GDPS® Overview & Recent Enhancements (Release 3.7 and 3.8)

Alan McClure
IBM Corporation

August 10, 2011
Session 9636

Trademarks

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

IBM*	ESCON*	Redbooks*
IBM (logo)*	FlashCopy*	Sysplex Timer*
Ibm.com*	GDPS*	System p*
AIX*	HyperSwap	System z*
DB2*	IBM*	Tivoli*
DS6000	IBM logo*	z/OS*
DS8000	Parallel Sysplex*	z/VM*
Dynamic Infrastructure*	POWER5	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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.

Agenda

- GDPS® Overview
- Enhancements (V3R8 GA and V3R7 SPE)
  - Availability
  - Scalability
  - Simplified system management
  - New tools
- Summary and Miscellaneous

What is GDPS® ?

- Manages application and data availability in and across sites
  - Monitors systems, disk & tape subsystems
  - Manages planned and unplanned activities
    System/disk maintenance / failure or Site maintenance / failure
- Integrated / Automated solution
- Builds on proven high availability technologies
  - Clustering
  - Remote copy (disk and tape)
  - Automation
- Easy to use interface
  - Intuitive GUI / panel interfaces
  - Policy based commands

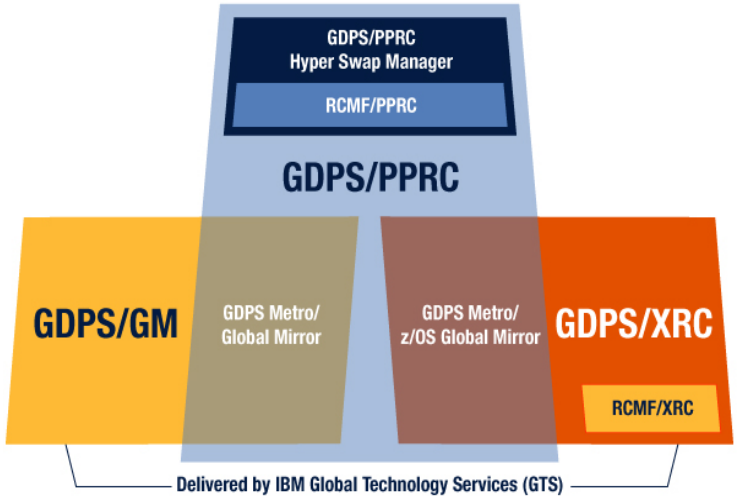

Continuous Availability & Disaster Recovery

- Mature
  - Nearly 600 implementations world-wide since 1998
  - Many customer references
- Flexible
  - Synchronous and/or Asynchronous remote copy
  - 1, 2, or 3 site
  - Easily customized automation
  - Multi-vendor support

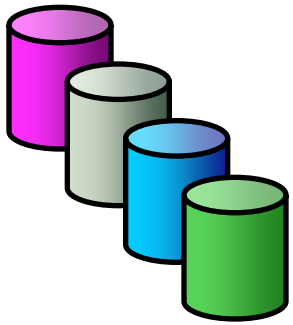

The GDPS® Family of Offerings

GDPS/PPRC HM	GDPS/PPRC	GDPS/GM & GDPS/XRC	GDPS/MGM & GDPS/MzGM	GDPS/Active-Active
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**	**Continuous Availability, Disaster Recovery, and Cross-site Workload Balancing at Extended Distance**
Single Data Center **Applications remain active** **Continuous access to data in the event of a storage subsystem outage** 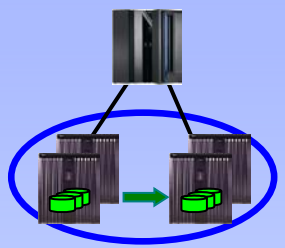 **RPO=0 & RTO=0**	**Two Data Centers** **Systems remain active** **Multi-site workloads can withstand site and/or storage failures** 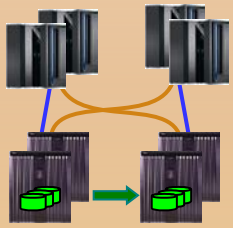 **A/S RPO=0 & RTO<1 hr or** **A/A RPO=0 & RTO mins**	**Two Data Centers** **Rapid Systems Disaster Recovery with "seconds" of Data Loss** **Disaster recovery for out of region interruptions** 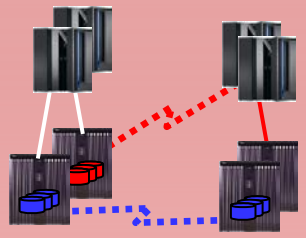 **RPO secs & RTO <1 hr**	**Three Data Centers** **High availability for site disasters** **Disaster recovery for regional disasters** 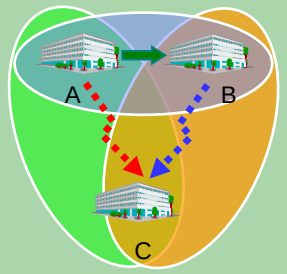 **A/S RPO=0 & RTO<1 hr or** **A/A RPO=0 & RTO mins** **and RPO secs & RTO <1 hr**	**Two or More Data Centers** **All sites active** 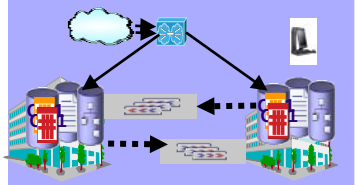 **RPO secs & RTO secs**

Components

Tivoli – NetView, SAz **STG** – System z, DS8K, PPRC **GTS** – GDPS control code, Services	**Tivoli** – NV, SAz, SA-MP, AppMan **STG** – System z, DS8K, VTS, PPRC **GTS** – GDPS control code, Services	**Tivoli** – NV, SAz **STG** – System z, DS8K, Global Mirror, XRC **GTS** – GDPS control code, Services	**Tivoli** – NV, SAz **STG** – System z, DS8K, MGM, MzGM **GTS** – GDPS control code, Services	**Tivoli** – SA, NetView **AIM -** Multi-site Workload Lifelife **IM -** DB2 &IMS replication **STG** – System z, DS8K, Global Copy **GTS** – GDPS control code, Services

RPO – recovery point objective RTO – recovery time objective

Heterogeneous Data/Platform Management

- Open LUN management (GDPS/PPRC, GDPS/GM)
- Multiplatform Resiliency for System z (GDPS/PPRC xDR)
  - z/VM & Guest (CKD disks)
  - Native Linux on z (CKD and FB disks)
- Distributed Cluster Manager (DCM)
  - VCS: GDPS/PPRC, GDPS/XRC, and GDPS/GM
  - SA AppMan: GDPS/PPRC

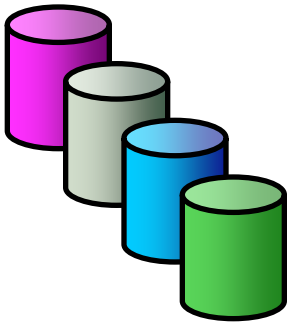

Agenda

- GDPS® Overview
- **Enhancements** (V3R8 GA and V3R7 SPE)
  - Availability
  - Scalability
  - Simplified system management
  - New tools
- Summary and Miscellaneous

Availability

- Simplified Freeze/Swap policy
- Reduced impact of false freezes
- STP recovery and network management
- xDR - HyperSwap support for fixed block disk
- xDR - Extended monitor function
- CPC / LPAR awareness for reset
- IPL protection

Simplified Freeze/Swap Policy
GDPS/PPRC and GDPS/PPRC HM

- Splits Freeze policy into two policy statements for action on
  - Mirroring problem (Freeze trigger)
  - Disk problem (HyperSwap trigger)

- Allows consistent handling of triggers based on policy
- Resultant actions easier to understand

- Freeze trigger specification

GO	Continue to update the primary disks after Freeze
STOP	Quiesce production systems after Freeze
COND	Query determines status of secondary disks If secondary disk failure GO, else STOP

- GDPS V3.7 SPE (PM24170)

Increased potential for zero data loss with reduced risk of availability loss

Reduced impact of False Freezes
GDPS/PPRC and GDPS/PPRC HM

- False Freeze: PPRC is interrupted to secondary disk
  - Loss of connectivity to secondary disk
  - Secondary disk subsystem failure.
  - With “Freeze=STOP” option, GDPS initiates a takeover
- Primary disk can continue to be updated if not the start of rolling disaster
- New query function available in DS8700/ DS8800 microcode provides status of secondary disks.
- If secondary disks have failed then primary can continue to be updated.
- GDPS V3.7 SPE (PM24170)

Availability

STP and Sysplex Timer Support – 1
GDPS/PPRC, GDPS/PPRC HM

Existing support

- When server that runs K-sys becomes unsynchronized…
- z/OS can recognize if an LPAR is a K-sys

- Puts K-sys into local timing mode up to 80 minutes
  - Allows K-sys to complete any in progress Freeze
  - Facilitates situation analysis before replying to WTORs
    - Review status of PPRC and other resources.
    - Remove ‘guessing game’ about if safe and which systems to restart.

STP and Sysplex Timer Support – 1
GDPS/PPRC, GDPS/PPRC HM

Enhancement

- GDPS automates 'RETRY' response to sync WTORs on spinning systems
  - STP (IEA394A) and ETR (IEA015A)
  - Allows re-synchronization if server in an ETR or STP network
  - Recovers from unsynchronized to a synchronized state,
  - Recovery can be automatic or via operator-initiated reconfiguration.
  - Uses the BCP Internal Interface
- GDPS V3.7 SPE (PM10041)

Improved Recovery Time
Reduced Operator Intervention
Helps prevent subsystem and application time outs

STP and Sysplex Timer Support – 2
GDPS/PPRC, GDPS/PPRC HM

- Automatically reconfigure an STP Coordinated Timing Network (CTN)
  No need for Operations to manually use HMC panels
- Reply to WTOR as part of script / commands
- Planned actions
  - Deactivate of the server hosting the current time server
  - Move current time server to server on the other site
- Unplanned actions
  - Define a server to be the clock source in the recovery site

- Benefits
  - Improves recovery times
  - Simplifies operations

Improved RTO / Availability

Fixed Block disk support
GDPS/PPRC with xDR

- Native Linux systems using SCSI attached FB disk can now be managed by xDR.
- Builds on existing GDPS Open LUN management capability (PPRC and Freeze).
- Unplanned and planned HyperSwap capability
- xDR system management functions
  - Cluster, node and hardware LPAR management
  - Scripted/coordinated planned/unplanned event management

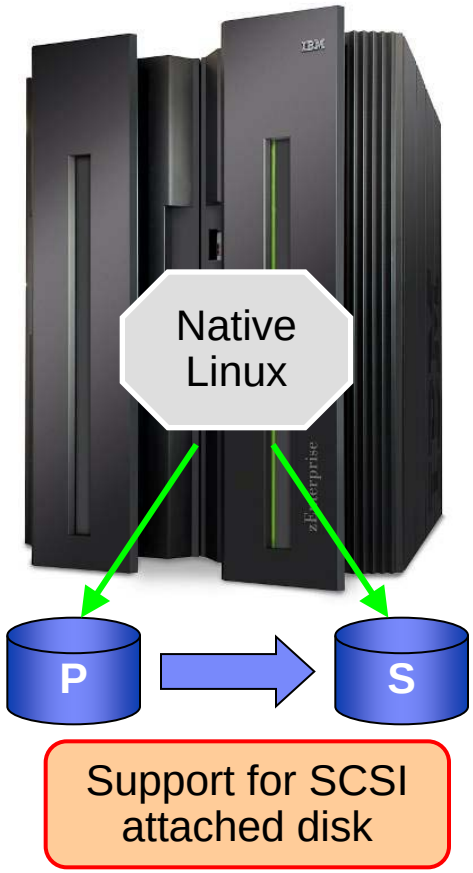

Heterogeneous data support
Flexible configuration options

xDR extended monitor function
GDPS/PPRC with xDR

- GDPS monitors native Linux and z/VM guest xDR environments for conditions that would prevent a HyperSwap from succeeding.
  - Orphaned pointers
  - Channel paths to secondary disks online (Native & Guest)
  - Proxy system disk not PPRCed (Guest)
  - Minidisk caching must not be enabled (Guest)
  - Mdraid disks must be synchronized (Native)
  - Root file system must not be read-only (Native)

- Extends GDPS HyperSwap monitoring to xDR environment.
- Raises alerts so problems can be fixed before they impact a HyperSwap.
- Autonomically corrects certain ‘broken’ conditions.
- Helps avoid having xDR systems reset after unplanned HyperSwap.
- Prevents planned HyperSwap when non-ready systems would be reset.

Heterogeneous data support
Identifies conditions where HyperSwap will fail

CPC and LPAR awareness
GDPS/PPRC (including xDR z/VM), GDPS/PPRC HM, GDPS/XRC

- Previous implementation for system resets:
  - GDPS relied on customer-maintained LPAR pointer (IPLTYPE)
  - IPLTYPE LPAR pointer and actual system LPAR could be different
  - GDPS Reset of IPLTYPE LPAR could result in:
    - Desired LPAR (and system) not getting reset
    - Unintended target LPAR (and system) getting reset

- GDPS now tracks actual CPC/LPAR location for GDPS systems
  - All GDPS Reset actions (autonomic or user initiated) reset actual location

- Avoids accidental reset of wrong system
- Avoids possible unpleasant outcome of intended system not being reset
- Eliminates dependence on manually maintained IPLTYPE pointer

- GDPS V3.7 SPE (PM17076)

Availability and Reliability

IPL Protection
GDPS/PPRC, GDPS/PPRC HM

- Prevents a system to be IPLed on the wrong set of disks
  - After a HyperSwap followed by DASD RECOVER or DELPAIR
  - Disk on both sites are “PPRC primary Suspended”
  - Checks if using the right Subchannel Set

- IPL done only from GDPS primary site
  - Helps prevent overlaying production data when mirroring restarted

- Environments
  - GDPS/PPRC HyperSwap Manager
    Verify Load and IODF volume addresses are on GDPS primary disk
    Quiesce system if IPLed from wrong set of disk

  - GDPS/PPRC IPL protection
    Enhanced protection prevents loading from wrong set of disk.
    Checks z/OS, z/VM and native zLinux systems

Protects environment after HyperSwap
Helps prevent data overlays

Scalability

- GDPS/XRC UCB Constraint Relief
- z/VM XRC time stamping support
- SNA communications

UCB Constraint Relief
GDPS/XRC, GDPS/MzGM

- Removes restriction that all devices in a GDPS/XRC configuration be defined and identified with unique device numbers.

- GDPS can optionally bypass checking for unique device numbers
  - Enables device number reuse on different SDM systems.
  - Increases the number of volumes that can be managed in a configuration.
  - Provides checking at physical device level to avoid duplicate usage.

- Additional UCB constraint relief provided with
  - Option to allow FlashCopy target devices to be defined only to the GDPS Controlling system and not to the SDM systems.
  - Option to allow specification of FlashCopy targets by LSS and CCA, so they need not be defined to any systems in the GDPS.

Support very large XRC configurations

Time stamping support for z/VM guests
GDPS/XRC, GDPS/MzGM

- Allows z/VM and guests to participate in GDPS/XRC sessions without the ‘stall’ problems and consistency exposures traditionally associated with non-timestamped writes.
- Relieves related scalability constraints.
- Allows XRC to provide consistency across z/OS data and z/VM data, including guests (such as zLinux, z/VSE and z/TPF)
- Exploits Server Time Protocol (STP) facility to ensure time consistency.
- Support provided by VM APAR VM64816

Extends XRC environment

GDPS SNA Communications Scalability
GDPS/MzGM

- MzGM - can manage multiple GDPS/PPRC sysplexes using a single GDPS/XRC
- The number of SNA nodes that can be defined to the GDPS/MzGM SNA Communication Facility has been increased from 10 to 15.
- The number of GDPS/PPRC environments (each with two control systems) that can be managed by a single GDPS/XRC environment has been increased from 4 to 7.
- GDPS V3.7 SPE (PM09738)

Scalability

Simplified Systems Management

- Incremental Zero Suspend FlashCopy
- Persistent reserve clean-up for Open LUNS
- Concurrent Copy session clean-up
- GM Monitor integrated into GDPS/GM
- Installation verification program for xDR
- GDPS/GM Query Services

Incremental Zero Suspend FlashCopy
GDPS/XRC, GDPS/MzGM

- GDPS supports XRC Zero Suspend FlashCopy
  - FlashCopy of XRC secondary while volume pairs remain duplex
  - Minimal impact to RPO achieved by the SDMs

- Incremental FlashCopy provides these benefits
  - Only copies tracks changed since the previous copy was flashed.
  - Results in faster completion of FlashCopy background copy tasks.
  - Reduces XRC session impact.
  - Allows more flexibility with FlashCopy scheduling.

- GDPS 3.7 SPE (PM19622)

Operational Flexibility

Persistent reserve clean-up for Open LUNs
GDPS/PPRC and GDPS/PPRC HM

- Applications on distributed systems may use Persistent Reserves on the Open LUNs (FB disks) as a means of serialization.
- Such reserves will be orphaned on the former primary disks if an unplanned disk switch occurs.
- Error occurs when attempting to resynchronize PPRC back to these former primary devices, using them as the target PPRC devices.

- New RPR keyword on DASD START SECONDARY and HYPERSW RESTORE optionally automates reset of persistent reserve on the former primary disks.
- GDPS V3.7 SPE (PM24877)

Heterogeneous data support
Simplifies removal of persistent reserves
Improves reliability of post-HyperSwap resynchronization

Concurrent Copy session clean-up
GDPS/PPRC and GDPS/HM

- Concurrent Copy (CC) allows point-in-time backup of data sets or volumes with a minimal update serialization window.
- Session information is kept in storage subsystem cache
- CC sessions can be orphaned on former primary disk after HyperSwap or LPAR failure.
- Orphaned sessions cause errors when attempting to resynchronize PPRC using the former primary devices as targets.

- New CCTERM keyword on DASD START SECONDARY and HYPERSW RESTORE optionally automates cleanup of CC sessions on the former primary disks.

Simplifies removal of CC sessions
Improves reliability of post-HyperSwap resynchronization

Global Mirror Monitor integrated into GDPS/GM
GDPS/GM and GDPS/MGM

- Integrates existing Global Mirror Monitor tool into GDPS/GM
  - Fully supported with GDPS
- Capabilities include:
  - Generates alerts and messages based on defined policy
  - Take automatic actions based on event or situation
    - Pause and resume Global Mirror based on policy
  - Writes detailed performance and behavioral data to SMF
    - Can integrate into existing diagnosis / performance analysis tools
  - View recent performance data for a Global Mirror session
- Both CKD and FB disk supported

Operational simplicity

Installation Verification Program
GDPS/PPRC with xDR

- Checks the setup of the xDR environment
- Helps ensure that installation and customization tasks were performed correctly.

- Operator initiated
- Recommended at installation time and for environment changes

- Works for xDR native Linux and guest Linux on System z environments.

Heterogeneous data support
Improved System Management and Availability

Query Services
GDPS/GM

- Enables retrieval of information from GDPS internal variables via user written REXX code
- Allows users to write tools, extend GDPS automation etc.
- Already available for GDPS/PPRC and GDPS/HM
- Now available for GDPS/GM
  - Queries:
    - ENVIRONMENT - information on GDPS/GM environment
    - MONITORS - information on monitors including GM Monitor
    - GDASD - information on GDPS managed disk
  - All queries available for THREESITE=NO and CASCADE
    - Only MONITORS & ENVIRONMENT for THREESITE=MGM

Allows clients to extend GDPS automation

New Tools

- GDPS/MGM Incremental Resynchronization
- Preserve Mirror Tool support for Alternate Subchannel Set
- GDPS/PPRC Configuration Checker
- GDPS Console Interface tool

Incremental Resynchronization Tool
GDPS/MGM

V3R7 - distributed “as is”, installed separately

Already existing function:
Incremental reintroduction of ‘intermediate’ disk

- When running A → MM → B → GM → C
- Planned / unplanned outage of B disk occurs
- GM is incrementally resynchronized A→C
- B disk is returned (intact)
- Incrementally reintroduce B to return to original configuration

Function added in 4Q2010:
Planned toggle between A and B disk

- When running A → MM → B → GM → C
- Planned HyperSwap A to B disk
- Incrementally resynchronize to B → MM → A → GM → C
- And same process to go back to A → MM → B → GM → C
- Nicely integrates into GDPS as USERPROC script statements

Improved availability and disaster recovery
Improved solution support

Incremental Resynchronization Tool
GDPS/MGM

V3R8 - shipped with GDPS

Supported as a standard GDPS component
- OCO
- SMP/E installed as part of GDPS code
- Standard support process
- PMRs, APARs, PTFs

Supports previously existing tool functions
- Intermediate disk reintroduction
- ABC/BAC planned toggle

New function to support incremental 'Go Home' when running A →B → C or B → A → C

- Planned/Unplanned outage of both A/B disks
- Recover on C
- When A/B are available, return to original configuration incrementally
- Requires that A/B are back 'intact'

Improved availability and disaster recovery
Improved solution support

Preserve Mirror Tool support for Alternate Subchannel Set
GDPS/PPRC and GDPS/PPRC HM

- The Preserve Mirror tool is designed to:
  - Start PPRC
  - Bring new devices to duplex state
  - Add them to the GDPS configuration
  - Monitor progress
- Keeps disk Freeze and HyperSwap ready during an initial copy.
- Now supports PPRC secondary devices defined in MSS1
- Facilitates GDPS configuration changes when exploiting alternate sub-channel set.
- “As is” tool for GDPS 3.6 and higher

Availability

Configuration Checker tool
GDPS/PPRC and GDPS/PPRC HM, GDPS/MzGM, GDPS/MGM

- Identifies devices defined in production systems that are not under GDPS management control.
- Devices may be accidentally left out from the GDPS configuration
  - Helps installations catch and correct errors.
- Helps verify entire production environment is protected
- Configurations:

  (2-site) GDPS/PPRC, GDPS/HM

  (3-site) GDPS/MzGM or GDPS/MGM running with GDPS/PPRC or GDPS/HM

Availability
Avoids data loss

New GDPS Console Interface (GCI) tool
GDPS/PPRC and GDPS/XRC

- Allows operators to issue GDPS script commands or execute GDPS scripts from the systems console.
- Provides an additional alternative when performing GDPS operations.

Simplified Operations

Agenda

- GDPS® Overview
- Enhancements (V3R8 GA and V3R7 SPE)
  - Availability
  - Scalability
  - Simplified system management
  - New tools
- **Summary and Miscellaneous**

GDPS V3.8 – Summary of Enhancements

- Availability
  - Reduced impact of false freezes
  - STP and Sysplex Timer recovery
  - HyperSwap support for FB disk
  - xDR extended monitoring
  - LPAR and CEC Awareness
  - IPL protection
- Scalability
  - UCB constraint relief for GDPS/XRC
  - z/VM XRC time stamping support
  - SNA communications
- Simplified systems management
  - Incremental Zero Suspend FlashCopy
  - Automated cleanup of persistent reserve and Concurrent Copy sessions
  - GM Monitor integrated into GDPS/GM
  - Installation verification program for xDR
  - GDPS/GM Query Services
- New & Improved GDPS tools
  - GDPS/MGM Inc. Resync tool – Phase 3
  - Preserve Mirror Tool – Alt. Subch. sets
  - GDPS/PPRC Configuration Checker
  - GDPS Console Interface tool

ibm.com/systems/z/gdps

GDPS
The Enterprise-wide Continuous Availability and Disaster Recovery Solution

Discontinued Support

- GDPS/RCMF 3.5 support ended with release 3.8 GA
  - System Automation 3.1 no longer supported (EOS 9/2010)
  - z/OS 1.9 no longer supported (EOS 9/2010)

- The following were removed in release 3.8
  - P/DAS support in GDPS/PPRC and RCMF/PPRC
  - GDPS network management functions for SNA
  - ACTIVATE command for IPL'ing systems.
    (LOAD remains as the only supported method to IPL systems)

- Note - release 3.8 is the last to support the 'policy switching' method for Coupling Facility management.
  - Clients are encouraged to move to using single policy method
  - GDPS/PPRC and GDPS/XRC

Currently in the Oven (upcoming 3.8 SPEs)

- GDPS/GM support for CopyOnce volumes
  - Such as Work volumes, Page Volumes
  - Only need volume and dataset structure existing in recovery site
  - Current data not required for recovery/restart
  - Continuous mirroring not required
- GDPS/XRC support for automated VOLSER refresh
  - Updates GDPS configuration when SDM detects label change
  - Driven by ANTI8027II messages for primary volumes
- New option for GDPS/XRC ADDPAIR
  - Add volumes that are in configuration but not in session

Additional Information

- **Web sites:**

GDPS	**www.ibm.com/systems/z/gdps**
Parallel Sysplex	**www.ibm.com/systems/z/pso**
Bus Resiliency z	**www.ibm.com/systems/z/resiliency**
Bus Resiliency	**www.ibm.com/systems/business_resiliency**
System z	**www.ibm.com/systems/z/hardware**
Storage	**www.ibm.com/systems/storage**
Redbooks®	GDPS Family: An Introduction to Concepts and Capabilities **www.redbooks.ibm.com/abstracts/sg246374.html?Open**

- **GDPS Web Site White Papers and Presentations**
  - GDPS: The Enterprise Continuous Availability / Disaster Recovery Solution
  - IBM Implementation Services for GDPS/Global Mirror
  - GDPS Business Continuity Solutions
  - Consistency Groups in a Nutshell
  - DS6000™ / DS8000™ Data Replication
  - GDPS Solutions

- **e-mail:** **gdps@us.ibm.com**

Session Suggestion

9666 Beyond Disaster Recovery: Taking Your Enterprise from High Availability to Continuous Availability

Thursday, August 11, 2011: 4:30 PM-5:30 PM Oceanic 5

Speaker: Karen Durward (IBM Corporation)

Questions?

email:
amcclure@us.ibm.com
gdps@us.ibm.com

Please remember your session evaluation.
Your Feedback is Important to Us !