

Updates to IBM's GDPS Improve Data Storage Protection

Josh Krischer

IBM's Geographical Dispersed Parallel Sysplex is a fully automated disaster recovery product for z/OS mainframes. Recent enhancements give users greater flexibility, availability and functionality.

WHAT YOU NEED TO KNOW

GDPS is an IBM mainframe automated disaster recovery technique. GDPS/PPRC provides resource sharing, workload balancing, continuous availability and disaster recovery protection, in addition to easing and reducing the time required for site reconfigurations in planned outages. It allows a reduction of manpower and skill required for these functions. It also supports storage in multiple platform environments. Companies reliant on mainframes and planning restoration after disasters with short recovery time objectives and recovery point objectives should investigate this option.

STRATEGIC PLANNING ASSUMPTION(S)

Till the end of 2007 the number of GDPS deployments will double (0.7 probability)

ANALYSIS

Geographical Dispersed Parallel Sysplex (GDPS) is a disaster recovery system delivered by IBM Global Services. It includes resource and data sharing, workload balancing, continuous availability and data mirroring for data availability — all provided by Parallel Sysplex clustering technology.

GDPS is a multi-site application that offers fast recovery time and highly automated control. It manages application availability in and across sites for planned (for example, site maintenance) or unplanned reconfigurations, like site failure or disaster. It supports both synchronous Peer-to-Peer Remote Copy (PPRC) and asynchronous Extended Remote Copy (XRC). GDPS also supports the Peer-to-Peer Virtual Tape Server (PtP VTS) form of remote copying tape data. It is independent of disk vendors, as long as the vendor meets the required level of compatibility for PPRC and XRC architectures.

The following details the major enhancements in GDPS/PPRC.

Open Logical Unit Number Management

The PPRC Level 4 architecture defines the capability for zSeries to issue PPRC commands against open system logical unit numbers (LUNs). It also includes the capability for the disk subsystem to report mirroring errors via simple network management protocol (SNMP) traps.

For zSeries to issue PPRC commands against open system LUNs, there must be some CKD formatted capacity available on the disk subsystem. The GDPS/PPRC open LUN management capability allows PPRC management for both zSeries and open systems data from a single point of control. The GDPS/PPRC freeze function is implemented for open LUNs and is designed to guarantee PPRC secondary data consistency. This, in turn, will allow the open systems to be "restarted" using the frozen secondary data, avoiding lengthy recovery actions. The system can ensure secondary data consistency for either zSeries or open data, or for customers that have cross-platform data dependencies, by freezing open and zSeries data together. Although GDPS can automate the restart of zSeries systems using the secondary copy of data, the manual restart of other platform servers is required. GDPS provides a complete data recovery solution to manage PPRC and secondary data consistency on behalf of heterogeneous systems.

The "GDPS/PPRC Multiplatform Resilience For zSeries" lite function (xDR) expands GDPS/PPRC functionality to zLinux System Automation for Linux "Tivoli System Automation for Multiplatforms" to clusters, initially running as zVM guests. In the future, it will also support Linux running natively in logical partitions.

This will provide an end-to-end solution for customers with distributed hybrid applications that span z/OS and zLinux (for example, WebSphere or SAP).

One key function provided by xDR is the coordinated HyperSwap for both z/OS GDPS sysplex data and for the zLinux cluster data. The functions included in the list, together with HyperSwap capability, improves zLinux availability for two-tier business-critical applications like SAP. It provides a single methodology for managing data recovery and coordinated disaster recovery for what is in effect a single application running over two platforms.

GDPS/PPRC Distance Extended to 100 Km

In the past, the maximum distance between GDPS/PPRC sites was limited to 40 km. Support for External Time Reference (ETR) and Inter System Channel 3 (ISC-3) links has been extended to 100 km. The ETR links attach the 9037 Sysplex Timer to zSeries or S/390 9672 servers. The ISC-3 links, operating in "Peer" mode, and supported on all zSeries servers, connect z/OS or OS/390 systems to a coupling facility. A PPRC over Fibre Channel can link with Dense Wavelength Division Multiplexer (DWDM) to support an even greater distance between sites.

This enhancement looks good but is mainly academic because the maximum fiber distance between a pair of 9037 Sysplex Timers in an expanded availability configuration remains 40 km. Users planning to use the extended distance have to locate one of the Sysplex Timers in an intermediary site away from one of the two sites, while maintaining the maximum 40-km limit. A new Sysplex Timer to support 100 km distance is planned for 1H06.

HyperSwap

The HyperSwap function is probably the most important improvement for IBM mainframes. Disk subsystems contain very large amounts of data and, in an integrated and complex applications environment, the disk becomes a single point of failure for the entire Parallel Sysplex. The HyperSwap mechanism is similar to the "Swap" command used with tape technology to mount a tape on a second drive in the event of an I/O error. The HyperSwap function can help significantly reduce the time needed to switch disks between sites in case of a disaster or disk subsystem malfunction on the primary site.

Prior to HyperSwap, switching disks from the PPRC primary site to the secondary one required all systems using the disks to be taken down. The disks would then have to be switched and systems brought back again using the secondary disks. For the data sharing environment, this results in a sysplex-wide outage. HyperSwap eliminates the outage and reduces the application impact time of the disk switch down to seconds. This function is controlled by automation, allowing all aspects of the site switch to be controlled via GDPS. Two modes of operation are used — planned and unplanned:

? Planned HyperSwap provides the ability to transparently switch all primary PPRC disk subsystems with the secondary one for planned maintenance without requiring any applications to be lost. It has the capability to swap a large number of disk devices and the ability to re-synchronize incremental disk data changes, in both directions, between primary and secondary subsystems.

New "failover/failback" functionality is delivered with GDPS/PPRC 3.2. This enhancement exploits the PPRC failover/failback capability and significantly reduces the switch times. In an IBM test environment, the time to HyperSwap without failover/failback using 2900 pairs of mirrored devices was 90 seconds. With failover, this dropped to 18 seconds.

? The unplanned reconfiguration is activated when a local disk malfunctions or during a "disaster." The switch time in this case will typically be 50 percent of the planned time because the GDPS will not have to make housekeeping actions on the primary site. Statistically, most

outages are caused by disk subsystem failures and not as a result of a disaster. Therefore, this functionality significantly increases site availability. HyperSwap cannot swap Unix and Windows data and this data will be “frozen” to maintain data consistency.

Capacity Back-Up

Best practice in mainframe disaster recovery planning is to purchase provisioning for additional central processors which can be activated in an emergency. This capacity back-up (CBU) option is available on IBM mainframes and can provide sufficient processing power to run business critical applications should the primary site be knocked out. This can be done with reasonable service-level agreements without purchasing the additional power upfront but paying much lower price for the provision. The CBU feature has the ability to add temporary capacity when it is lost elsewhere in the organization. GDPS-CBU management automates the process of dynamically adding provisioned central processors, thereby helping to minimize the outage time. A typical time to issue an automated call to IBM service center, to verify the contract and to activate the CBU central processors, takes approximately five minutes. Both the GDPS/PPRC and GDPS/XRC configurations support this capability. CBUs are available for both z/OS and Linux.

Planned GDPS Enhancements

- ? GDPS single site HyperSwap to improve local availability. The official name is GDPS/PPRC HyperSwap Manager. It is intended as a solution for continuous data availability on a single site, or as an entry-level disaster recovery solution across two sites.
- ? Disk subsystem granularity in HyperSwap. The current plan is that single disk subsystem switch will be available with GDPS/PPRC HyperSwap manager intended for customers that use it on a single site to provide continuous data availability without having to invoke disaster recovery for a single component outage.
- ? New Sysplex Timer to allow practical distances of up to 100 km between two Sysplex Timers.
- ? “Freeze and Stop” command — improved synchronization between the storage and the coupling facility structures with the target to achieve no data loss.

Peer-to-Peer Virtual Tape Server Support

There is certain data on disk and on tape that are related to each other. For a PPRC suspension (freeze event), with a “Freeze and Go” policy, it is possible that the secondary copy of disk is frozen but tape continues to be used and the secondary copy of tape goes out of synch with disk. For example, a new tape volume could be written but the frozen database will not reflect this, causing the volume to appear as scratch. GDPS support of PtP VTS provides consistency across tape and disk so that data is not out of synch when systems need to be restarted using the secondary site. It also saves tape duplication and avoids the need for physical transport to remote sites.

Management of "Other" zSeries Operating Systems

GDPS can now manage a customer's other zSeries production operating systems — such as 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) 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 guest under z/VM (now supported) or in logical partition (to be supported in the future). zLinux guests under zVM is the better and preferred

solution because it elevates these systems to "equality" with z/OS systems running in the GDPS complex.

GDPS Installations

At the end of 2004, there were about 200 GDPS/PPRC systems installed, 68 percent in EMEA, 28 percent in North America and 4 percent in Asia/Pacific. In 2000, 90 percent of GDPS was deployed in EMEA. IBM has 57 GDPS/XRC, mainly in the United States and Asia/Pacific.

GDPS for z/OS requires PPRC level 3 and is supported by all major high-end disk subsystem vendors. Initially, almost all the disk subsystems taking part in GDPS/PPRC were Hitachi Data Systems. During the last few years, however, there are increasing numbers of IBM installations and several EMC installations as well. PPRC level 4 was licensed to StorageTek and EMC (expected support 1H05 — 0.8 probability). Hitachi carried out "reverse engineering" to implement it.

Key Issues

Who are the leading providers of business continuity planning products and services, and what are their strategies, key offerings and business practices?

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Level 7, 40 Miller Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Latin America Headquarters
Av. das Nações Unidas 12.551
9 andar—WTC
04578-903 São Paulo SP
BRAZIL
+55 11 3443 1509