



GDPS: The Enterprise Continuous Availability / Disaster Recovery Solution

David Clitherow
dave.clitherow@uk.ibm.com

David Petersen
petersen@us.ibm.com

David Raften
raften@us.ibm.com

Table of Contents

Introduction	2
IT Resilience	4
What is GDPS?	5
GDPS Suite of Offerings	6
HyperSwap Benchmark measurements	10
Need for Data Consistency	10
Need for Automation	10
GDPS Systems.....	12
Functions common to multiple GDPS solutions.....	12
FlashCopy support	12
GUI Interface	13
Support for Heterogeneous Environments.....	13
Management of IBM Z Operating Systems	13
GDPS Metro, GDPS HM, and GDPS Global – GM Open LUN Management	13
GDPS Metro Multiplatform Resiliency for IBM Z (xDR)	14
Distributed Cluster Manager.....	15
Recent Enhancements in GDPS	15
IBM Global Technology Services (GTS) Offerings	16
Technical Consulting Workshop (TCW)	16
IBM Installation Services for GDPS.....	16
Prerequisites.....	17
GDPS Metro at work in a real disaster	17
Summary	18
Additional Information	19

Introduction

How would a shutdown of your Information Technology (IT) systems affect your business? Do you put off system maintenance and upgrades to help minimize system downtime? Are your business-critical processing and data protected from a site disaster? A 2006 survey conducted by the Robert Frances Group, Figure 1, Cost of Downtime by Industry, shows the participating companies' responses to the revenue impact per hour of an outage.

Cost of Downtime by Industry	
<u>Industry/Sector</u>	<u>Revenue/Hour</u>
Financial	\$8,213,470
Telecommunications	\$4,611,604
Information Technology	\$3,316,058
Insurance	\$2,582,382
Pharmaceuticals	\$2,058,710
Energy	\$1,468,798
Transportation	\$1,463,128
Banking	\$1,145,129
Chemicals	\$1,071,404
Consumer Products	\$989,795

Source: Robert Frances Group 2006, "Picking up the value of PKI: Leveraging z/OS for Improving Manageability, Reliability, and Total Cost of Ownership of PKI and Digital Certificates."

Figure 1: Cost of Downtime by Industry

All enterprises have become much more dependent on Information Technology (IT) since the survey results were compiled in 2006. In fact, ITIC survey data indicates that the cost of hourly downtime has increased by 25% to 30% from 2008 to 2016.¹ Not only are direct customers affected by an outage, but they are also likely to let all their friends know through social media, affecting corporate reputation.

It has been observed that many companies have business continuance plans developed on the premise that back office and manual processes will keep the business running until computer systems are available. Characteristics of these recovery models may allow critical applications to recover within 24 to 48 hours, with data loss potentially exceeding 24 hours, and full business recovery taking days or weeks. As companies transform their business to compete in the business place, business continuity strategies and availability requirements should be reevaluated to determine if they are based on today's business objectives.

¹ ITIC study: <http://itic-corp.com/blog/2016/08/cost-of-hourly-downtime-soars-81-of-enterprises-say-it-exceeds-300k-on-average/>

Lessons learned about IT survival

The events of September 11, 2001 in the United States of America have underlined how critical it is for businesses to be ready for disasters. The Federal Reserve, the Office of the Comptroller of the Currency, the Securities and Exchange Commission, and the New York State Banking Department (the agencies) met with industry participants to analyze the lessons learned from the events of September 11. The agencies released an interagency white paper (referenced in the section, Additional Information) on practices to strengthen the resilience of the US financial system.

The following is a summary of lessons the agencies learned about IT service continuity:

- *Identify clearing and settlement activities in support of critical financial markets*
- *Determine appropriate recovery and resumption objectives for clearing and settlement activities in support of critical markets.*
- *Geographical separation of facilities and resources is critical to maintaining business continuity. Any resource that cannot be replaced from external sources within the Recovery Time Objective² (RTO) should be available within the enterprise, in multiple locations. This not only applies to buildings and hardware resources, but also to employees and data, since planning employee and data survival is very critical. Allowing staff to work out of a home office should not be overlooked as one way of being D/R ready.*
- *Depending on the RTO and Recovery Point Objective³ (RPO) – RTO and/or RPO are typically expressed in hours or minutes – it may be necessary for some enterprises to implement an in-house D/R solution. If this is the case, the facilities required to achieve geographical separation may need to be owned by the enterprise.*
- *The installed server capacity at the second data center can be used to meet normal day-to-day data processing needs and fallback capacity can be provided either by prioritizing workloads (production, test, development, data mining) or by implementing capacity upgrades based on changing a license agreement, rather than by installing additional capacity. Disk resources need to be duplicated for disk data that is mirrored.*
- *Recovery procedures must be well-documented, tested, maintained and available after a disaster. Data backup and/or data mirroring must run like clockwork all the time.*
- *It is highly recommended that the D/R solution be based on as much automation as possible. In case of a disaster, key skills may not be available to restore I/T services.*
- *An enterprise's critical service providers, suppliers and vendors may be affected by the same disaster, therefore, enter into a discussion with them about their D/R readiness.*

²Recovery Time Objective: a metric for how long it takes to recover the application and resume operations after an outage - planned or unplanned.

³Recovery Point Objective: a metric for how much data is lost, or the actual recovery point to which all data is current and consistent.

IT Resilience

IBM defines IT Resilience as the ability to rapidly adapt and respond to any internal or external opportunity, demand, disruption, or threat, and continue business operations without significant impact. This includes continuous / near-continuous application availability, keeping your applications up and running throughout the far more common planned and unplanned outages, and Disaster Recovery which concentrates on recovering from an unplanned event.

When investigating an IT Resilience solution, consider:

- *Support for data that spans more than one platform. Does the solution support data from more systems than just z/OS® such as z/VM®, Linux®, UNIX®, and Microsoft® Windows®? Does it support both FBA and ECKD™ data? Does it provide data consistency across all supported platforms, or only within the data from each platform? Does the solution enable a single D/R plan instead of multiple, independent, silo-based products?*
- *Support for multiple disk subsystem vendors. Does the solution allow you the flexibility to use multiple disk vendors at the same time, not tying you down to one vendor?*
- *Support for planned site outages. Does the proposed solution provide the ability to move the entire production environment (systems, subsystems, and data) from the production site to the recovery site? Does it provide the ability to move production systems back and forth between production and recovery sites with minimal or no manual intervention?*
- *Support for managing the remote copy environment. Does the solution provide an easy-to-use interface for monitoring and managing the remote copy environment? Will it automatically react to connectivity or other failures in the remote copy configuration?*
- *Support for data consistency. Does the solution provide data consistency across all remote copied volumes and disk subsystems? Does it provide support for protecting the consistency of the secondary volumes if it is necessary to re-synchronize the primary and secondary volumes?*
- *Support for continuous application availability. Does the solution support continuous application availability? From the failure of any component? From the failure of a complete site? Does the solution provide protection against planned events such as disk migrations or recycling software? Typically, solutions that provide continuous or near-continuous application availability impose fairly stringent limitations on the distance between the two sites.*
- *Support for hardware failures. Does the solution support recovery from a hardware failure? Is the recovery disruptive (re-IPL) or transparent (HyperSwap®, for example)?*
- *Support for monitoring the production environment. Does the solution provide monitoring of the production environment? Is the operator notified in case of a failure? Can recovery be automated?*

- *Automated and transparent failover. Does the solution provide for automatic detection and handling of failures and other disruptive events? Does it provide automated fail-over or recovery capabilities, with no or minimal human intervention?*
- *Automated failback. Does the solution automate the process to get back to the original configuration with no or minimal human intervention?*
- *Dynamic provisioning of resources. Does the solution have the ability to dynamically allocate resources and manage workloads? Will critical workloads continue to meet their service objectives, based on business priorities, in the event of a failure?*
- *Support for recovery across database managers. Does the solution provide recovery with consistency independent of the database manager? Does it provide data consistency across multiple database managers?*
- *End-to-end recovery support. Does the solution cover all aspects of recovery, from protecting the data via backups or remote copy, through to automatically bringing up the systems following a disaster?*
- *Are the applications cloned? Do your critical applications support data sharing and workload balancing, enabling them to run concurrently in more than one site? If so, does the solution support and exploit this capability?*
- *Support for recovery from regional disasters. What distances are supported by the solution? What is the impact on response times? Does the distance required for protection from regional disasters permit a continuous application availability capability?*

As these capabilities are the backbone of the GDPS® solution, GDPS is positioned to provide a total IT Resiliency solution **across your entire data center**.

What is GDPS?

GDPS is an integrated, automated application and data availability solution designed to provide the capability to manage the storage subsystem(s) and remote copy configuration across heterogeneous platforms, automate Parallel Sysplex® operational tasks, and perform failure recovery from a single point of control, thereby helping to improve application availability. GDPS is independent of the transaction manager (e.g., CICS® TS, IMS™, WebSphere®) or database manager (e.g., Db2®, IMS, and VSAM) being used, and is enabled by means of key IBM technologies and architectures:

- *Base or Parallel Sysplex®*
- *Tivoli® NetView® for z/OS*
- *System Automation for z/OS*
- *Disk control units such as IBM System Storage® DS8000® that support:*
 - *Metro Mirror architecture for GDPS Metro with single and dual leg configurations*
 - *Global Mirror and z/OS Global Mirror architecture for GDPS Global*
 - *Combinations of 3- and 4-site solutions*
- *Db2, IMS, and VSAM software replication for GDPS Continuous Availability*

GDPS Suite of Offerings

GDPS provides different solutions to meet different business requirements:

GDPS Metro HyperSwap Manager (GDPS HM) provides an entry level multisite disaster recovery solution at a cost-effective price. HyperSwap provides the ability to dynamically switch to secondary volumes without requiring applications to be quiesced. Typically done in 3–15 seconds in actual customer experience, this provides near-continuous data availability for planned actions and unplanned events. Using HyperSwap technology, GDPS HM provides disk remote copy management, data consistency and HyperSwap for data availability. This can be across a single site or multiple sites up to 100 km away, 200 km with RPQ. If used as part of a D/R solution across sites, then it is up to the customer to implement their own D/R procedures. One can migrate from a GDPS HyperSwap Manager implementation to the full-function GDPS Metro capability as business requirements demand shorter recovery time objectives.

GDPS Metro has all the function as GDPS HM and is also designed to fully automate the recovery at the remote site. This includes disk reconfiguration, managing servers, Sysplex resources, CBU, activation profiles, etc. GDPS Metro is designed to be a near-continuous availability and disaster recovery solution. GDPS Metro is application and data independent. It can also be used to provide a consistent recovery for z/OS as well as non-z/OS data running under Linux or Windows platforms. This is important for a common situation when a multi-tier application has dependencies upon multiple operating system architectures. It is not enough that z/OS data is consistent, but it needs to be consistent with non-mainframe data to allow rapid business resumption.

GDPS Metro also has the capability to manage the Multi-Target Metro Mirror configuration, extending PPRC management and HyperSwap capabilities to support two synchronous Metro Mirror relationships from a single primary volume. Each leg is tracked and managed independently. This provides additional data protection in the event of a disk subsystem failure or local disaster scenario.

GDPS Metro provides management extensions for heterogeneous platforms (xDR) to be able to fully manage either z/OS systems outside the GDPS sysplex, z/VM systems and their guests and KVM on IBM Z® environments providing full end-to-end support including disk management with freeze and planned/unplanned HyperSwap support, systems management and monitoring.

GDPS Metro also extends management for the disk mirroring used by systems outside of the Sysplex that are *not* under GDPS control. These systems can be other z/OS systems, VSE, z/VM and guests, KVM on IBM Z and guests, or Linux running in a native LPAR. These systems are known as “foreign systems.” Because GDPS manages PPRC for the disks used by these systems, these disks will be attached to the GDPS controlling systems. With this setup, GDPS can capture mirroring problems and

perform a freeze. All GDPS managed disks belonging to the GDPS systems and these foreign systems are frozen together, regardless of whether the mirroring problem is encountered on the GDPS system's disks or the foreign system's disks. This support also applies to z/VM and KVM on FB formatted disk. (GDPS Metro V4.1 requires SPE for this support)

If an unplanned HyperSwap occurs on a GDPS managed system, the foreign systems do not have the ability to HyperSwap to the secondaries. A long Extended Long Busy timeout (the maximum is 18 hours) is set for these systems so that when the GDPS managed systems swap, these systems hang. The ELB prevents these systems from continuing to use the former primary devices. GDPS automation can be used to reset these systems and re-IPL them using the swapped-to primary disks.

GDPS Metro is capable of:

- *Near continuous disk availability*
- *Near transparent D/R solution*
- *Recovery Time Objective (RTO) less than an hour*
- *Recovery Point Objective (RPO) of zero*
- *Protection against localized area disasters (distance between sites limited to 200 km fiber)*
- *Protection against multiple failures with three synchronous copies of the data*
- *End-End automation of Linux on Z environment*
- *Management of non-Z data*
- *Management of foreign system disk*

GDPS Global – XRC (GDPS XRC) is a highly scalable asynchronous remote copy solution for z/OS and Linux on IBM Z data. Based upon z/OS Global Mirror, it is a combined hardware and software asynchronous remote copy solution. Since z/OS Global Mirror uses asynchronous data replication, the secondary site can be thousands of miles from the primary site.

GDPS XRC is capable of:

- *End-end automated disaster recovery solution*
- *Support for z/OS, z/VM, and z/VM guests, including Linux on Z*
- *RTO between one and two hours*
- *RPO of seconds*
- *Protects against localized as well as regional disasters (distance between sites is unlimited)*
- *Minimal remote copy performance impact*

GDPS Global – GM (GDPS GM) is an asynchronous remote copy solution for z/OS and non-z/OS data. Based on the DS8000 Global Mirror, it is a hardware based asynchronous mirroring solution that is designed to maintain a consistent copy of data at virtually unlimited distances with minimal impact to application response time. Global Mirror is data independent. It can be used to provide a consistent recovery for z/OS, z/VM, Linux, as well as non-IBM Z data. This is important when a

multi-tier application has dependencies upon multiple operating system architectures.

GDPS GM is capable of:

- *End-end automated disaster recovery solution*
- *Disaster recovery solution for IBM Z and non-IBM Z data*
- *RTO between an hour to two hours*
- *RPO less than one minute*
- *Protects against regional disasters (distance between sites is unlimited)*
- *Minimal remote copy performance impact*

Three Site or Four Site Solutions are available for businesses requiring the benefits of both synchronous and asynchronous remote copy. Synchronous remote copy using GDPS Metro and GDPS HM provides benefits such as near-continuous availability using HyperSwap and the ability to configure for zero data loss. Asynchronous remote copy using GDPS XRC or GDPS GM provide benefits such as protection from regional disasters with little to no application impact. To provide for this requirement, GDPS supports three-site configurations to help provide maximum availability across the widest range of possible scenarios.

- *GDPS Metro or GDPS HM combined with GDPS XRC is called GDPS Metro Global – XRC, or GDPS MzGM.*
- *GDPS PPRC or GDPS HM combined with GDPS GM is called GDPS Metro Global – GM, or GDPS MGM. GDPS MGM is a solution that also provides a solution for both IBM Z and non-Z data.*

GDPS MGM and GDPS MzGM are capable of:

- *Down to zero data loss*
- *HyperSwap for disk availability*
- *Protection from regional events*
- *No impact to end user response time*
- *Disaster recovery solution for IBM Z and non-IBM Z data (with GDPS MGM)*

More information on GDPS is available in “GDPS Family – An Introduction to Concepts and Capabilities” at

<http://www.redbooks.ibm.com/redpieces/abstracts/sg246374.html?Open>

GDPS Continuous Availability (GDPS AA), previously called GDPS/Active-Active, is an asynchronous remote copy solution for select z/OS workloads. GDPS Continuous Availability is a software based asynchronous mirroring solution that is designed to maintain a consistent copy of data within IMS, within Db2, and within VSAM data at virtually unlimited distances, recovery time measured in seconds, with minimal impact to application response time. Since it is a software-based solution, GDPS Continuous Availability is disk independent. Both an Active-Standby configuration and an Active-Query configuration to offload read-only workload to the recovery site are supported.

GDPS AA works together with GDPS Metro to provide a zero data loss solution. GDPS AA working together with GDPS Metro to keep the non-database data (“flat files”) current at the recovery site.

GDPS Continuous Availability is capable of:

- *Continuous Availability and Disaster recovery solution for Db2, IMS, and VSAM data*
- *Application level granularity*
- *Site switch in seconds (once the event is detected)*
- *RPO in seconds (with the ability to issue reports that can be generated showing orphaned data)*
- *Protects against metro and regional disasters (distance between sites is unlimited)*
- *Minimal response time impact*

More information on GDPS/Active-Active is available in announcement 613-033, “IBM GDPS active/active continuous availability active query configuration,” dated October 22, 2013. Also, see “GDPS Family – An Introduction to Concepts and Capabilities” at <http://www.ibm.com/systems/z/advantages/gdps/resources.html>

Near Continuous Availability of data with HyperSwap

GDPS Metro supports HyperSwap. This function is designed to broaden the near continuous availability attributes of GDPS Metro by extending the Parallel Sysplex redundancy to disk subsystems. The HyperSwap function can help significantly reduce the time needed to switch to the secondary set of disks while keeping the z/OS systems active, together with their applications.

HyperSwap is designed to transparently switch to use secondary disk subsystems which contain mirrored data consistent with the primary data, in the event of unplanned outages of the primary disk subsystems or a failure of the site containing the primary disk subsystems (site 1).

HyperSwap can also provide the ability to transparently switch all primary disk subsystems with the secondary disk subsystems for planned reconfigurations. It can provide the ability to perform disk configuration maintenance and planned site maintenance without requiring any applications to be quiesced. The important ability to re-synchronize incremental disk data changes, in both directions, between primary and secondary disks is provided as part of this function.

HyperSwap provides support for:

- *Production systems to remain active during a disk subsystem failure. Disk subsystem failures will no longer constitute a single point of failure for an entire sysplex.*
- *Production systems to remain active during a failure of the site containing the primary disk subsystems (site 1) for many disaster scenarios, if applications are cloned and exploiting data sharing across the two sites.*
- *Support for very large disk configurations.*

HyperSwap Benchmark measurements

HyperSwap is a key part of the availability strategy and disaster recovery plan for many customers. As such, it is often tested by them. In a planned test, a customer was able to switch 20,670 device pairs across 106 LCUs in 15 seconds of user impact time with a planned HyperSwap, other customers had swapped in down to 3 seconds with smaller configurations. z/OS Processing then continued. An unplanned HyperSwap should be even faster.

Need for Data Consistency

Data consistency across all secondary volumes spread across any number of storage subsystems is essential in providing data integrity and the ability to do a normal database restart in the event of a disaster. The main focus of GDPS control software is whatever happens in site 1, to allow the secondary copy of the data in site 2 to be data consistent to allow rapid restart. Data consistent means that, from an application's perspective, the secondary disks contain all updates until a specific point in time, and no updates beyond that specific point in time.

Data recovery involves restoring image copies and logs to disk and executing forward recovery utilities to apply updates to the image copies. This process can take many hours or days. With GDPS applications only need to be restarted. An installation can be up and running quickly, even when the primary site (site 1) has been rendered totally unusable.

GDPS Metro uses a combination of storage subsystem and Sysplex technology triggers to ensure a data consistent secondary site (site 2) copy of the data. This is done using the PPRC freeze function. The freeze function, initiated by automated procedures, is designed to freeze the image of the secondary data at the very first sign of a disaster, even before any database managers are made aware of I/O errors. This can prevent the logical contamination of the secondary copy of data that would occur if any storage subsystem mirroring were to continue after a failure that prevents some, but not all secondary volumes from being updated.

Data consistency in the GDPS Global environments are provided by the underlying remote copy technologies themselves, although GDPS helps control and manage the environments and automate the recovery actions.

Need for Automation

Implementing remote copy, tape remote copy, FlashCopy®, and so on, are necessary prerequisites to be able to recover from a disaster given stringent Recovery Time and Recovery Point objectives. However, they are only enabling technologies. In order to achieve the stringent objectives, it is necessary to tie those technologies together with automation.

In an average computer room immediately following a basic system failure, all the phones are ringing, every manager within reach moves in to find out when everything will be recovered, the operators are frantically scrambling for procedures that are more than likely out of date, and the Systems Programmers are all vying with the operators for control of the consoles. In short - chaos!

Imagine instead a scenario where the only manual intervention is to confirm that one should proceed. From that point on, the system will recover itself using well tested procedures. It responds to messages at system speed. You don't need to worry about out of date procedures being used. The operators can concentrate on handing calls and queries from the assembled managers. And the Systems Programmers can concentrate on pinpointing the cause of the outage, rather than trying to get everything up and running again.

And all of this is just for a system outage. In a disaster recovery situation one also needs to invoke Capacity Back-Up (CBU), remove failed systems from the sysplex, switch disk to use secondaries, reverse the remote copy, clean up CF structures and switch policies, modify activation profiles to come up on the second site using the correct IPL volume, switch network resources, IPL failed systems, quiesce discretionary LPARs, and so on.

Training staff takes time. People come and go. You cannot be sure that the staff that took part in the last disaster recovery test will be on hand to drive recovery from this real disaster. In fact, depending on the nature of the disaster, your skilled staff may not even be available to drive the recovery.

Even for the day-to-day activities in setting up and modifying a system for planned activities, managing a remote copy environment is complicated. One needs to define the remote copy pairs, establish paths, establish pairs, constantly monitor if remote copy is ever broken, then re-synchronize and re-establish remote copy. This needs to be done for each of the thousands of volumes.

The use of automation removes these concerns as potential pitfalls to your successful recovery.

But GDPS day-to-day automation goes beyond just management of the remote copy environment. It includes various standard actions which can be initiated against a single system or a group of systems to:

- 1) *Quiesce a system's workload and remove the system from the Parallel Sysplex cluster (i.e., stop the system prior to a hardware change window);*
- 2) *IPL a system (i.e., start the system after a hardware change window); and*
- 3) *Quiesce a system's workload, remove the system from the Parallel Sysplex cluster, and re-IPL the system (e.g., recycle a system to pick up software maintenance).*

- 4) *Invoke HyperSwap to perform disk maintenance and planned site maintenance without requiring applications to be quiesced.*
- 5) *Manage CF structures*
- 6) *Manage z/OS resources such as Couple Data Sets or JES2 Checkpoint data sets*
- 7) *Invoke Capacity BackUp (CBU) or On/Off Capacity Upgrade on Demand (OOCUOD) policy*
- 8) *Customizable scripting capability for user defined actions*

All GDPS functions can be performed from a single point of control using a GUI interface. This can help simplify system resource management.

GDPS Systems

The controlling system coordinates GDPS processing. By convention all GDPS functions are initiated and coordinated by the controlling system.

All GDPS systems run GDPS automation based upon Tivoli NetView for z/OS and Tivoli System Automation for z/OS. Each system can monitor the sysplex cluster, Coupling Facilities, and storage subsystems and maintain GDPS status. GDPS automation can coexist with an enterprise's existing automation product.

Functions common to multiple GDPS solutions

The following functions are supported by multiple GDPS solutions:

FlashCopy support

FlashCopy, available on the IBM System Storage DS Family is designed to provide an “instant” point-in-time copy of the data for application usage such as backup and recovery operations. FlashCopy can enable you to copy or dump data while applications are updating the data. FlashCopy before resynchronization is automatically invoked (based upon policy) whenever a resynchronization request is received. This function provides a consistent data image to fall back to, in the rare event that a disaster should occur while resynchronization is taking place. FlashCopy can also be user-initiated at any time. Customers can then use the tertiary copy of data to conduct D/R testing while maintaining D/R readiness, perform either test/development work, shorten batch windows, etc. GDPS automation is designed to help initiate and manage different flavors of FlashCopy data, including COPY, NOCOPY, and Incremental.

GUI Interface

GDPS offers a graphical user interface (GUI) to simplify management of the GDPS and remote copy environment. This allows many capabilities such as alert monitoring, remote copy views and commands with drill-down capability to device level, standard Actions such as LOAD, STOP, or RESET an LPAR, view and execute GDPS policies, sysplex management activities such as couple data set management, and many other capabilities. The GUI is designed to have the same look and feel as used by the IBM storage products.

Support for Heterogeneous Environments

Management of IBM Z Operating Systems

In addition to managing images within the base or Parallel Sysplex cluster, GDPS can now also manage a customer's other IBM Z production operating systems and data – these include z/OS, Linux for IBM Z, z/VM, KVM, and VSE/ESA™. For example, if the volumes associated with the Linux images are mirrored using PPRC, GDPS can restart these images as part of a planned or unplanned site reconfiguration. The Linux for IBM Z images can either run as a logical partition (LPAR) or as a guest under z/VM. The operating systems must run on servers that are connected to the same Hardware Management Console (HMC) Local Area Network (LAN) as the GDPS control system

GDPS Metro, GDPS HM, and GDPS Global – GM Open LUN Management

GDPS Metro, GDPS HM, and GDPS GM technology have been extended to manage a heterogeneous environment of IBM Z and distributed systems Logical Unit Numbers (LUNs), also known as Fixed Block disk, or FB disk. If installations share their disk subsystems between the IBM Z and distributed systems platforms, GDPS Metro, GDPS HM, and GDPS GM can manage the Metro Mirror and Global Mirror remote copy configurations, as well as FlashCopy for distributed systems storage. GDPS Metro and GDPS Global – GM are also designed to provide:

- *A single point of control and management for both FB and ECKD disk replication*
- *Freeze capability to protect the consistency group for synchronous replication*
- *Data consistency across the Z and distributed applications for applications that span multiple tiers*
- *HyperSwap capability for FB devices with native Linux on Z*

Open LUN Management requires a small CKD access device for each FB cluster being managed. This is so the GDPS control system can receive status and send remote copy commands to the disk devices. On GDPS V4.1, GDPS Metro can use the DS8000 zFBA support and not require a separate access device.

Using Open LUN Management support allows GDPS to be a single point of control to manage business resiliency across multiple tiers in the infrastructure, improving cross-platform system management and business processes.

GDPS Metro Multiplatform Resiliency for IBM Z (xDR)

GDPS Metro provides a function called “GDPS Metro Multiplatform Resiliency for IBM Z,” also referred to as cross-platform disaster recovery, or xDR. This function is especially valuable for customers who share data and storage subsystems between z/OS and Linux z/VM guests on IBM Z or SUSE Linux running native on IBM Z LPARs. For example, an application server running on Linux on IBM Z and a database server running on z/OS.

With a multi-tiered architecture, there is a need to provide a coordinated near Continuous Availability/Disaster Recovery solution for both z/OS and Linux on IBM Z. GDPS Metro can provide this capability when Linux is running as a z/VM guest or native. Using the HyperSwap function so that the virtual device associated with one real disk can be swapped transparently to another disk, HyperSwap can be used to switch to secondary disk storage subsystems mirrored Metro Mirror. If there is a hard failure of a storage device, GDPS coordinates the HyperSwap with z/OS for continuous availability spanning the multi-tiered application. HyperSwap is supported for ECKD and xDR managed FB disk.

For site failures, GDPS invokes the Freeze function for data consistency and rapid application restart, without the need for data recovery. HyperSwap can also be helpful in data migration scenarios to allow applications to migrate to new disk volumes without requiring them to be quiesced.

When using ECKD formatted disk, GDPS Metro can provide the reconfiguration capabilities for the Linux on IBM Z servers and data in the same manner as for z/OS systems and data. To support planned and unplanned outages these functions have been extended to KVM on IBM Z with GDPS V4.1. GDPS provides the recovery actions such as the following examples:

- *Re-IPL in place of failing operating system images*
- *z/VM Live Guest Relocation management*
- *Manage z/VM LPARs and z/VM guests, including Linux on Z*
- *Heartbeat checking of Linux guests*
- *Disk error detection*
- *Data consistency with freeze functions across z/OS and Linux*
- *Site takeover/failover of a complete production site*
- *Single point of control to manage disk mirroring configurations*
- *Coordinated recovery for planned and unplanned events*

Additional support is available for Linux running as a guest under z/VM. This includes:

- *Re-IPL in place of failing operating system images*
- *Ordered Linux node or cluster start-up and shut-down*

- *Coordinated planned and unplanned HyperSwap of disk subsystems, transparent to the operating system images and applications using the disks*
- *Transparent disk maintenance and failure recovery with HyperSwap across z/OS and Linux applications*

GDPS XRC support is provided by z/VM. By inserting timestamps on I/Os for z/VM and its guests, virtually any z/VM guest operating system data can take part of a GDPS XRC configuration, even if the data spans multiple Logical Subsystems (LSSs).

Distributed Cluster Manager

Distributed Cluster Management (DCM) enables coordination of planned and unplanned actions between IBM Z and distributed servers clustered using clustering software. A GDPS DCM agent running in each distributed cluster will provide advisory and coordination functions between GDPS and one or more distributed clusters. The advisory functions will provide the capability of continuous heartbeat and status gathering to alert the support staff about any events that may prevent recovery at the time of an outage. The coordination functions will allow workflow integration for takeover and recovery testing, cross-platform monitoring to maintain recovery capability and cross-platform recovery management to provide an automated enterprise-level rapid recovery in the case of an outage. DCM support is provided for Veritas® Infoscale Availability with GDPS Metro and GDPS Global. The distributed cluster manager can monitor the status of applications on multiple operating systems such as IBM AIX®, Sun-Solaris, HP-UX, Linux, Microsoft Windows, and VMware based systems. It can then automatically move them to another server in the event of a fault. This is designed to be coordinated with GDPS from a central point of control.

Recent Enhancements in GDPS

A summary of the enhancements announced with GDPS is listed in the web site:

www.ibm.com/systems/z/advantages/gdps/whatsnew.html

IBM Global Technology Services (GTS) Offerings

The following GDPS services and offerings are available from IBM Global Services.

Technical Consulting Workshop (TCW)

TCW is a two day workshop where IGS specialists work with your representatives to understand your business objectives, service requirements, technological directions, business applications, recovery processes, cross-site and I/O requirements. High-level education on GDPS is provided, along with the service and implementation process. Various remote and local data protection options are evaluated.

IGS specialists present a number of planned and unplanned GDPS reconfiguration scenarios, with recommendations on how GDPS can assist you in achieving your objectives. At the conclusion of the workshop, the following items are developed: acceptance criteria for both the test and production phases, a high level task list, a services list, and project summary.

IBM Installation Services for GDPS

- *Assists in planning, configuring, and automation code customization*
- *Provides onsite assistance*
- *Provides an automated, cross-platform disaster recovery solution (GDPS Metro, GDPS HM, GDPS Global, GDPS Continuous Availability, as well as the 3-site and 4-site solutions combining GDPS Metro and GDPS Global)*
- *Includes onsite delivery, configuration, implementation and testing*
- *Provides training for your support staff*
- *Provides centralized management of your data replication and recovery environment leveraging automated technologies to help provide an end-to-end disaster recovery solution*

The services also include project management and support throughout the engagement, and assistance to help you implement any prerequisite software.

Prerequisites

For IGS to provide these services, you must have certain prerequisite hardware and software. These are listed in the GDPS Web site: <https://www.ibm.com/it-infrastructure/z/technologies/gdps>

GDPS Metro at work in a real disaster

How well does GDPS perform in a real disaster such as a fire? GDPS Metro was put to the test in an actual disaster incident, and the results convinced VPC, the first customer to implement GDPS Metro, that it really works.

VPC AB, a security depository and clearing (CSD) organization, has a GDPS Metro configuration in production – a 3-way Parallel Sysplex cluster. At the time 100 volumes were being mirrored using Metro Mirror between two sites separated by less than 10 km.

In the middle of the night, the operator on call received a GDPS TAKEOVER alert. Since an attempt to call the data center was unsuccessful, two operators traveled to the data center, and verified that there had been a power loss in the primary site (site 1) due to a cable fire in an infrastructure support area. The fire had been put out by the security personnel stationed in the building.

As soon as the real disaster was verified, a decision was made to execute the site TAKEOVER. A short time later, production applications were up and running in site 2.

Summary

GDPS is designed to provide not only resource sharing, workload balancing, and near continuous availability benefits of a Parallel Sysplex environment, but it can enhance the capability of an enterprise to recover from disasters and other failures and to manage planned exception conditions. GDPS can allow a business to achieve its own continuous availability and disaster recovery goals. Through proper planning and exploitation of IBM's GDPS technology, enterprises can help protect their critical business applications from an unplanned or planned outage event.

GDPS is application independent and, therefore, can cover the customer's comprehensive application environment. Note that specific software subsystem solutions such as IMS Remote Site Recovery are very effective, but applicable to IMS applications only. When comparing GDPS with other near continuous availability and D/R solutions, the following factors must be considered:

- *Do you want to improve your application availability?*
- *Does the solution handle both planned and unplanned outages? (Refer to Figure 1: Cost of Outage/Hour for the potential impact of outages.)*
- *Which solution meets the RTO of your business? Note that you may have different RTOs for the different applications in your organization. RTO for your critical applications should be as small as possible.*
- *Which solution meets the RPO of your business? Note that you may have different RPOs for the different applications in your organization. Data loss for your critical applications should be none or minimal when there is an outage or disaster.*
- *Do you want to minimize the cost of taking repetitive volume dumps, transporting the cartridges to a safe place, keeping track of which cartridges should be moved to which location and at what time?*
- *What is the cost of disaster recovery drills?*

The ease of planned system, disk, Remote Copy and site reconfigurations offered by GDPS may allow your business to reduce on-site manpower and skill required for these functions. GDPS can enable a business to control its own near continuous availability and disaster recovery goals.

Additional Information

GDPS home page:

<https://www.ibm.com/it-infrastructure/z/technologies/gdps>

IBM Implementation Services for GDPS:

<https://www.ibm.com/us-en/marketplace/gdps-implementation-services>

For an overview of IBM Z Parallel Sysplex clustering technology and how it can enable your business achieve near continuous availability, refer to

<https://www.ibm.com/it-infrastructure/z/technologies/parallel-sysplex>

For an overview of Server Time Protocol (STP) and how it can help in a GDPS environment, refer to

<https://www.ibm.com/it-infrastructure/z/technologies/parallel-sysplex-stp>

For Interagency White Paper on Sound Practices to strengthen the resilience of the US. Financial System, refer to: [sec.gov/news/studies/34-47638.htm](https://www.sec.gov/news/studies/34-47638.htm)

For Summary of "Lessons Learned" from Events of September 11 and Implications for Business Continuity prepared by the Securities and Exchange Commission, refer to: [sec.gov/divisions/marketreg/lessonslearned.htm](https://www.sec.gov/divisions/marketreg/lessonslearned.htm)

For complete results of the survey conducted in 2001 by Contingency Planning Research, refer to: [Contingencyplanningresearch.com/2001%20Survey.pdf](https://contingencyplanningresearch.com/2001%20Survey.pdf)

GDPS Family - An Introduction to Concepts and Capabilities, SG24-6374, at www.redbooks.ibm.com/abstracts/sg246374.html

For additional information on GDPS, contact your IBM representative or email gdps@us.ibm.com.



Copyright IBM Corporation 2018
IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.

04/2018

IBM, IBM eServer, IBM (logo), AIBM Z, IX, CICS, Db2, DS8000, ECKD, Enterprise Storage Server, FlashCopy, GDPS, HyperSwap, IMS, NetView, Parallel Sysplex, Sysplex Timer, System Storage, Tivoli, Veritas, VSE/ESA, WebSphere, z/OS, z/VM and z/VSE are trademarks or registered trademarks of the International Business Machines Corporation.

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 therefrom.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds 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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

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

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

ZSW01920-USEN-18