 instance?

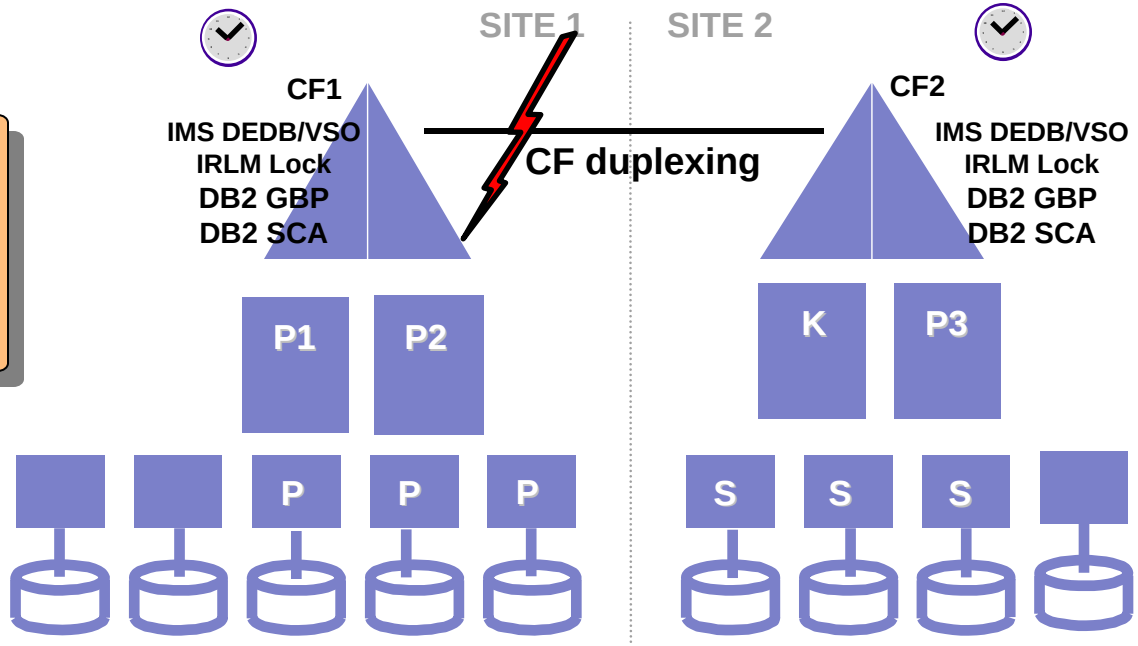

- System Managed CF structure duplexing required
- Freeze policy must be Freeze and STOP
- In the event of CF duplexing failure which instance should we keep?
- GDPS Controlling System will 'hint' that structures in the CFs in site where secondary disks are located should be used

No special recovery actions (eg GRECP) required
Facilitates faster application restart (improved RTO)
Provides consistent recovery time

Summary

- Flexible configuration options to meet a wide-range of Business Continuity requirements
  - Near-Continuous Availability of data within a single site – HyperSwap Manager
  - Solutions to handle distributed applications
    - GDPS/PPRC Open LUN Management
    - GDPS/PPRC Multi Platform Resiliency for zSeries
- In case of disaster
  - Designed to enable data consistency and integrity
    - **No data loss (GDPS/PPRC) or**
    - **Minimal data loss (GDPS/XRC, GDPS/GM)**
  - Offers prompt, responsive disaster recovery through end-to-end automation
- GDPS/PPRC addresses both Disaster Recovery and Near-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 functions will continue to be enhanced to support the On Demand Business Resiliency Framework

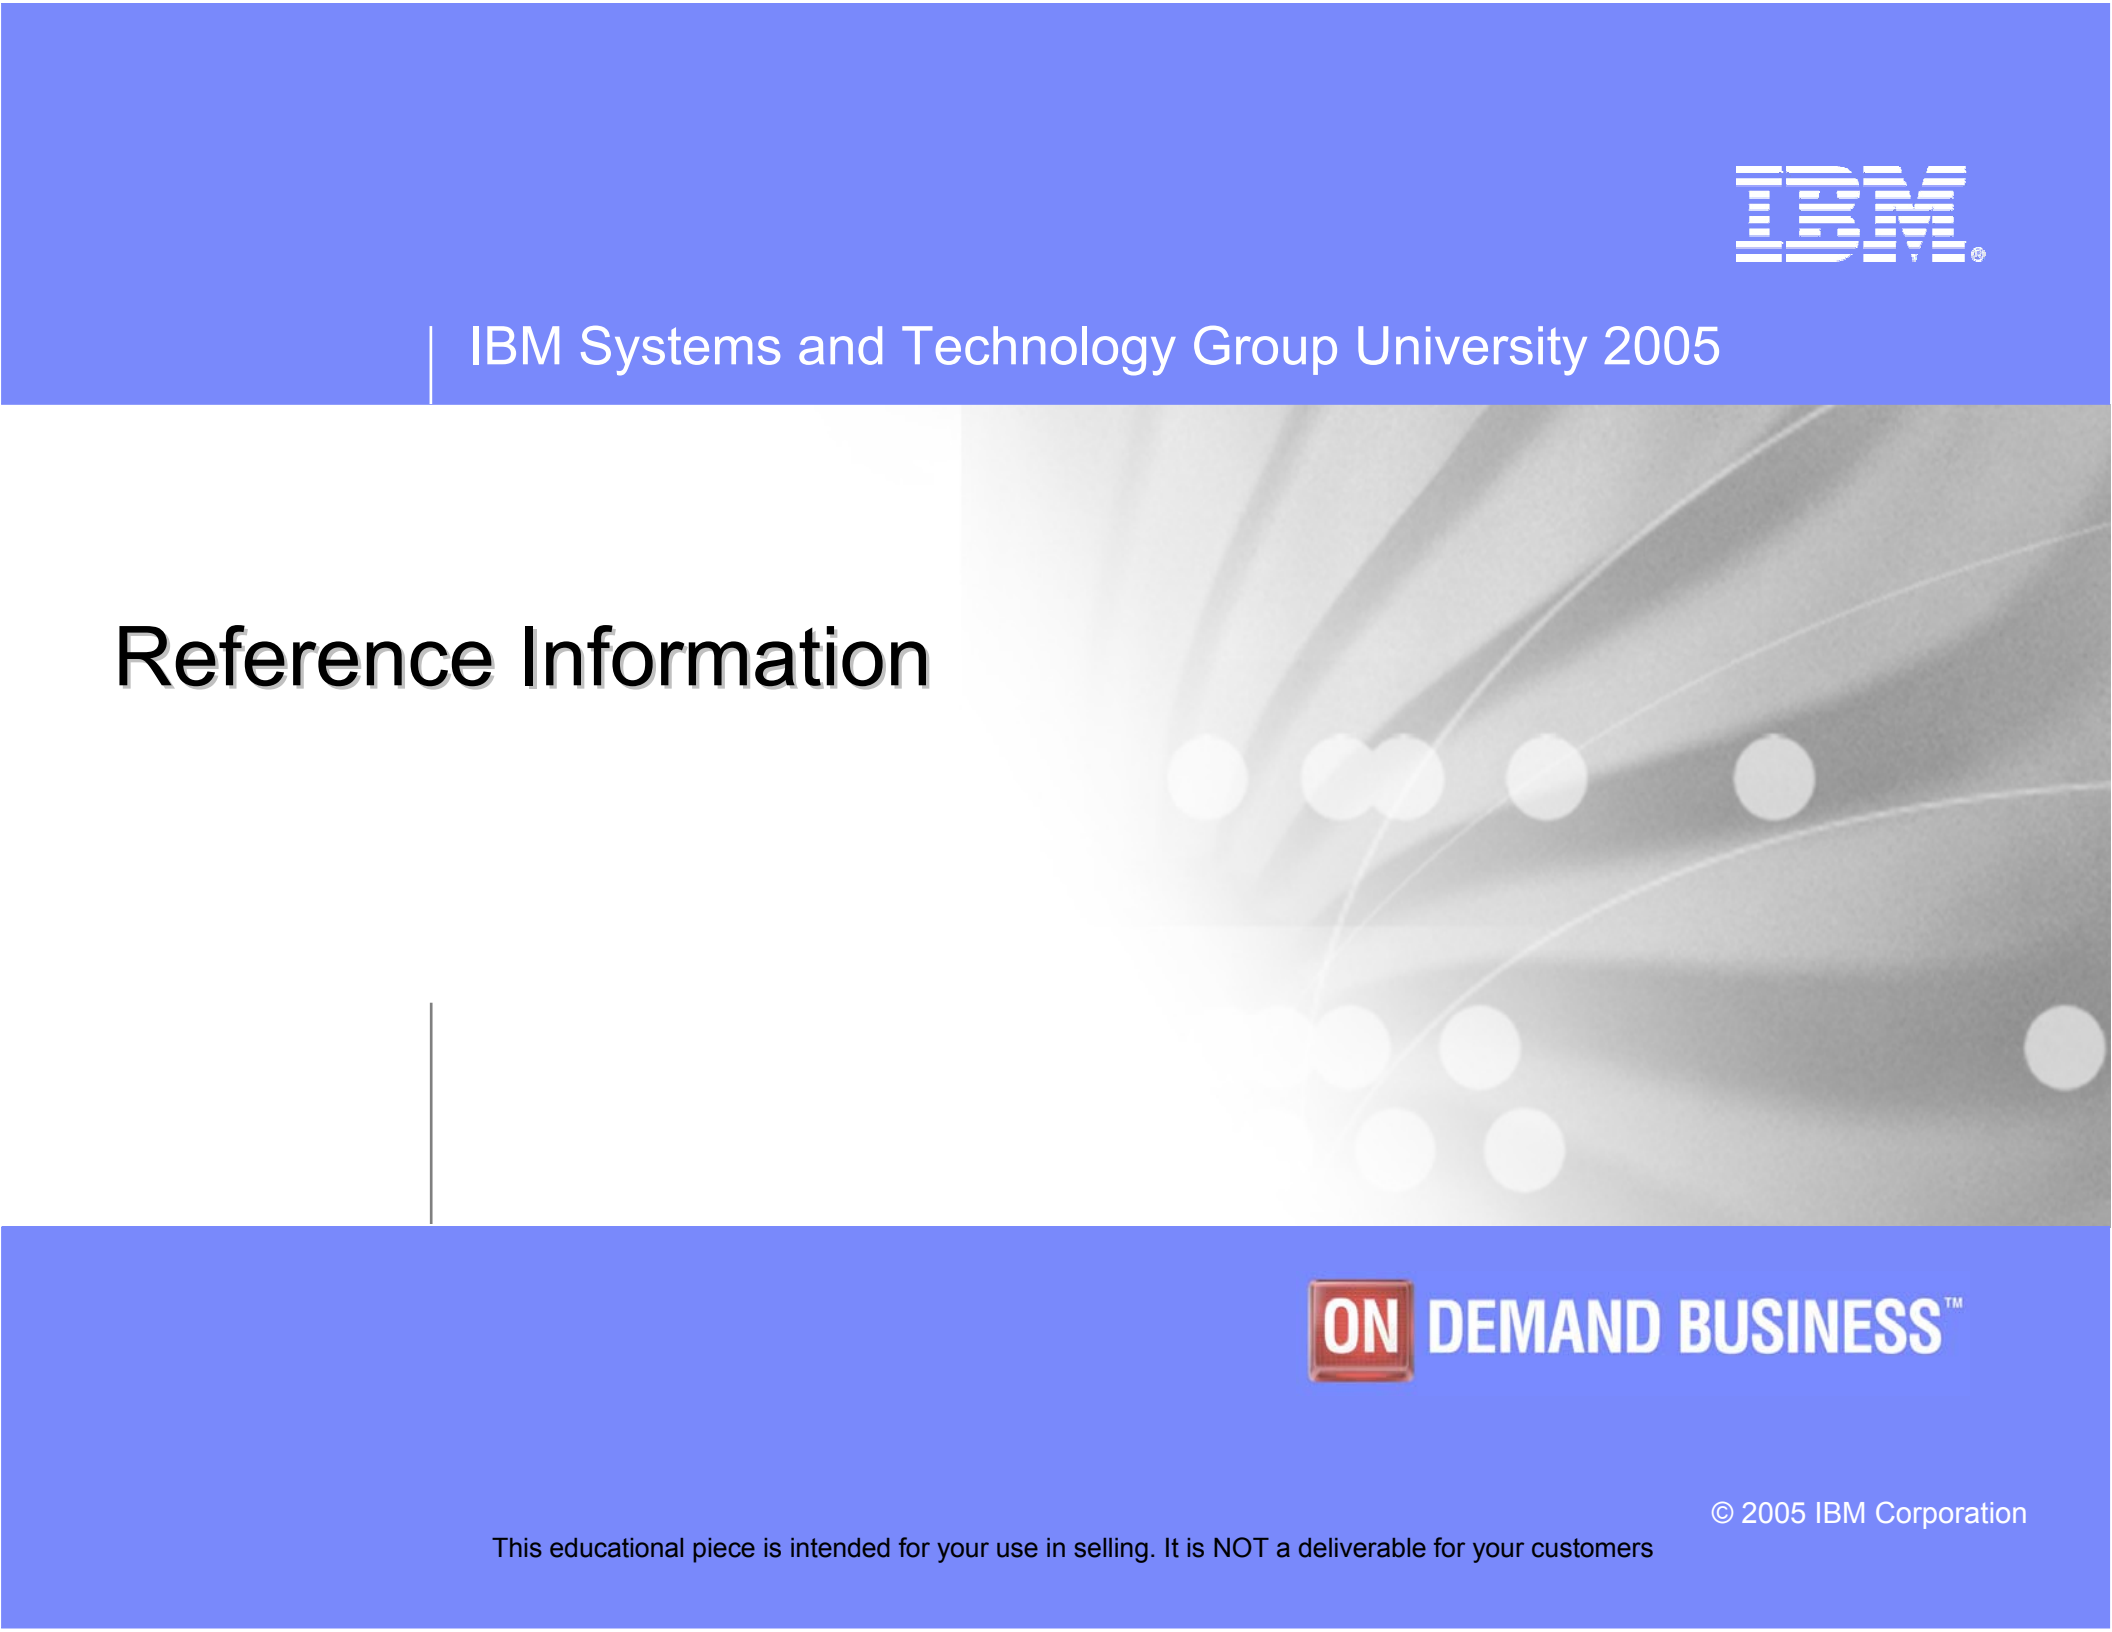

Business Continuity Services Offerings

- GDPS Technical Consulting Workshop (TCW)
  - **Designed to determine that the GDPS Availability & Recovery solution can meet the Client's business requirements as they relate to near-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
- Web Pages References
  - GDPS Home Page
  - Installation Services for GDPS
  - Business Continuity and Recovery Services
  - ITS Global Brand Services – GDPS (Internal)
  - ITS GDPS Home (Internal)
  - Customer Reference data base
- TotalStorage Business Continuance Sales Kit
- 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

GDPS ITS Contacts Worldwide

Americas	
USA -Financial Sector	Brian Teichman/Dallas/IBM@IBMUS
USA -Distribution Sector	George Westerman/Southfield/IBM@IBMUS
USA -All other sectors	Robert Pitcole/Southfield/IBM@IBMUS
Canada	Joe Murphy/Markham/IBM@IBMCA
Latin America	Patricia Elizagaray/Ecuador/IBM@IBMEC

EMEA	
EMEA	Rudi Rauch/Germany/IBM@IBMDE
North Region	Carol Ashcroft/UK/IBM@IBMGB
West Region	Philippe Bienne/France/IBM@IBMFR
Nordics Region	Knud-Erik Skelmose/Denmark/IBM@IBMDK
South Region	Massimiliano Linetti/Italy/IBM@IBMIT
Central Region	Thomas Heizler/Germany/IBM

Asia Pacific	
AP	Sumiko Fukaya/Japan/IBM@IBMJP
ANZ	Robert Boucher/Australia/IBM@IBMAU
ASEAN	Hwee Kwang Tan/Singapore/IBM@IBMSG
GCG -Taiwan	Jeff Kuo/Taiwan/IBM@IBMTW
GCG -China	Tao Tao/China/IBM@IBMCN
GCG - Hong Kong	Sammy Kwok/Hong Kong/IBM@IBMHK
Japan	Hirotsugu Yamanaka/Japan/IBM@IBMJP
Korea	JongChul Choi/Korea/IBM@IBMKR

IBM Systems and Technology Group University 2005

Please Complete the Brief Survey on This Session

Your Feedback Helps Us Provide You Right Education

Appendix

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

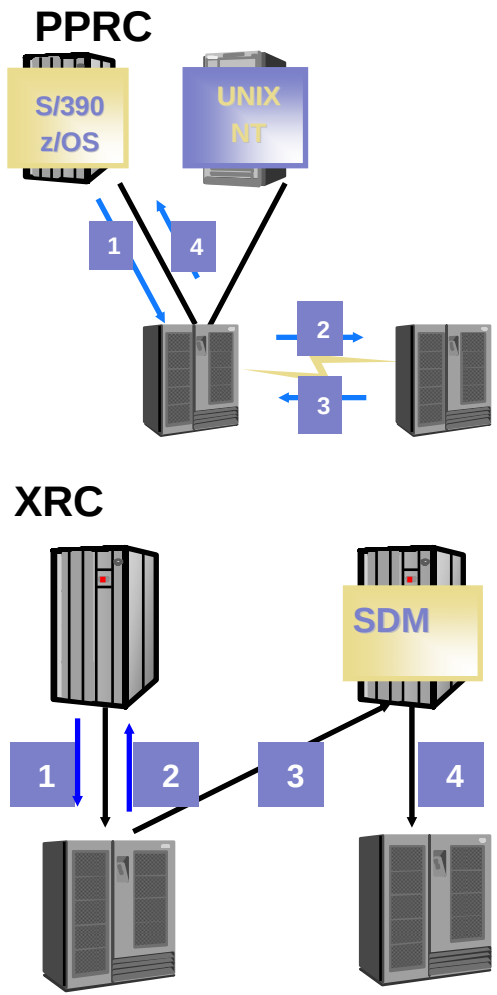

XRC Intersite Connectivity Options

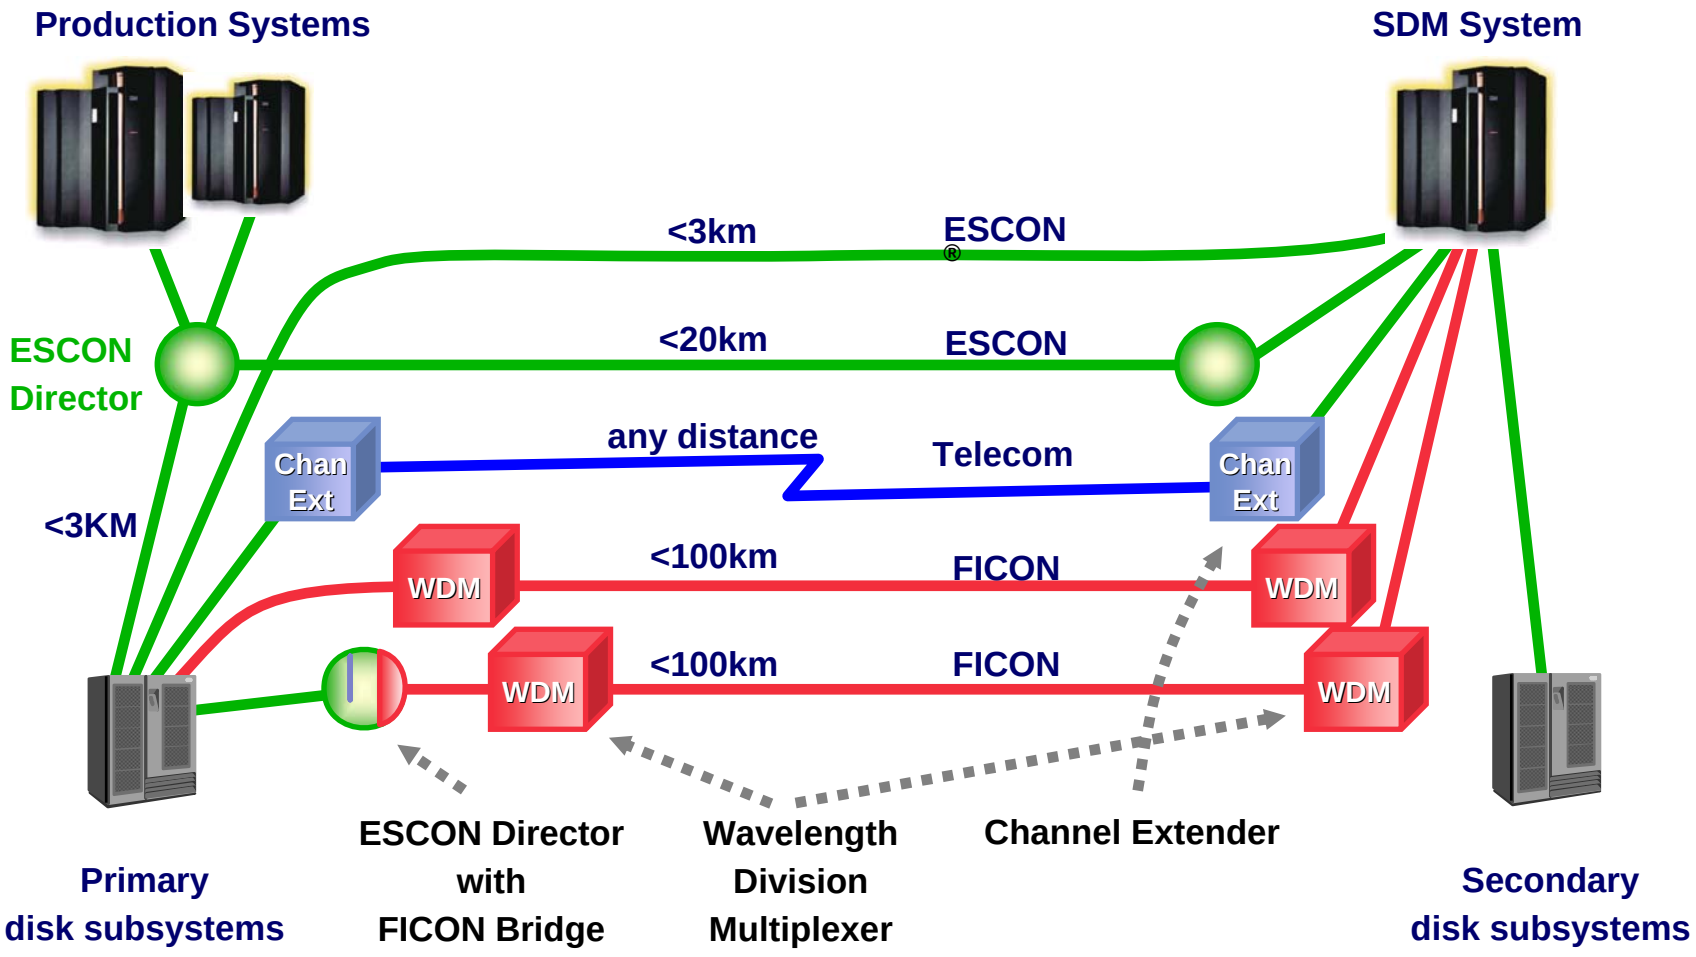

Failover / Failback

- Failover
  - Issue "failover" command to site B volumes
    - reverses direction of PPRC pairs
      - establish pairs (will succeed even if B>A paths down)
  - Start production on secondary site (B) volumes
  - Bit maps at secondary site (B) keep track of all changed tracks since failover
- Failback
  - Issue "failback" command to site B volumes
    - establishes B>A paths
    - resyncs B>A (full duplex)
  - Quiesce production IO to B
  - Remove B>A paths & establish A>B paths
  - Issue "failover" command to site A volumes
    - terminates B>A relationships
    - establishes A>B relationships
  - Issue "failback" command to site A volumes
  - Restart production IO on primary site volumes

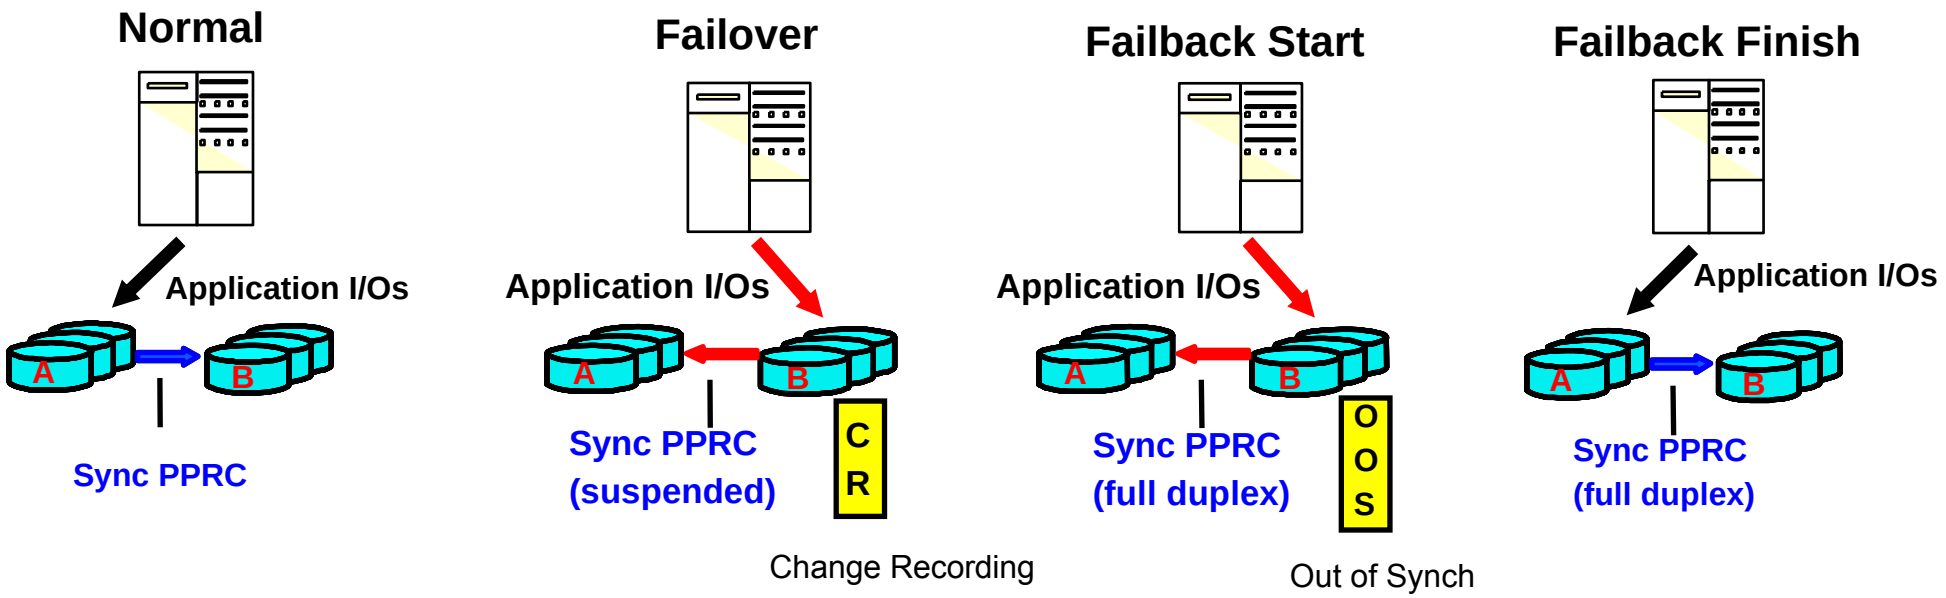

GDPS/PPRC

GDPS/PPRC across 100 km (example)

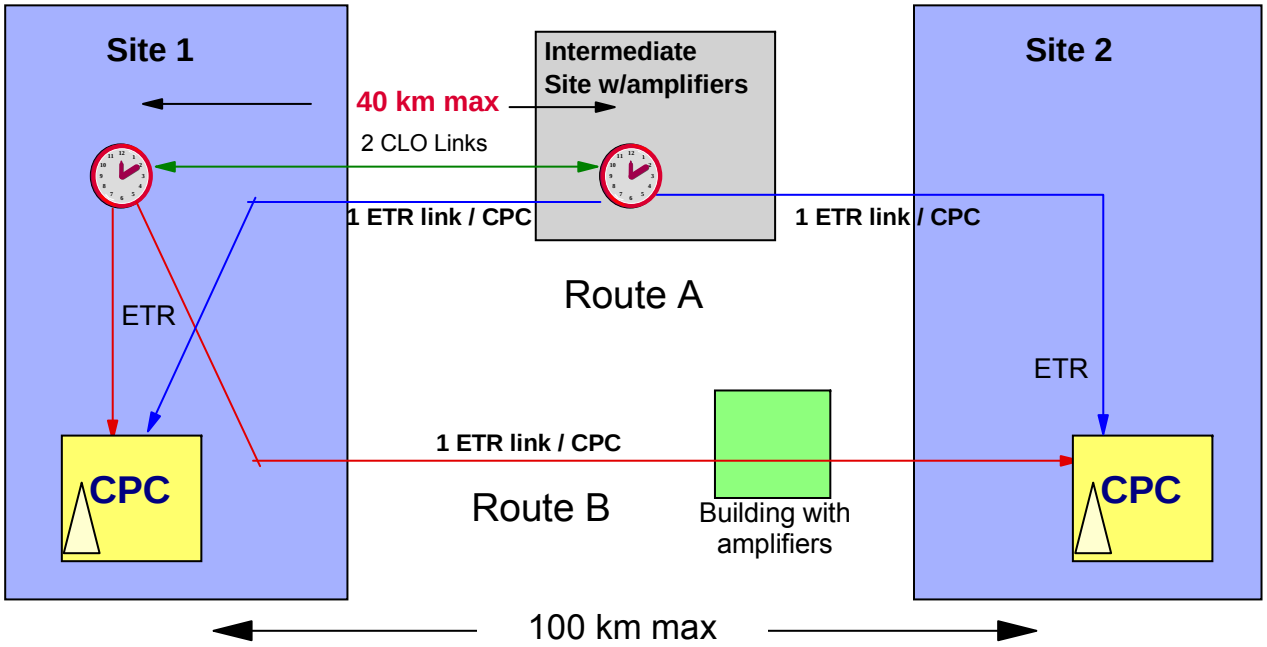

Redundant ISC3 cross site links from Site 1 to Site 2 also extended to up to 100 km max

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 enable 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

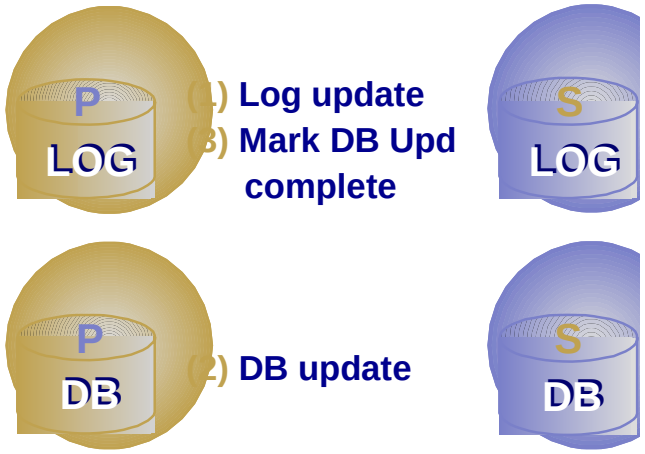

GDPS/PPRC

Unplanned Cross-site Reconfiguration

- Sysplex
- Disk subsystem
- Tape subsystem

FREEZE

Tape and disk subsystems
I/O consistent secondaries
- Across multiple CUs
- Multiple host systems

Customer business policies
- Freeze & Go, Freeze & Stop, Freeze & Stop Conditional
- Swap [GO|STOP]

ANALYSIS

Present actual status to operator
Request takeover authorization

SITUATION MANAGEMENT

- Remove systems from Sysplex
- Perform disk reconfiguration
- Perform tape reconfiguration
- Perform CBU activation
- Perform CF reconfiguration
- Perform CDS reconfiguration
- Acquire processing resources and IPL systems into Sysplex
- Initiate application startup

HyperSwap allows production systems to remain active
IBM provides a complete solution !

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
LINUX is a registered trademark of Linux Torvalds
UNIX is a registered trademark of The Open Group 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.
Intel is a registered trademark of Intel Corporation
*** 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.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.