Executive Summary
March 2005

GDPS: The e-business Availability Solution

Noshir Dhondy
dhondy@us.ibm.com

David Petersen
Petersen@us.ibm.com

David Raften
raften@us.ibm.com

Table of Contents

Introduction . . . . . Page 2
Lessons learned about IT survival . . . . . Page 6
What is a GDPS? . . . . . Page 7
Need for Data Consistency . . . . . Page 7
GDPS Systems . . . . . Page 8
GDPS/PPRC . . . . . Page 9
Topology . . . . . Page 9
Near-Continuous Availability with HyperSwap . . . . . Page 11
Planned Reconfiguration support . . . . . Page 11
Benchmark measurements using HyperSwap for Planned Reconfiguration . . . . . Page 12
Unplanned Reconfiguration support . . . . . Page 13
Benchmark measurements using HyperSwap for Unplanned Reconfiguration . . . . . Page 13
GDPS/PPRC HyperSwap Manager . . . . . Page 14
Near continuous availability of data within a single site . . . . . Page 15
Near continuous availability of data/Disaster Recovery solution at metropolitan distances Page 15
GDPS/PPRC HyperSwap Manager Prerequisites . . . . . Page 16
GDPS/XRC . . . . . Page 17
GDPS/XRC Topology . . . . . Page 17
Planned Reconfiguration support . . . . . Page 18
Coupled System Data Mover support . . . . . Page 19
Functional Highlights (GDPS/PPRC and GDPS/XRC) . . . . . Page 20
Peer-to-Peer Virtual Tape Server (PtP VTS) support . . . . . Page 20
FlashCopy support . . . . . Page 20
Recent Performance Enhancements . . . . . Page 21
GDPS Support for Heterogenous Environments . . . . . Page 22
Management of zSeries Operating Systems . . . . . Page 22
GDPS/PPRC Management for Open Systems LUNs (Logical Unit Number): . . . . . Page 22
GDPS/PPRC Multi-Platform Resiliency for zSeries . . . . . Page 22
IBM Global Services (IGS) Offerings . . . . . Page 24
Technical Consulting Workshop (TCW) . . . . . Page 24
Remote Copy Management Facility™ (RCMF) . . . . . Page 24
GDPS/PPRC HyperSwap Manager . . . . . Page 24
GDPS . . . . . Page 24
Prerequisites . . . . . Page 25
GDPS/PPRC at work in a real disaster . . . . . Page 25
Summary . . . . . Page 27
Additional Information . . . . . Page 28

Introduction

How would a shutdown of your z/OS® system affect your business? Do you put off system maintenance and upgrades to help minimize system downtime? Is your business-critical processing and data protected from a site disaster? In 2001, a survey[1] was conducted as a joint effort between Contingency Planning Research, a division of Eagle Rock Alliance, and Contingency Planning & Management magazine. Figure 1, Cost of Outage/hour, shows the participating companies' response to the hourly impact of an outage.

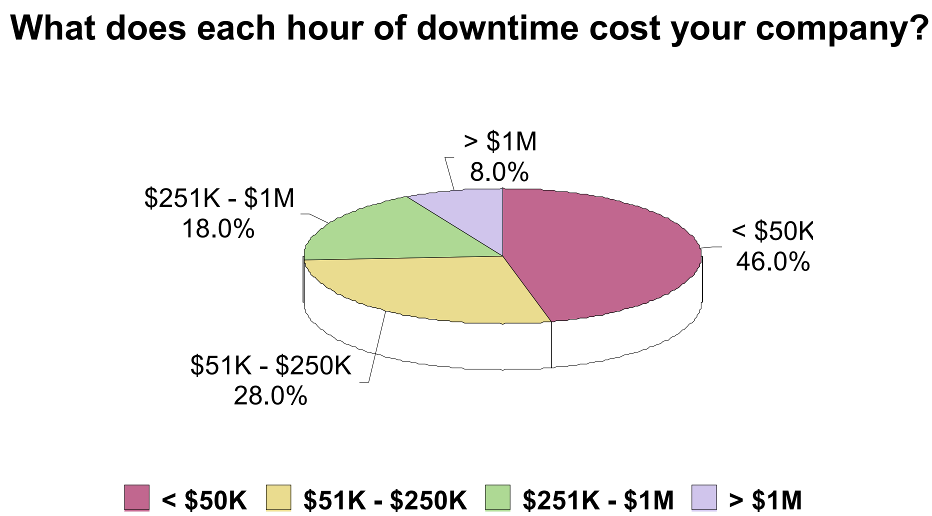

Figure 1: Cost of Outage/hour

Source: 2001 Survey by Eagle Rock Alliance (www.eaglerockalliance.com)

All enterprises have become much more dependent on Information Technology (IT) since the survey results were compiled. As a result, the average hourly cost of an outage could be significantly higher in 2005 than shown in Figure 1.

[1] The survey was intended to provide a limited updating on certain elements of the original "1996 Cost of Downtime Research Study" conducted by Contingency Planning Research (CPR). The survey questions were aimed at benchmarking the changes that have occurred in the business continuity landscape during the five years since the full research study was last conducted.

***Survey Methodology** Contingency Planning Research and Contingency Planning & Management magazine presented business continuity professionals an opportunity to participate in this survey at www.ContingencyPlanning.com. Registered website visitors were directed from the home page to an online questionnaire for completion and submission. This survey was available online from March to May 2001, at which time 163 tabulations were received. Further information and demographics are available upon request from Contingency Planning Research at cpr@eaglerockalliance.com or 800-CPR-5511.

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 e-marketplace, business continuity strategies and availability requirements should be reevaluated to determine if they are based on today's business objectives.

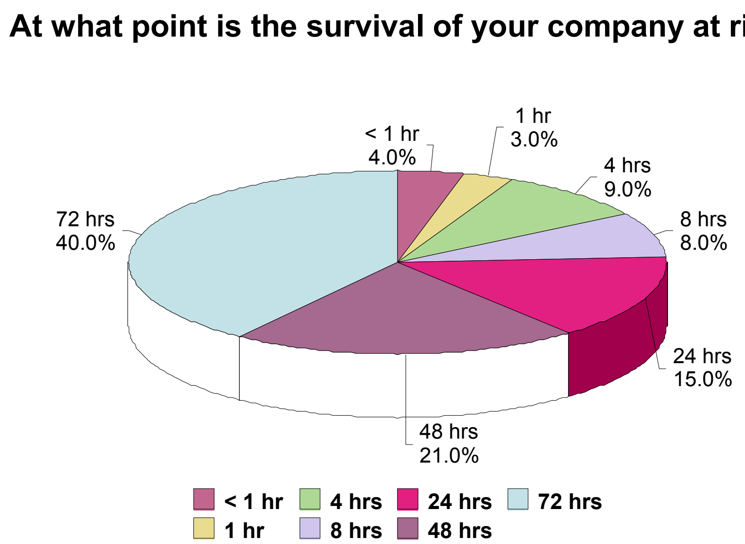

Figure 2: Survival of company at risk

Source: 2001 Survey by Eagle Rock Alliance (www.eaglerockalliance.com)

In the survey referenced above, respondents were also asked at what point is the survival of their company at risk. The results are summarized in Figure 2, Survival of company at risk. Just as explained above for the cost of outages, it may also be expected that in the e-marketplace of 2005, a larger percentage of businesses may not survive an outage of even a limited duration of 4 to 8 hours.

In e-business, two important objectives for survival are systems designed to provide continuous availability and near transparent disaster recovery (DR). Systems that are designed to deliver continuous availability combine the characteristics of high availability and near continuous operations to deliver high levels of service – targeted at 24x7. High availability is an attribute of a system that provides service at agreed upon levels and can mask **unplanned outages** from end users. Near continuous operations, on the other hand, is the attribute of a system designed to continuously operate and mask **planned outages** from end users. To attain high levels of continuous availability and near-transparent DR, the solution should be based on geographical clusters and data mirroring. These technologies are the backbone of the **GDPS®** solution. The GDPS solution, based on Peer-to-Peer Remote Copy (PPRC, recently renamed to IBM TotalStorage® Metro Mirror), is referred to as GDPS/PPRC, and the GDPS solution based on Extended Remote Copy (XRC, recently renamed to IBM TotalStorage z/OS Global Mirror), is referred to as GDPS/XRC.

GDPS/PPRC is designed with the attributes of a continuous availability and disaster recovery solution. Metro Mirror (PPRC) is a hardware solution that synchronously mirrors data residing on a set of disk volumes, called the primary volumes, to secondary disk volumes in a second system. Only when the primary storage subsystem receives "write complete" from the secondary storage subsystem is the application I/O signaled completed.

In GDPS/PPRC, since IBM Parallel Sysplex® clustering technology is designed to enable resource sharing and dynamic workload balancing, enterprises can now dynamically manage workloads across multiple sites which can enable them to achieve high levels of availability. With the introduction of GDPS/PPRC HyperSwap™ Manager, described later in this paper, Parallel Sysplex availability can now be extended to disk subsystems, even if multiple sites are not available and the Parallel Sysplex is configured in one site. GDPS/PPRC complements a multisite Parallel Sysplex implementation by providing a single, automated solution to dynamically manage storage subsystem mirroring (disk and tape), processors, and network resources designed to help a business to attain "continuous availability" and "near transparent business continuity (disaster recovery)" with no or minimal data loss. GDPS/PPRC is designed to minimize and potentially eliminate the impact of any failure including disasters, or a planned outage.

It is designed to provide the ability to perform a controlled site switch for both planned and unplanned site outages, with no or minimal data loss, maintaining full data integrity across multiple volumes and storage subsystems and the ability to perform a normal Data Base Management System (DBMS) restart – not DBMS recovery – in the second site. GDPS/PPRC is application independent and therefore can cover the customer's complete application environment.

On the other hand, GDPS/XRC has the attributes of a Disaster Recovery solution. z/OS Global Mirror (XRC) is a combined hardware and software asynchronous remote copy solution. The application I/O is signaled completed when the data update to the primary storage is completed. Subsequently, a DFSMSdfp™ component called System Data Mover

(SDM), typically running in the recovery site (site 2), asynchronously offloads data from the primary storage subsystem's cache and updates the secondary disk volumes. Refer to the section, **Need for Data Consistency**, for details on how the SDM provides data update sequence consistency for all volumes participating in the XRC session.

In GDPS/XRC, the production system(s) located in site 1 can be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster[2]. GDPS/XRC provides a single, automated solution, designed to dynamically manage storage subsystem mirroring (disk and tape) to allow a business to attain "near transparent" disaster recovery with minimal data loss. GDPS/XRC is designed to provide the ability to perform a controlled site switch for an unplanned site outage, maintaining data integrity across multiple volumes and storage subsystems and the ability to perform a normal Data Base Management System (DBMS) restart – not DBMS recovery – in the recovery site. GDPS/XRC is application independent and therefore can cover the customer's complete application environment.

For the remainder of this paper, the terms GDPS/PPRC and GDPS/XRC will be used when there is a need to differentiate the two solutions. The term GDPS will be used to discuss items common to both solutions. The differences of the GDPS/PPRC and GDPS/XRC implementation should be carefully reviewed before deciding which implementation best fits your business objectives.

A number of GDPS customers have reported that when running D/R tests, they have experienced significant reductions to the recovery time window and have experienced no data loss or minimal data loss after the recovery. GDPS/PPRC works even in a real disaster. A customer's data center experienced a fire, which generated a GDPS TAKEOVER alert. A short time after the decision was made to execute the site TAKEOVER script, the production applications were up and running in site 2. Refer to the section, **GDPS/PPRC at Work in a Real Disaster** for a more detailed description.

---

[2] The Parallel Sysplex cluster does not span across multiple sites - physical or logical.

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) have met with industry participants to analyze the lessons learned from the events of September 11. The agencies have released an interagency white paper (referenced in the section, **Additional Information**) on sound practices to strengthen the resilience of the US financial system.

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

- *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 Objective*[3] *(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 Objective*[4] *(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.*

GDPS, based on geographical separation and automation, is clearly positioned to provide a total business continuity solution for the entire IBM @server® zSeries® platform.

---

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

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

What is a GDPS?

GDPS is an integrated, automated application and data availability solution designed to provide the capability to manage the remote copy configuration and storage subsystem(s), 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*
- *IBM TotalStorage DS6000 and DS8000 series and Enterprise Storage Server® (ESS)*
- *Peer-to-Peer Virtual Tape Server (PtP VTS)*
- *Optical Dense or Coarse Wavelength Division Multiplexer (DWDM or CWDM)*
- *Metro Mirror architecture for GDPS/PPRC*
- *z/OS Global Mirror architecture for GDPS/XRC*
- *Virtual Tape Server Remote Copy architecture*

GDPS supports both the synchronous (Metro Mirror) as well as the asynchronous (z/OS Global Mirror) forms of remote copy. GDPS also supports Peer-to-Peer Virtual Tape Server (PtP VTS) form of remote copying tape data. The GDPS solution is a nonproprietary solution, working with IBM as well as Other Equipment Manufacturer (OEM) disk vendors, as long as the vendor meets the specific functions of the Metro Mirror and z/OS Global Mirror architectures required to support GDPS functions as documented in the section, **Prerequisites**.

Need for Data Consistency

Data consistency across all primary and 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 automation is whatever happens in site 1, to allow the secondary copy of the data in site 2 to be data consistent (the primary copy of data in site 1 will be data consistent for any site 2 failure). 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.

Time consistent data in the secondary site allow applications to restart in the secondary location without having to go through a lengthy and time-consuming data recovery process. 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. Since 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/PPRC uses a combination of storage subsystem and Parallel Sysplex technology triggers to capture, at the first indication of a potential disaster, a data consistent secondary site (site 2) copy of the data, 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 a GDPS/XRC environment is provided by the Consistency Group (CG) processing performed by the System Data Mover (SDM). The CG contains records that have their order of update preserved across multiple Logical Control Units within a storage subsystem and across multiple storage subsystems.

Providing data consistency enables the secondary copy of data to perform normal restarts (instead of performing database manager recovery actions). This is the essential design element of GDPS in helping to minimize the time to recover the critical workload, in the event of a disaster in site 1.

GDPS Systems

GDPS consists of production systems and controlling systems. The production systems execute the mission critical workload. There must be sufficient processing resource capacity (typically in site 2), such as processor capacity, main storage, and channel paths available that can quickly be brought on-line to restart a system's or site's critical workload (typically by terminating one or more systems executing expendable [non-critical] work and acquiring its processing resource). The Capacity BackUp (CBU) feature, available on the IBM @server zSeries server could provide additional processing power, which can help you to achieve cost savings. The CBU feature has the ability to increment capacity temporarily, when capacity is lost elsewhere in the enterprise. CBU adds Central Processors (CPs) to the available pool of processors and is activated only in an emergency. GDPS-CBU management automates the process of dynamically adding reserved Central Processors (CPs), thereby helping to minimize manual customer intervention and the potential for errors. The outage time for critical workloads can potentially be reduced from hours to minutes. Similarly, GDPS-CBU management can also automate the process of dynamically returning the reserved CPs when the temporary period has expired.

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.

GDPS/PPRC

GDPS/PPRC is designed to manage and protect IT services by handling planned and unplanned exception conditions, and maintain data integrity across multiple volumes and storage subsystems. By managing both planned and unplanned exception conditions, GDPS/PPRC can help to maximize application availability and provide business continuity.

GDPS/PPRC is capable of the following attributes:

- *Near continuous Availability solution*
- *Near transparent D/R solution*
- *Recovery Time Objective (RTO) less than an hour*
- *Recovery Point Objective (RPO) of zero (optional)*
- *Protects against localized area disasters (distance between sites limited to 100 km fiber)*

Topology

The physical topology of a GDPS/PPRC, referring to Figure 3: GDPS/PPRC, consists of a base or Parallel Sysplex cluster spread across two sites (known as site 1 and site 2 in this paper) separated by up to 100 kilometers (km) of fiber – approximately 62 miles – with one or more z/OS systems at each site.

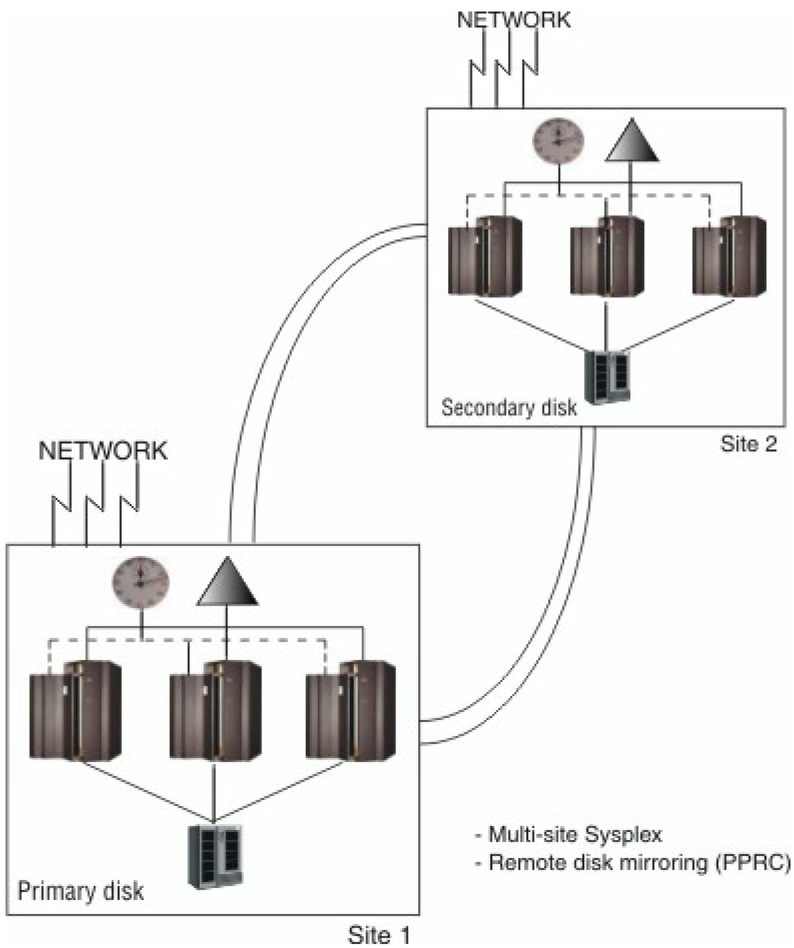

Figure 3: GDPS/PPRC

The multisite sysplex cluster must be configured with redundant hardware (for example, a Coupling Facility and a Sysplex Timer® in each site) and the cross-site connections must be redundant. All critical data resides on storage subsystem(s) in site 1 (the primary copy of data) and is mirrored to the storage subsystem(s) in site 2 (the secondary copy of data) via PPRC synchronous remote copy.

Customers have the capability to configure GDPS/PPRC with up to 100 km of fiber between two sites. An immediate advantage of this extended distance is to potentially decrease the risk that the same disaster will affect both sites, thus permitting customers to recover their production applications at another site.

GDPS/PPRC supports Metro Mirror over Fiber Channel Protocol (FCP). Since Metro Mirror over FCP requires only one protocol exchange compared to two or three exchanges when using Metro Mirror over ESCON®, it is expected that the distance between sites can be increased while maintaining acceptable application performance. The efficiency of the FCP protocol is also expected to lower the total cost of ownership, since two Metro Mirror FCP links between each pair of ESS disk subsystems are considered sufficient for most workloads, allowing a reduction in cross-site connectivity.

Near Continuous Availability of data with HyperSwap

Exclusive to GDPS in the PPRC environment is HyperSwap. This function is designed to broaden the near continuous availability attributes of GDPS/PPRC 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.

With the release of GDPS/PPRC V3.2, the HyperSwap function has been enhanced to exploit the Metro Mirror Failover/Failback (FO/FB) function. For planned reconfigurations, FO/FB may reduce the overall elapsed time to switch the disk subsystems, thereby reducing the time that applications may be unavailable to users. This is demonstrated by the benchmark measurements discussed below. For unplanned reconfigurations, Failover/Failback allows the secondary disks to be configured in the suspended state after the switch and record any updates made to the data. When the failure condition has been repaired, resynchronizing back to the original primary disks requires only the changed data to be copied, thus eliminating the need to perform a full copy of the data. The window during which critical data is left without Metro Mirror protection following an unplanned reconfiguration is thereby minimized.

Planned Reconfiguration support

GDPS/PPRC planned reconfiguration support automates procedures performed by an operations center. These include standard actions to: a) quiesce a system's workload and remove the system from the Parallel Sysplex cluster (e.g., stop the system prior to a hardware change window); b) IPL a system (e.g., start the system after a hardware change window); and c) 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). Standard actions can be initiated against a single system or a group of systems. With the introduction of HyperSwap, you now have the ability to perform disk maintenance and planned site maintenance without requiring applications to be quiesced. Additionally, GDPS/PPRC provides customizable scripting capability for user defined actions (e.g., planned disk maintenance or planned site switch in which the workload is switched from processors in site 1 to processors in site 2.

All GDPS functions can be performed from a single point of control, which can help simplify system resource management. Panels are used to manage the entire remote copy configuration, rather than individual remote copy pairs. This includes the initialization and monitoring of the remote copy volume pairs based upon policy and performing routine operations on installed storage subsystems – disk and tape. GDPS can also perform standard operational tasks, and monitor systems in the event of unplanned outages.

The Planned HyperSwap function is designed to provide the ability to transparently switch all primary disk subsystems with the secondary disk subsystems for planned reconfigurations. During a planned reconfiguration, HyperSwap can provide the ability to perform disk configuration maintenance and planned site maintenance without requiring any applications to be quiesced. Large configurations can be supported, as HyperSwap is designed to provide capacity and capability to swap large number of disk devices very quickly. The important ability to re-synchronize incremental disk data changes, in both directions, between primary / secondary disks is provided as part of this function.

Benchmark measurements using HyperSwap for Planned Reconfiguration

The following table lists the experiences of some GDPS/PPRC reference customers who have implemented HyperSwap for planned reconfigurations. The table also includes the results obtained for one of the configurations at the IBM Test facility.

Reference Customer	Configuration (note 1)	Switch Time (w/o FO/FB)	Switch Time (w/ FO/FB) (note2)
ARZ (Austria)	2300 volume pairs (14 TB)	82-84 seconds	
Postbank (Germany)	1800 volume pairs (32 TB**)	80-84 seconds	
iT Austria (Austria)	650 volume pairs (12 TB**)	32-36 seconds	
iT Austria (Austria)	4200 volume pairs (24 TB)	75 seconds	
iT Austria (Austria)	4500 volume pairs (76 TB**)	75 seconds	
IBM test facility	2900 volume pairs (4.6 TB)	93 seconds	18 seconds

** 3390-9 device type volumes

Note 1: HyperSwap prerequisites are described in section, Prerequisites.

Note 2: Failover/Failback is a function in GDPS R3.2 available March 31, 2005. At the time of the writing, there were no customer benchmark data available.

"Using the GDPS/PPRC HyperSwap technology is a significant step forward in achieving continuous availability. The benefits in our GDPS environments are that a site switch by means of HyperSwap (i.e., dynamic switching of the disk configuration) takes only a few

minutes without application outage. Without HyperSwap we had to shutdown the GDPS/Sysplex, switch the disk configuration, and restart systems and applications. The elapsed time was almost two hours for our largest GDPS with 11 systems and over 11,000 PPRC volume pairs(note) ."

Wolfgang Dungl, Manager of Availability, Capacity and Performance Management
Wolfgang Schott, GDPS Project Manager, iT-AUSTRIA

(Note: original configuration of 11,000 volume pairs based on 3390-3 device type migrated to 4,500 volume pairs based on 3390-9 device type as shown in table above).

Unplanned Reconfiguration support

GDPS/PPRC unplanned reconfiguration support not only can automate procedures to handle site failures, but can also help minimize the impact and potentially mask a z/OS system, processor, Coupling Facility, disk or tape failure, based upon GDPS/PPRC policy. If a z/OS system fails, the failed system and workload can be automatically restarted. If a processor fails, the failed system(s) and their workload can be restarted on other processors.

The Unplanned HyperSwap function 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).

With Unplanned HyperSwap support:

- *Production systems can remain active during a disk subsystem failure. Disk subsystem failures will no longer constitute a single point of failure for an entiresysplex.*
- *Production systems can remain active during a failure of the site containing the primary disk subsystems (site 1), if applications are cloned and exploiting data sharing across the 2 sites. Even though the workload in site 2 will need to be restarted, an improvement in the Recovery Time Objective (RTO) is accomplished.*

Benchmark measurements using HyperSwap for Unplanned Reconfiguration

An unplanned disk reconfiguration test using HyperSwap with failover/failback, conducted at the GDPS solution center demonstrated that the user impact time was only 15 seconds to swap a configuration of 2900 volumes of ESS disks while keeping the applications available, compared to typical results of 30-60 minutes without HyperSwap.

What this benchmark does not show is the failover/failback capability to only copying the changed data instead of the entire disk during the resynchronization process. This can save significant time and network resources.

GDPS/PPRC HyperSwap Manager

GDPS/PPRC HyperSwap Manager (GDPS/PPRC HM) expands zSeries Business Resiliency to customers by providing a single-site near continuous availability solution as well as a multi-site entry-level disaster recovery solution.

Within a single site, GDPS/PPRC HyperSwap Manager extends Parallel Sysplex availability to disk subsystems by masking planned and unplanned disk outages caused by disk maintenance and disk failures. It also provides management of the data replication environment and automates switching between the two copies of the data without causing an application outage, therefore providing near-continuous access to data. Figure 4 shows an example of a GDPS/PPRC HM configuration.

In the multisite environment, GDPS/PPRC HyperSwap Manager provides an effective entry-level disaster recovery offering for those zSeries customers that have the need for very high levels of data availability. Value is further enhanced by being able to use specially priced Tivoli System Automation and NetView products. In addition, a customer can migrate to the full function GDPS/PPRC capability across multiple sites as business requirements demand shorter Recovery Time Objectives provided by a second site. The initial investment in GDPS/PPRC HM is protected when customers choose to move to full-function GDPS/PPRC by leveraging the existing GDPS/PPRC HM implementation and skills.

GDPS/PPRC HM simplifies the control and management of the Metro Mirror (PPRC) environment for both z/OS and Open Systems data. This reduces storage management costs while reducing the time required for remote copy implementation.

GDPS/PPRC HM provides support for FlashCopy®. GDPS/PPRC HM can be set up to automatically take a FlashCopy of the secondary disks before resynchronizing the primary and secondary disks following a Metro Mirror suspension event, ensuring a consistent set of disks are preserved should there be a disaster during the re-synch operation.

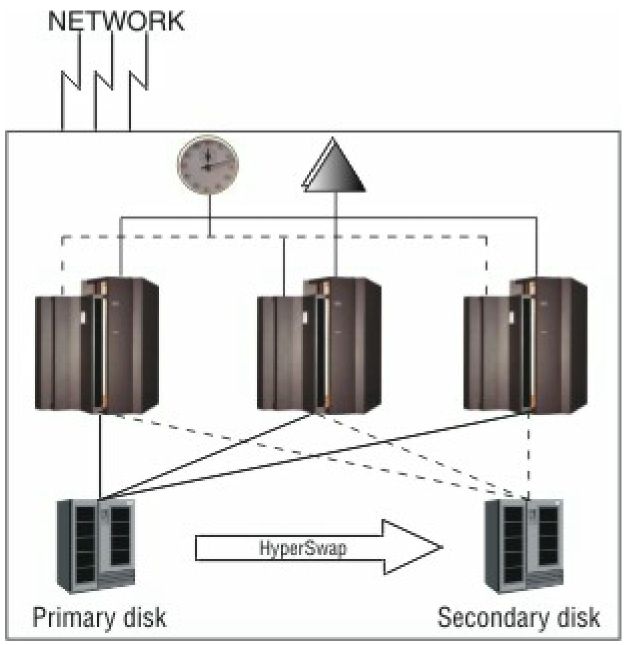

Figure 4: GDPS/PPRC HyperSwap Manager Configuration

Near continuous availability of data within a single site

A Parallel Sysplex environment has been designed to reduce outages by replicating hardware, operating systems and application components. In spite of this redundancy having only one copy on the data is an exposure. GDPS/PPRC HyperSwap Manager is designed to provide continuous availability of data by masking disk outages caused by disk maintenance and/or failures. For example, if normal processing is suddenly interrupted when one of the disk subsystems experiences a hard failure, thanks to GDPS the applications are masked from this error because GDPS detects the failure and autonomically invokes HyperSwap. The production systems continue using data from the mirrored secondary volumes. Disk maintenance can also be similarly performed without application impact by executing HyperSwap command.

Near continuous availability of data/Disaster Recovery solution at metropolitan distances

In addition to the single site capabilities, in a two site configuration GDPS/PPRC HyperSwap Manager provides an entry-level disaster recovery capability at the recovery site. GDPS/PPRC HM uses the Freeze function described in the section, Need for Data Consistency. The Freeze function is designed to provide a consistent copy of data at the recovery site from which production applications can be restarted. The ability to simply restart applications helps eliminate the need for lengthy database recovery actions. Automation to stop and restart the operating system images available with the full-function GDPS/PPRC is not included with GDPS/PPRC HyperSwap Manager.

GDPS/PPRC HyperSwap Manager Prerequisites

The GDPS/PPRC HyperSwap manager can use as prerequisites IBM Tivoli System Automation for GDPS/PPRC HyperSwap Manager with NetView, V1.1, which provides the System Automation and NetView requirements, or IBM Tivoli System Automation for GDPS/PPRC HyperSwap Manager, V1.1 together with the full-function IBM Tivoli NetView for z/OS product. These new customized products together with the GDPS/PPRC HM offering are designed to bring new levels of affordability to customers who do not require the full-function products. This makes GDPS/PPRC HM an excellent entry-level offering for those single-site or multisite installations that need higher levels of IT availability.

GDPS/XRC

Extended Remote Copy (XRC, recently renamed to IBM TotalStorage z/OS Global Mirror) is a combined hardware and z/OS software asynchronous remote copy solution. Consistency of the data is maintained via the Consistency Group function within the System Data Mover. GDPS/XRC includes automation to manage remote copy pairs and automates the process of recovering the production environment with limited manual intervention, including invocation of CBU, thus providing significant value in reducing the duration of the recovery window and requiring less operator interaction.

GDPS/XRC is capable of the following attributes:

- *Disaster recovery solution*
- *RTO between an hour to two hours*
- *RPO less than two minutes, typically 3-5 seconds*
- *Protects against localized as well as regional disasters (distance between sites is unlimited)*
- *Minimal remote copy performance impact*

GDPS/XRC Topology

The physical topology of a GDPS/XRC, referring to Figure 5, consists of production system(s) in site 1. The production systems could be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster[5]. Site 2, (the recovery site) can be located at a virtually unlimited distance from site 1 (the production site).

[5] Note that if there is a base or Parallel Sysplex cluster in the production site, it does not span across sites 1 and 2.

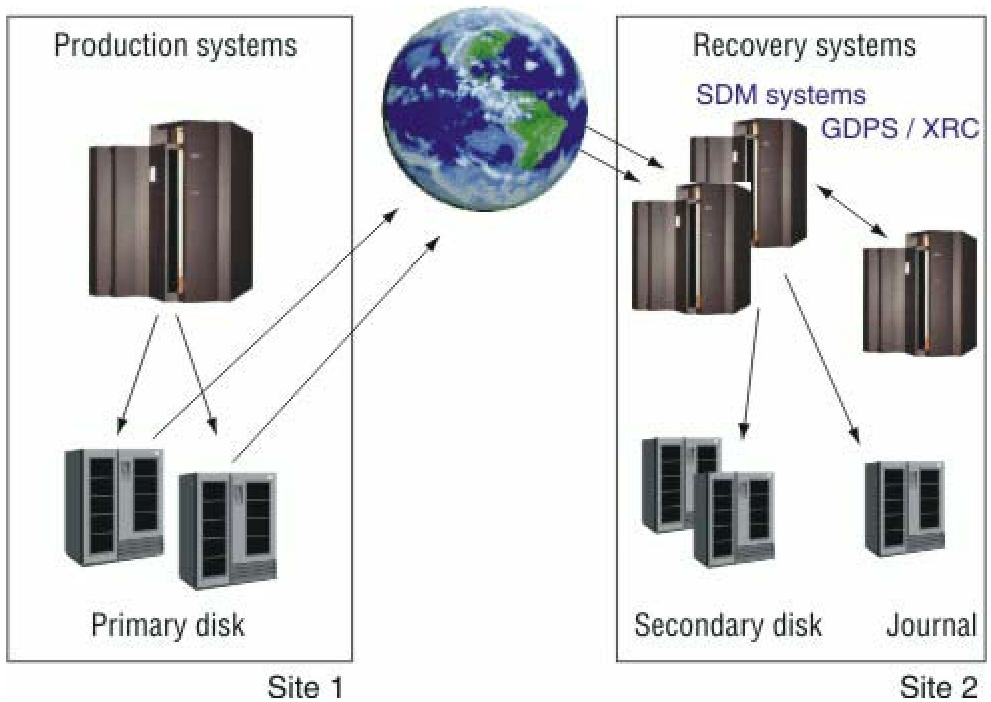

Figure 5: GDPS/XRC

During normal operations, the XRC System Data Mover (one or more) execute in site 2 and are in a Base Sysplex with the GDPS controlling system (refer to section, **GDPS Systems** for a definition of the GDPS controlling system). All critical data resides on storage subsystem(s) in site 1 (the primary copy of data) and is mirrored to the storage subsystem(s) in site 2 (the secondary copy of data) via XRC asynchronous remote copy.

Planned Reconfiguration support

All the planned reconfiguration actions described in the section, GDPS/PPRC are provided by GDPS/XRC for the System Data Mover (SDM) Sysplex in site 2. For example, GDPS/XRC will manage the temporary relocation of the SDM, if it is needed. By managing the SDM Sysplex, GDPS/XRC can also manage the z/OS Metro Mirror remote copy configuration.

As noted in the section, **Lessons learned about IT survival**, it is recommended that the D/R solution be based on as much automation as possible to minimize the dependency on key skills being available to recover from a disaster. GDPS/XRC is designed to automate the process of recovering the production environment with minimal manual intervention, which can provide significant value in minimizing the duration of the recovery window.

Coupled System Data Mover support

Coupled Extended Remote Copy (CXRC) expands the capability of XRC, so that customers who have configurations consisting of thousands of primary volumes can recover all their volumes to a consistent point in time. A single SDM can typically manage approximately 1000 to 2000 volume pairs (based on the write I/O rate), CXRC can provide the scalability that is required to support larger XRC configurations. CXRC allows customers the capability of "coupling" multiple XRC sessions together into a master session. CXRC coordinates the consistency of data for "coupled" sessions in a master session, which can allow recovery of data for all the volumes in the "coupled" sessions to a consistent point in time.

Commands are now capable of being executed in parallel across multiple SDMs in a GDPS/XRC configuration, which can provide improved scalability. The parallelism is across multiple SDMs, provided there is only one SDM per z/OS image. If there are multiple SDMs per z/OS image, processing is sequential by SDM within the z/OS image.

Functional Highlights (GDPS/PPRC and GDPS/XRC)

The following functions are supported by both GDPS/PPRC and GDPS/XRC:

Peer-to-Peer Virtual Tape Server (PtP VTS) support

GDPS also supports Peer-to-Peer Virtual Tape Server. By extending GDPS support to data resident on tape, the GDPS solution is intended to provide continuous availability and near transparent business continuity benefit for both disk and tape resident data. Enterprises may no longer be forced to develop and utilize processes that create duplex tapes and maintain the tape copies in alternate sites. For example, previous techniques created two copies of each DBMS image copy and archived log as part of the batch process and manual transportation of each set of tapes to different locations.

Operational data, or data that is used directly by applications supporting end users, is normally found on disk. However, there is another category of data that 'supports' the operational data, which is typically found on tape subsystems. Support data typically covers migrated data, point in time backups, archive data, etc. For sustained operation in the recovery site, the support data is indispensable. Furthermore, several enterprises have mission critical data that only resides on tape.

The PtP VTS provides a hardware-based duplex tape solution and GDPS can automatically manage the duplexed tapes in the event of a planned site switch or a site failure. Control capability has been added to allow GDPS to "freeze" copy operations, so that tape data consistency can be maintained across GDPS managed sites during a switch between the primary and secondary VTSs.

FlashCopy support

FlashCopy, available on the IBM TotalStorage DS Family and IBM TotalStorage Enterprise Storage Server (ESS), 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. Prior to the release of FlashCopy v2 in 2003, both source and target volumes had to reside on the same logical subsystem. Since this constraint has been removed with FlashCopy v2, GDPS will now allow a FlashCopy from a source in one LSS to a target in a different LSS within the same disk subsystem.

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.

FlashCopy can operate in either of two modes, the COPY mode which runs a background copy process and the NOCOPY mode that suppresses the background copy. Previously GDPS/PPRC and GDPS/XRC have provided support for both COPY and NOCOPY.

With the release of GDPS/PPRC V3.2 and GDPS/XRC V3.2, two FlashCopy enhancements are now available. The first enhancement is support for NOCOPY2COPY which allows changing an existing FlashCopy relationship from NOCOPY to COPY. This gives you the option of always selecting the NOCOPY option of FlashCopy and then converting it to the COPY option when you want to create a full copy of the data in the background at a non-peak time. Another FlashCopy enhancement available with GDPS V3.2 is support for Incremental FlashCopy. This provides the capability to refresh a volume in a FlashCopy relationship and reduce background copy time when only a subset of the data has changed. With Incremental FlashCopy, the initial relationship between a source and target is maintained after the background copy is complete. When a subsequent FlashCopy establish is initiated, only the data updated on the source since the last FlashCopy is copied to the target. This reduces the time needed to create a third copy, thus giving you the option to perform a FlashCopy on a more frequent basis.

Recent Performance Enhancements

- *Concurrent activation of Capacity Backup (CBU) can now be performed in parallel across multiple servers which results in an improved RTO.*
- *Analysis has shown that PPRC and XRC commands issued by GDPS can generate a large number of Write to Operator messages (WTOs) that may cause WTO buffer shortages and temporarily adversely impact system performance. The Message Flooding Automation function can substantially reduce the WTO message traffic and help improve system performance by suppressing redundant WTOs.*
- *With the release of GDPS/PPRC V3.2, the HyperSwap function has been enhanced to exploit the PPRC Failover/Failback function. For more details, refer to HyperSwap section.*

GDPS Support for Heterogeneous Environments

Management of zSeries Operating Systems

In addition to managing images within the base or Parallel Sysplex cluster, GDPS can now also manage a customer's other zSeries production operating systems – these include z/OS, Linux for zSeries, z/VM®, and VSE/ESA™. The operating systems have to run on servers that are connected to the same Hardware Management Console (HMC) Local Area Network (LAN) as the Parallel Sysplex cluster images. 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 zSeries images can either run as a logical partition (LPAR) or as a guest under z/VM.

GDPS/PPRC Management for Open Systems LUNs (Logical Unit Number):

GDPS/PPRC technology has been extended to manage a heterogeneous environment of z/OS and Open Systems data. If installations share their disk subsystems between the z/OS and Open Systems platforms, GDPS/PPRC can manage the Metro Mirror and FlashCopy for open systems storage. GDPS/PPRC is also designed to provide data consistency across both z/OS and Open Systems data. This 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/PPRC Multi-Platform Resiliency for zSeries

GDPS/PPRC has been enhanced to provide a new function called "GDPS/PPRC Multi-Platform Resiliency for zSeries." This function is especially valuable for customers who share data and storage subsystems between z/OS and z/VM Linux guests on zSeries – for example, an application server running on Linux on zSeries 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 zLinux. GDPS/PPRC can now provide that. z/VM 5.1 provides a 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 by Peer-to-Peer Remote Copy (PPRC, or 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. 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.

GDPS/PPRC will provide the reconfiguration capabilities for the Linux on zSeries servers and data in the same manner as for z/OS systems and data. To support planned and unplanned outages, GDPS provides the recovery actions such as the following examples:

- *Re-IPL in place of failing operating system images*
- *Site takeover/failover of a complete production site*
- *Coordinated planned and unplanned HyperSwap of disk subsystems, transparent to the operating system images and applications using the disks.*
- *Linux node or cluster failures*
- *Transparent disk maintenance or failure recovery with HyperSwap across z/OS and Linux applications*
- *Data consistency with freeze functions across z/OS and Linux*

IBM Global Services (IGS) Offerings

The following GDPS services and offerings are provided by 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.

Remote Copy Management Facility™ (RCMF)

With this service, the RCMF/PPRC or RCMF/XRC automation to manage the remote copy infrastructure will be installed, the automation policy customized, and the automation verified along with providing operational education for the enterprise.

GDPS/PPRC HyperSwap Manager

IBM Implementation Services for GDPS/PPRC HyperSwap Manager helps simplify implementation by working with you to get GDPS/PPRC HyperSwap Manager and its prerequisites up and running with limited disruption to your business. On-site planning, configuration, implementation, testing, and education will help make your new IBM Implementation Services for GDPS/PPRC HyperSwap Manager solution accessible in the most efficient manner.

GDPS

IBM Implementation Services for GDPS/PPRC or GDPS/XRC will assist you with planning, configuration, automation code customization, testing, onsite implementation assistance, and training in IBM's GDPS solution. Either option supports Peer-to-Peer Virtual Tape Server (PtP VTS) form of tape data mirroring.

Prerequisites

For IBM to perform these services, you must have the following elements. The prerequisites listed may not contain all of the requirements for this service. For a complete list of prerequisites, consult your sales representative.

Prerequisites	RCMF	GDPS
Supported version of z/OS or z/OS.e z/VM V5.1 or higher (note 1)	✓	✓
IBM Tivoli System Automation for Multiplatforms V1.2 or higher (note 1)		✓
IBM Tivoli System Automation for z/OS V2.2 or higher (note 2)	✓	✓
IBM Tivoli NetView V5.1 or higher (note 2)	✓	✓
Storage subsystem with PPRC Freeze function (CGROUP Freeze/RUN) (note 3 + note 4)		✓
XRC support with Unplanned Outage support (note 4)		✓
Multisite Base or Parallel Sysplex (GDPS/PPRC)		✓
Common Timer Reference (Sysplex Timer) for XRC	✓	✓

Note 1: z/VM is a prerequisite if GDPS/PPRC Multiplatform Resiliency is required

Note 2: For GDPS/PPRC HyperSwap Manager, the following software products are required:

- *IBM Tivoli System Automation for GDPS/PPRC HyperSwap Manager with NetView, V1.1 or higher, or*
- *IBM Tivoli NetView for z/OS V5.1 or higher together with one of the following:*
  - *IBM Tivoli System Automation for GDPS/PPRC HyperSwap Manager, V1.1 or higher, or*
  - *IBM Tivoli System Automation for z/OS V2.2 or higher*

Note 3: GDPS/PPRC HyperSwap requires PPRC support for Extended CQuery. GDPS/PPRC Management of Open Systems LUNs requires support for Open PPRC, management of Open PPRC via CKD device addresses, and Open PPRC SNMP alerts.

Note 4: GDPS FlashCopy support requires FlashCopy V2 capable disk subsystems

In addition, visit **ibm.com**/servers/storage/support/solutions/bc/index.html for a listing of all recommended maintenance that should be applied, as well as review Informational APAR II12161

GDPS/PPRC at work in a real disaster

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

VPC AB, a security depository and clearing (CSD) organization, has a GDPS/PPRC configuration in production – a 3-way Parallel Sysplex cluster with 100 volumes being mirrored using PPRC (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:
ibm.com/servers/eserver/zseries/gdps

zSeries Business Resiliency Web site:
ibm.com/servers/eserver/zseries/resiliency/features.html

For an overview of how planning, implementation and proper exploitation of zSeries Parallel Sysplex clustering technology can enable your business achieve near continuous availability, refer to "Five Nines / Five Minutes – Achieving Near Continuous Availability" at:
ibm.com/servers/eserver/zseries/pso/

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

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

For complete results of survey conducted in 2001 by Contingency Planning Research, refer to:
contingencyplanningresearch.com/2001%20Survey.pdf

For additional information on GDPS, contact your IBM representative or e-mail
gdps@us.**ibm.com**

IBM Corporation
Software Communications
Route 100
Somers, NY 10589
U.S.A.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Information concerning non-IBM products was obtained from the suppliers of their products or their published announcements. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

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.

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.

The information in this document is intended to provide guidance for those implementing GDPS. It discusses findings based on a solution that was created and tested under laboratory conditions. These findings may not be realized in all customer environments. Implementation in such environments may require additional steps, configurations, and performance analysis. The information herein is provided "AS IS" with no warranties, express or implied. This information does not constitute a specification or form part of the warranty for any IBM or products. The users of this document should always check the latest release information in the product Readme file(s) and check the product Web pages for the latest updates and findings.

GF22-5114-06