Executive Summary
September, 2011

GDPS Active-Active Executive Overview

David Raften
raften@us.ibm.com

Noshir Dhondy
dhondy@us.ibm.com

Sim Schindel
schindel@us.ibm.com

Table of Contents

Introduction 2
GDPS Suite of offerings 2
Configurations 5
Active-Standby deep dive 6
GDPS Automation 7
Prerequisites 8
Futures 8
Additional Information 9

Introduction

Every company that is serious about availability has a disaster recovery plan. This can be to recover using tape backup with typically up to three days recovery time objective (RTO) to bring applications online with up to 24 hours of data loss, or recovery point objective (RPO). The other option is through disk remote copy with typically RTO of six hours, 1-2 hours with Geographically Dispersed Parallel Sysplex™ (GDPS®) automation, and RPO of seconds to minutes, depending upon bandwidth and update activity. But there is a growing demand for the recovery time objective to be seconds, especially with volatile financial markets. For example, on August 18, 2011, the Dow Jones average fell 400 points in the first hour of trading. If a financial company had an outage at that time, it could have been devastating.

GDPS/Active-Active provides a rapid recovery disaster recovery solution. It is designed to provide an automated disaster recovery solution meeting a recovery time objective of seconds, not hours, at virtually unlimited distances.

GDPS Suite of offerings

The GDPS suite of offerings provides different solutions to meet different business requirements. All of which are based on disk replication techniques. GDPS automation was then added on top of that to help manage the remote copy environment. The different GDPS offering options includes:

GDPS/PPRC HyperSwap™ manager, based on synchronous PPRC disk replication, provides an entry level multisite disaster recovery solution. HyperSwap provides the ability to dynamically switch to secondary volumes without requiring applications to be quiesced. Typically done in 3–30 seconds, this provides near-continuous data availability for planned actions and unplanned events. Using HyperSwap technology, GDPS/PPRC HyperSwap Manager provides disk remote copy management, data consistency, and HyperSwap for data availability. This can be across a single or multiple sites up to 200 km away. 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/PPRC HyperSwap Manager implementation to the full-function GDPS/PPRC capability as business requirements demand shorter recovery time objectives.

GDPS/PPRC, based on synchronous PPRC disk replication, has all the function as GDPS/PPRC HM, and also is designed to fully automate the recovery at the remote site. This includes disk reconfiguration, managing servers, Sysplex resources, CBU, activation profiles, etc. GDPS/PPRC is designed to be a near-continuous availability and disaster recovery solution. GDPS/PPRC is application and data independent. When managing the data it can be used to provide a consistent recovery for z/OS® as well as non-z/OS data. 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-IBM System z® data to allow rapid business resumption.

GDPS/PPRC 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)*

GDPS/XRC, based on asynchronous z/OS Global Mirror disk replication, is a highly scalable remote copy solution for z/OS and Linux® on System 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:

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

GDPS/Global Mirror, based on asynchronous Global Mirror disk replication, is a remote copy solution for z/OS and non-z/OS data. Global Mirror is a solution 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-System z data. This is important when a multi-tier application has dependencies upon multiple operating system architectures.

GDPS/Global Mirror is capable of:

- *Disaster recovery solution for System z and non-System 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 Solutions are available for businesses requiring the benefits of both synchronous and asynchronous remote copy. Synchronous remote copy using GDPS/PPRC and GDPS/PPRC 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/Global Mirror provides benefits such as protection from regional disasters with little

to no application impact. To provide for this requirement, GDPS supports two solutions based upon a three site configuration. These solutions help provide maximum availability across the widest range of possible scenarios.

- *GDPS/PPRC or GDPS/PPRC HM combined with GDPS/XRC is called GDPS Metro/ z/OS Global Mirror, or GDPS/MzGM. The GDPS/MzGM solution is a multi-target replication solution.*
- *GDPS/PPRC or GDPS/PPRC HM combined with GDPS/Global Mirror is called GDPS Metro/Global Mirror, or GDPS/MGM. GDPS/MGM is a cascading replication solution that also provides a solution for both System z and non-z data.*

GDPS MGM and GDPS MzGM are capable of:
- *Potential for zero data loss*
- *HyperSwap for disk availability*
- *Protection from regional events*
- *No impact to end user response time*
- *Disaster recovery solution for System z and non-System z data (with GDPS MGM)*

With GDPS/PPRC, it is possible that a site failure can be recovered from in seconds. This requires:
- *Production workload actively running on two sites (multi-site configuration)*
- *All critical data is remote copied and HyperSwap enabled*
- *All critical CF structures are duplexed*
- *Applications are Parallel Sysplex® datasharing enabled*
- *The site failure first affects the disk, resulting in a HyperSwap*

In a multi-site workload configuration, the signal latency between sites will potentially impact online workload throughput and batch duration resulting in the sites typically being separated by no more than approximately 20 km fiber distance.

GDPS disk remote copy solutions generally have a minimum recovery time of approximately one hour after a site failure from the point when the operator initiates the takeover action.

GDPS/Active-Active is a fundamental paradigm shift from a failover model to a continuous availability model. GDPS/Active-Active as a software based asynchronous mirroring solution for select z/OS workloads, designed to maintain a consistent copy of data within IMS™ and within DB2® at virtually unlimited distances, recovery time measured in seconds, with minimal impact to application response time. GDPS/Active-Active is disk independent. Current support is for the Active-Standby configuration.

GDPS/Active-Active is capable of:

- *Disaster recovery solution forDB2 and IMS data*
- *Application level granularity*
- *Site switch in seconds (once the event was detected)*
- *RPO in seconds(with the ability to issue reports that can be generated showing orphaned data)*
- *Protects against metro and regional disasters (unlimited distance between sites)*
- *Minimal response time impact*

The goal is to have two sites, separated by virtually unlimited distances, running the same applications and having the same data to provide cross-site workload balancing and Continuous Availability / Disaster Recovery.

GDPS automation, primarily operating at a workload level, helps manage the availability of selected workloads and the routing of transactions for these workloads. The GDPS/Active-Active product, a component of the GDPS/Active-Active solution, acts primarily as the coordination point or controller for these activities; including being a focal point for operating and monitoring the solution and readiness for recovery.

More information on GDPS/Active-Active is available in announcement 611-023, "IBM GDPS active/active continuous availability," dated May 24, 2011. Also, see "GDPS Family – An Introduction to Concepts and Capabilities" at
http://www.redbooks.**ibm.com**/redpieces/abstracts/sg246374.html?Open

Configurations

The first release of GDPS/Active-Active delivers the Active/Standby configuration. At any given point in time, transactions for the workload are only being routed to one site which is, at that point in time, the Active instance of the workload. The workload runs in the other site in Standby mode.

IBM has made a Statement of Direction to deliver support for GDPS/Active-Active workloads that run in Active/Query configuration. Here, a given workload is actually running in both sites. Any transactions that have the potential of updating the data belonging to a workload (read/write transactions) are routed to the site with the active instance of the workload. However, transactions that only query the data (read-only transactions) can be routed to be processed either on the active site or on the query-only site of the workload to provide a degree of workload balancing between the two sites.

Active-Standby deep dive

The GDPS/Active-Active environment is comprised of two production Sysplexes (sites) in different locations. For each workload that is to be managed by GDPS/Active-Active, at any given point in time one of the Sysplexes will be the active sysplex and the other will act as standby. When there are multiple GDPS-managed workloads, a given sysplex can be the active sysplex for one workload while it is standby for another. It is the routing for each workload that determines which sysplex is active and which sysplex is standby for a given workload. As such, in environments where there are multiple workloads, there is no concept as Active system or sysplex. There is a system or sysplex that is the currently active one for a given workload.

Both the active and the standby production instances are IPLed and running. What makes a Sysplex active or standby is whether transactions are currently being routed to that Sysplex for that workload. The standby Sysplex is waiting for work, and would be able to process work at any time should there be a planned or unplanned workload switch, resulting in transactions being routed to this Sysplex. If there is a workload switch, the standby Sysplex will become the active Sysplex for the given workload.

On the active Sysplex are one or more replication capture engines. This is software that captures all updates to the databases used by the GDPS-managed workload and forwards them to the standby Sysplex. On the standby Sysplex is the apply engine which receives the updates sent by the capture engine and it immediately applies them to its copy of the database.

The data replication in the GDPS/Active-Active environment is asynchronous. The workload can perform a database update and this write operation can complete independent of the replication process. When a transaction commits changes, the replication software (InfoSphere™ Replication Server for z/OS, or the InfoSphere IMS Replication for z/OS) captures the updates as soon as they are committed and sends them via MQSeries® or TCP/IP to the standby site where it gets applied in real time to a running copy of the database. If replication is disrupted for any reason, when restored, the replication engines have logic to know where they left off and are able to transmit only those changes made after the disruption.

Planned workload switches can take place without any data loss by quiescing the workload on the sending side, draining all updates, and applying them on the receiving side. Because replication is asynchronous, unplanned switches will typically have some data captured but not yet transmitted and therefore not yet applied on the target Sysplex. This is the latency, or

RPO. With a correctly-sized, robust transmission network, during normal operations the RPO is expected to be as low as just a few seconds.

It is possible to run other workloads not managed by GDPS/Active-Active on the same production systems that run Active-Active workloads.

GDPS 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. The Systems Programmers can concentrate on pinpointing the cause of the outage, rather than trying to get everything up and running again.

All of this is just for a system outage. In a disaster recovery situation one also needs to verify status of both sites, identify any issue that could prevent a switch to the standby site, modify the routing of transactions to the standby site, invoke Capacity Back-Up (CBU), quiesce discretionary LPARs, verification, 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.

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 replication environment. It includes various standard actions which can be initiated against a single system or a group of systems to:

1) *Providing a central point of monitoring for the GDPS/A-A environment with a GUI interface including status indicators.*
2) *Starting and stopping the replication for each of the workloads*

3) *Managing the load balancer, IBM Multi-site Workload Lifeline,*
4) *Managing the servers at each site, including managing IPL parameters, IPLing, and invoking Capacity BackUp (CBU) or On/Off Capacity Upgrade on Demand (OOCoD).*
5) *Providing the ability to perform a controlled workload site switch for both planned and unplanned workload interruptions, with minimal or no data loss.*
6) *Quiesce a system's workload and remove the system from the Parallel Sysplex cluster (stop the system prior to a hardware change window);*
7) *IPL a system (start the system after a hardware change window)*
8) *Customizable scripting capability for user defined actions*

Prerequisites

The required software prerequisites are listed in the GDPS Web site:
ibm.com/systems/z/gdps/getstarted

Futures

IBM announced a "Statement of Direction" indicating plans to build upon the Active/Standby configuration with the Active/Query configuration. This will provide better the ability to selectively query data in either site. It can make better use of existing capacity at the remote site to process read-only transactions that can tolerate data that may be slightly old. If the latency, the time it takes for an update on the "Active" site to be replicated on to the "Standby" site, is within the user-defined toleration limit, then a read-only transaction can be routed to the Standby site to execute. If the latency gets longer, maybe because of network bottlenecks, then the read-only transactions will be routed to the Active site until the latency value drops again.

Additional Information

GDPS home page:
ibm.com/systems/z/gdps

IBM eServer™ zSeries® Business Resiliency Web site:
ibm.com/systems/z/resiliency

For an overview of System z Parallel Sysplex clustering technology and how it can enable your business achieve near continuous availability, refer to **ibm.com**/systems/z/pso/

For an overview of Server Time Protocol (STP) and how it can help in a GDPS environment, refer to **ibm.com**/systems/z/pso/stp.html

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 11th and Implications for Business Continuity prepared by the Securities and Exchange Commission, refer to: 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

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.

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
10/2011

ZSW03194-USEN-00