THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2011030 August 24, 2011

IBM GDPS Actively Advances — Better To Be Up Than Down

Analyst: Stephen D. Bartlett

Management Summary

In the Information Age, everybody has to be concerned about preservation of important and essential data. We all want to be assured that the photos of our precious children or grandchildren or favorite pet will be there when we next look for them. We also want to be assured that the several years of accumulated financial data will be accessible whenever we need it and also that we will be able to find that warranty from that scurrilous vendor from whom we bought what now appears to be a lemon. Unfortunately, nine times out of ten, if you ask your average adult what they do to keep their precious data safe, more than likely the response is a blank stare. I willingly concede that there are lots of smart people out there who back up their data and system images at least semi-regularly, either locally or maybe to a remote location in the cloud. They pretty much keep these important parts of their lives in order. However, when these same smart people operate in the business world, what attitudes do they bring with them? Are they making a conscientious effort to make sure that their essential business data they create and use is protected from loss, and does their IT department (which may be small or large, but almost always overburdened) provide them with the means and the incentive to do so without any annoying obstacles?

Sooner or later, the average person or a business's IT department with poor or inadequate data recovery processes will rediscover *Murphy's Law* and its corollaries: *If anything can go wrong, it will; it will occur at the most inopportune time; and it's all your fault.* One day the average person's PC will refuse to start up normally, or the C-drive will crash, or some malicious software will creep into the innards and render it mostly useless. Even the most rudimentary IT department knows that similar disasters will occur, and the probability increases dramatically as the volume of data increases. The only question remains, accepting the inevitable: *What has been done to ensure the protection of business systems and all operational and historical data?* The average person will be frustrated and perhaps angry for a short while – usually directed at their poor dumb PC or its vendor. The businessperson, however, may find themselves unable to conduct business for a prolonged period of time or even worse.

What I am driving at here – and if you are reading this I am preaching to the proverbial choir – is that every business large or small must have a data protection, disaster recovery, and a business continuity plan. It must be appropriate for the risks, must be implementable and practical, and it must be affordable; in other words, one size does not fit all. It would be safe to say that IBM has some experience here and that its most comprehensive solutions involve *System z* technology. Building on over 12 years in this arena, on May 24th IBM took another step forward with the announcement of *Geographically Dispersed Parallel Sysplex (GDPS) Active/Active Continuous Availability*. Highlighted in IBM's most recent *zEnterprise System* announcement, the focus has switched for what has been primarily a *failover* model to a *continuous availability* model, and now includes IBM's roadmap for future Active/Active development. This new emphasis and the rationale that drives it are so significant that they warrant discussion and elaboration. To learn more, please read on.

IN THIS ISSUE

- Context and History 2
- Driving Toward the Perfect Solution 5
- Active/Standby — How It Can Work For You 6
- Conclusion 8

Context and History

This issue is about newly enhanced capabilities for Systems z, IBM's mainframe, using *GDPS (Geographically Distributed Parallel Sysplex*). It is about improving the quality of data protection at a greater level than was previously possible. If this is a problem that describes your operational battleground, read on to learn the details, even if you don't (yet) have a mainframe. You may wish that you did!

Understanding the Problem

Allow me a moment to relate a personal experience that awakened me to the consequences of an inadequate data protection and backup planning. It was the late 1960's and *System/360* mainframes were being installed at a high rate. In this small state government's IT operation in question, batch processing was the norm for its mostly-COBOL applications. The Department of Motor Vehicles regularly sent out renewal notices to its citizens based on their birthdates. On this particular cycle, the office was receiving an abnormally high number of calls inquiring as to when they should expect to receive their renewal, as they appeared to be late. "Soon" was the initial response by the DMV operators. At some point, it became obvious that something was amiss – just too many renewals were not reaching their citizens, and *then* the inevitable witch-hunt began.

Naturally, the DMV's programmer/analysts jumped into the problem very early. A pattern was established as to the missing renewal notices and the magnetic tapes that presumably contained the client records were reviewed. The records they were looking for were *missing*. Initially, fingers were pointed at IBM's tape drives and the possibility that they had, quite randomly, lost several hundred license-renewal records. As one of the supporting IBM systems engineers, this caused some very uneasy moments as we explained that this was highly unlikely, approaching zero probability. When the department's processes were reviewed, it was revealed that the tapes containing the records in question had been over-written in a prior cycle. A procedure had been put in place to save a few bucks on media by reusing the backup tapes every third of fourth cycle. Needless to say, the cost to the state to recreate the records far, far exceeded the cost of even a bushel of tape reels, not to speak of the embarrassment to the city officials, and the inconvenience to its citizens.

This little anecdote may not seem very germane in this modern era, but the *penny-wise, pound foolish* argument is still very much the issue when considering the value of operational and historical data and the costs of not having it readily available when it is needed. Moving forward to the present day when most business is being conducted online and in real time, the value of information currency is absolutely essential, almost to the preclusion of any other consideration. Therefore, for the remainder of this paper, I will assume that your business has a data backup and recovery plan and that you strive to stay as close as possible to continuous availability of your key processes, mission-critical applications, and the data that supports them.

Whether we are discussing disaster recovery or continuous availability, and whether we are considering planned or unplanned outages, the primary metrics for measuring the unavailability of a system, an essential process, or a key application are *Recovery Point Objective (RPO)* and *Recovery Time Objective (RTO)*. RPO is the quantity of data (number of transactions, records, storage units) you are willing to recreate following an outage. RTO is the duration of the outage, be it seconds, minutes, hours, or days. *Had an unhappy unplanned outage recently? How long did it last? How much did it cost? What could have been done to minimize it, or even avoid it entirely? Do you wish that your planned outages, say for maintenance activities, weren't taking so long and, therefore, are they being deferred or just not done?*

Every plan or process in place to ensure recovery and availability must be evaluated with respect to these metrics. The question that must be asked and answered is: *What level of RPO and RTO is appropriate for my business? Moreover, what level of RPO and RTO is necessary for each of the business's key processes?* Furthermore, it also is quite appropriate that each of the evaluated processes will have different RPO/RTO criteria and hence, different recovery and continuity solutions would be indicated. Also, consider the notion that *global enterprises* should not be construed to be limited to just mega-corporations. A global enterprise is *any* business that desires to operate 24x7, across time zones, or even across continents – this can apply equally well to even very small businesses and, therefore, affects how they should evaluate their recovery and continuity processes.

Exhibit 1 (at the top of the next page) describes the primary characteristics and RPO/RTO objectives that generally can be achieved from a broad range of solutions. Where are you in this hierarchy? This paper will focus on the seventh and eighth tiers; they define the highest and most rigorous tier for disaster recovery (DR) and continuous

Exhibit 1 — Levels of Disaster Recovery / Continuous Availability

Tier	Description	Typical Recovery Point Objective (RPO)	Typical Recovery Time Objective (RTO)
1	Point-in-Time (PiT)	Days since last backup	Days
2	PiT – Hot Site	Days since last backup	Days
3	Electronic Vaulting	Hours	Hours to Days
4	Active replication to Remote Site	Seconds to Minutes	Hours to Days
5	Active Storage replication to Hot Site	Near zero to Minutes	1 or more Hours
6	Disk Mirror at Hot Site	Near zero to Minutes (fewer than Level 5)	1 or more Hours
7	Disk Mirror at Dedicated Hot Site, Automated Take-over	Zero or Near Zero	Minutes to <2 Hours
8	Active / Active	<3 Seconds	<1 Minute

Source: IBM

availability (CA) solutions.

The Drivers

The primary factors that drive business's need for a robust disaster recovery and high availability solution would include the following.

- **Increased need for continuous availability** of key IT processes that serve customers and markets.
- **Risks of financial loss**, including lost revenues, punitive penalties, or legal actions resulting from business disruption.
- **Security-related incidents** that cause severe business impacts.
- **The effect on market reputation and brand image** by prolonged business disruptions.
- **Increasing regulatory requirements**.[1]

On the issue of financial losses and impacts, studies commissioned by IBM have estimated the business costs associated with IT service outages can range from $1M/hour in the consumer products industry to in excess of $8M/hour for financial services companies, and these costs are only rising. As this data clearly indicates, the costs of not having an appropriate disaster recovery and continuity plan can be staggering. The events of September 11, 2001, alerted us in real time what the risks of sudden catastrophic events were and how broad the effects could be. The very survival of a number

[1] For instance, see the *Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System*, FRB Docket No. R-1128, April 7, 2003, where the Federal Reserve prescribes recovery and continuity objectives for critical financial markets. Check it out at http://www.federalreserve.gov/boarddocs/press/bcreg/2003/20030408/attachment.pdf.

Exhibit 2 — A Parallel Sysplex Refresher

System z Parallel Sysplex is an innovative multisystem technology that facilitates data sharing, without compromising data integrity, provides multisystem clustering (up to 32 nodes are supported) with near linear scalability. It has been evolving for over 20 years and is the fundamental processing node for all GDPS solutions.

Every server in a Parallel Sysplex cluster has access to all the data resources so that all applications can run on any server, and is known as a *shared-everything* model. In addition, requests associated with a single workload, such as a business transaction or a database query, can be distributed dynamically for parallel execution on multiple nodes in the Sysplex cluster. In the event of a hardware or software failure, either planned or unplanned, workloads can be redirected dynamically to available System z servers thus providing near continuous application availability.

In summary, Parallel Sysplex is a means of managing multiple systems that can provide a user with:

- Continuous availability
- Very large total processing capacity
- Dynamic workload balancing
- Single system image
- Non-disruptive change or growth
- Application compatibility
- Ease of use

Exhibit 3 — The Demographics of GDPS (1H2011)

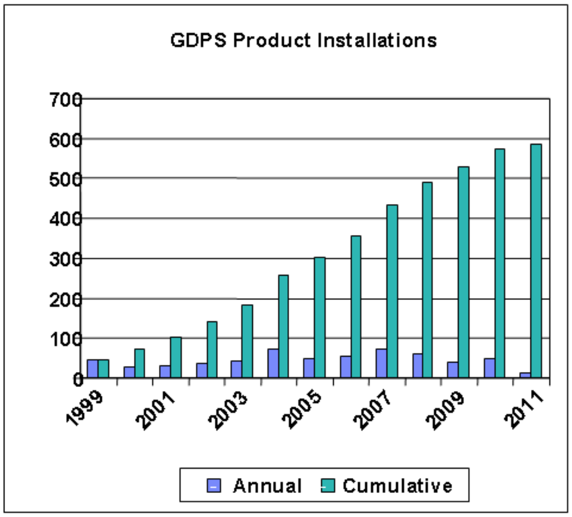

One or two site GDPS installations by product type

product	installs	percentage
RCMF/PPRC	52	8.8%
RCMF/XRC	15	2.5%
GDPS/PPRC HM	76	12.8%
GDPS/PPRC	302	50.9%
GDPS/XRC	97	16.4%
GDPS/GM	51	8.6%
Totals	593	100.0%

Three site GDPS installations by product type

product	installs
GDPS/MzGM *	36
GDPS/MGM **	27

GDPS solution by Industry sector

sector	installs	percentage
Communications	36	6.1%
Distribution	28	4.7%
Finance	432	72.8%
Industrial	30	5.1%
Public	52	8.8%
Internal IBM	10	1.7%
SMB	5	0.8%
Total	593	100.0%

GDPS solution by geography

major geo	installs	percentage
AG	155	26.1%
AP	70	11.8%
EMEA	368	62.1%
Totals	593	100.0%

* GDPS/MzGM consists of GDPS/PPRC HM or GDPS/PPRC and GDPS/XRC. 22 of 36 have PPRC in the same site. The licenses are counted in the prior table.
** GDPS/MGM consists of GDPS/PPRC HM or GDPS/PPRC and GDPS/GM. 11 of 28 have PPRC in the same site. The licenses are counted in the prior table.

Source: IBM

of financial, communications, transportation, and information technology systems, as well as a number of enterprises, were in question after that day.

The Solutions in Hand

This paper is about using Parallel Sysplex for maintaining geographically-distributed copies of data; thus, the GDPS acronym. For a quick refresher on Parallel Sysplex, see Exhibit 2, on the previous page.

Here is how IBM defines GDPS:

GDPS is a multi-site or single-site end to end application availability solution that provides the capability to manage remote copy configuration and storage subsystems (including IBM TotalStorage Enterprise Storage Server), to automate Parallel Sysplex operation tasks and perform failure recovery from a single point of control. GDPS helps automate recovery procedures for planned and unplanned outages to provide near-continuous availability and disaster recovery capability.[2]

As enterprises, large and small, become globalized, the necessity for maintaining a 24x7 presence on the web becomes more and more acute. As Exhibit 3, above, clearly illustrates, enablement of some form of GDPS solution has continued to grow at a steady pace across many industrial sectors. The financial sector, driven by consolidations and the regulatory environment, particularly in Europe, now leads all others by an order of magnitude.

Within IBM's System z range of GDPS initiatives, there are several solutions that have been implemented by their customers (shown in Exhibit 4, on the next page, and occupying the seventh tier, as referenced in Exhibit 1).

GDPS/PPRC HyperSwap Manager

RPO=0
RTO=0

This is a near-Continuous Availability (CA) solution for a single sysplexed site. It is based on the same technology as GDPS/PPRC but does not include much of the systems automation capability that makes GDPS/PPRC a more complete DR solution.

GDPS/PPRC (Peer to Peer Remote Copy)

RPO=0
RTO=<1 hour

This is a near CA-DR solution across two sysplexed sites separated by metropolitan distances, defined as a Fibre Channel (FC) connection not to exceed 200 km., approximately 124 miles. This solution encompasses both System z and open systems storage and is based on the IBM PPRC synchronous disk mirroring technology.

GDPS/GM (Global Mirror) and GDPS/XRC (Extended Remote Copy)

RPO=Seconds
RTO=<1 hour

These are DR solutions across two sysplexed sites separated by virtually unlimited distance. The *Global Mirror (GM)* solution encom-

[2] See http://www-03.ibm.com/systems/z/advantages/gdps/.

Exhibit 4 — GDPS Active/Standby Configuration (Before Failover)

- Static routing
- Automatic failover

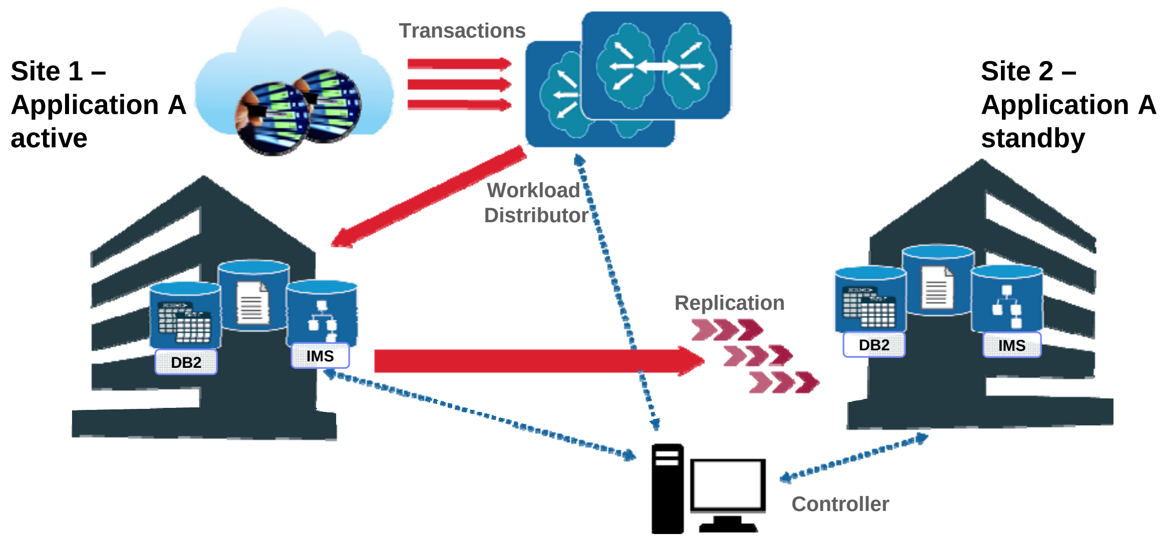

Source: IBM

passes both System z and open systems storage and is based on the IBM System Storage Global Mirror technology, a disk subsystem-based asynchronous form of remote copy, which is point-in-time (PiT) specific. The XRC solution encompasses System z (CKD) data only and is based on the IBM XRC asynchronous disk mirroring technology, and relies on cache and non-volatile storage to ensure consistency.

GDPS/MGM (Metro Global Mirror) and GDPS/MzGM (Metro z/OS Global Mirror)

RPO=Seconds
RTO=<1 hour

These are three-site solutions that provide CA across two sites within metropolitan distances and DR to a third site at virtually unlimited distances. *Metro Global Mirror (MGM)* is based on a cascading mirroring technology that combines PPRC and Global Mirror. *Metro z/OS Global Mirror (MzGM)* is based on a multi-target mirroring technology that combines PPRC and XRC.

Driving Toward the Perfect Solution

The primary distinguishing feature of the multi-site solutions is the distance between the sites that can be accommodated while meeting the business's RPO objectives. In order to achieve an RPO=0, the distance between data centers can be the limiting factor. It may work for a New York City to a metropolitan New Jersey or Philadelphia connection, but won't cut it for a second site in Boston, or Washington, D.C. Moreover, in many transaction processing and batch applications, signal latency will affect throughput significantly, thus limiting site separation to no more than 20 to 25 km. Thus, the starting point is GDPS/PPRC. Those who have implemented this solution in their respective enterprises[3] are the most intent on overcoming some of the limitations of that solution, driving toward an RPO and RTO of zero.

The *perfect solution*, if there is to be one, would have to embody all of the following characteristics:

- Unlimited distance between sites.
- Requires no changes to applications.
- Focus on continuous business operations, not just the host technology.
- Recovery process must be automated, not limited by operator skills.
- Provide dynamic workload distribution between sites.
- Provide application (workload) level granularity to recognize critical versus non-critical requirements.
- Provide for planned as well as unplanned recovery scenarios.
- Replace or totally avoid the need for RYO (Roll Your Own) applications.

At this point, the last requirement needs some elaboration, as it is an important principle that distinguishes IBM's System z GDPS solutions from most others. To wit, recovery and continuous availability solutions should be driven by standardized

[3] GDPS/PPRC represents more than half of the current GDPS installations (as of 1Q2011), as was shown in Exhibit 3.

processes and built into the host's hardware, system software, and middleware, or as IBMers' sometimes refer to it – *in the stack*. When implemented in this manner these facilities:

- Can be accessible to all,
- Likely to be more efficient and reliable,
- Of higher performance,
- Allow applications to be much less complex with their design focused on business processes and not computer processes, and
- Facilitate portability and platform independence.

The alternative would require inclusion of all the recovery and continuity logic completely within the application code along with the business logic, and for each and every application that warrants this discipline. For even the largest and most sophisticated installation, this becomes extremely challenging and costly.

Active/Standby — How It Can Work For You

IBM's solution to this was recently announced and is known as *GDPS Active/Active Continuous Availability*[4], the next evolution of the GDPS solution. Its objectives encompass application-level workloads[5], high availability, automated monitoring, automatic workload routing and recovery, and utilizes asynchronous replication between two sites. The objectives for RPO and RTO are less than three seconds and less than one minute, respectively (which was shown as Tier 8 in Exhibit 1 on page 3.) Multiple configurations of this solution are intended (see below) but the first configuration to be announced and supported is called *Active/Standby*, which became generally available on June 30, 2011.

A GDPS Active/Standby configuration includes the following elements.

1. **Sites 1 and 2 each defined as a Parallel Sysplex** with processors, storage, and replication software; all CA workloads (software, data, and network) are fully mapped to each, though the *standby* sysplex need only be in a ready state.
2. **Workload Distributor (WD) –** The WD is one or more front-end switches to the "cloud"; they are available from a number of vendors.[6] The WD communicates with the z/OS workload management functions in each Sysplex using the *SASP (Server/Application State Protocol)*[7]. The WD is programmed with the policy and rules for failover, and can be either automatically invoked or require operator (manual) intervention. Several would be located at each site or a neutral location to eliminate any single point of failure. GPDS supports up to 128 WDs.
3. **Controller** – The Controller is a z/OS server, preferably isolated from the Site 1 and 2 sysplexes. It includes a suite of management and monitoring software that monitors each site's status – a *heartbeat monitor*. The controller is used to start/stop sysplexes or individual workloads. It notifies the WD when new routing policy is set, either automatically or by manual intervention.

For the sake of illustration of an unplanned workload interruption, let's assume that the Sysplex at Site 1 in Chicago is active with "Application A" while the Site 2 Sysplex in Paris is in standby with that application, as shown in Exhibit 4 on the previous page. A transaction (unit-of-work) enters from the "cloud" to the Workload Distributor. The current policy rules send the transaction to Chicago for processing, a completion confirmation is returned, and then the altered data in the DB2 or IMS databases are replicated asynchronously to the matching data in Paris. The data latency will be about 2 seconds, or so, between these location distances. This continues as the steady state until the Controller detects an LPAR or server failure that "Application A" has been interrupted. Within seconds, a new routing policy is initiated. It may be an *automated* response or may be referred to the key operator for action, a *manual* response. Also, at this time additional processor capacity may be brought online at the standby site, if required, assuming provision has been made for standby capacity backup (CBU). In this case, the policy switches control of "Application A" from Chicago, Site 1, to Paris, Site 2, and all transactions for that application are now routed to this location as the

[4] From IBM U.S Announcement Letter 611-023, May 24, 2011: IBM GDPS active/active continuous availability is the next generation of GDPS and a fundamental paradigm shift for near continuous availability solutions.

[5] The workload within the scope of this solution is the aggregation of software: user-written applications and the middleware runtime environment (e.g., COBOL programs and CICS regions, respectively). The data is the related set of objects that must preserve transactional consistency and optionally referential integrity constraints (e.g. DB2 tables). The network connectivity is one or more TCP/IP addresses and ports (e.g., 01.10.10.1:80).

[6] Typical examples include Cisco *Catalyst 6500 Series*, Citrix *NetScaler Appliance*, F5 *Big IP Switch*, Radware *Alteon Switch*, and others, frequently referred to as *application delivery controllers (ADC)*. They perform services such as load balancing, caching, compression, and secured socket layer (SSL) functions.

[7] For more detailed information about SASP, see http://www.ietf.org/rfc/rfc4678.txt#.

Exhibit 5 — GDPS Active/Standby Configuration (After Failover)

- Static routing
- Automatic failover

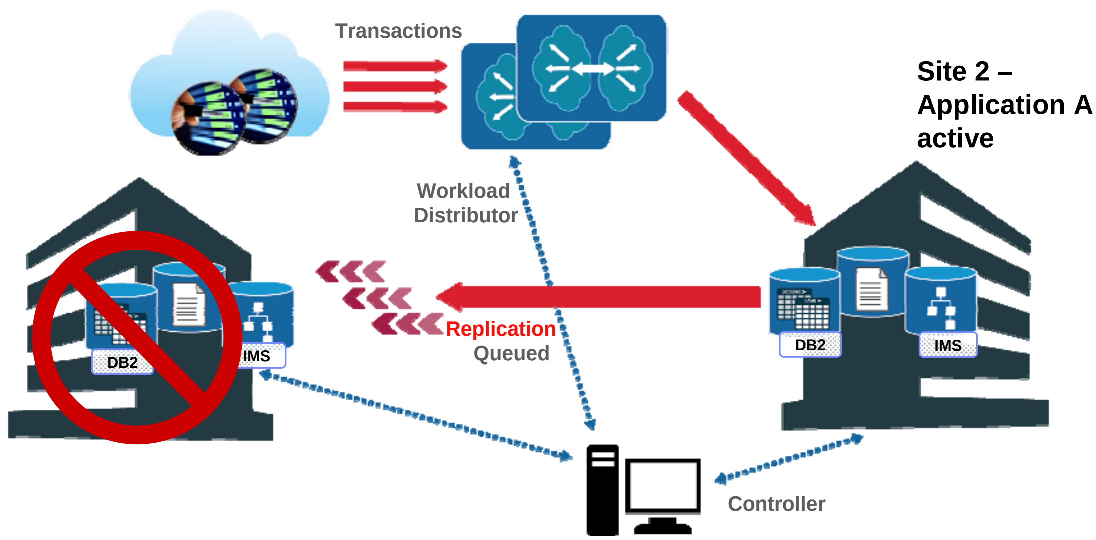

Source: IBM

primary site. (See Exhibit 5, above.)

Action must then be taken in Chicago to reinstate (fix and restart) the application environment in order for it to take its place as the new standby to the lead server for this application in Paris. Please note here that a complete switchover of all the applications (i.e., the whole site) in Chicago would not be necessary to institute the recovery. Chicago might very well remain the primary site for the remaining suite of applications, "B to Z", so to speak.

Single-workload switching, planned or unplanned would not be possible using traditional disk replication technology. The RPO would be unacceptable. When "Application A" is again available in Chicago, all of the data that changed while it was down is now replicated from Paris to return then to a steady state. The RPO in this example would be in the range of about 2 seconds, consistent with the replication latency.[8]

The scenario above would also apply in the event of a whole site failure as well – in this case, the entire application suite in the primary site would have to be switched, but the performance parameters would be essentially identical. For a planned switchover, whether it be for a workload or the entire site, RTOs could be in the tens of seconds with an RPO=0, no data losses, because the transition would be completely synchronous.

Early Testing Results – Does it Work?

Development is on-going, as one might expect, and the center of this activity is IBM's Montpellier Laboratory in France. The four contexts being tested are both workload and site switching, either planned or unplanned, with a simulated unlimited distance separating the two sites.

Currently, the Lab has achieved 20-second planned workload[9] switches and 120-second unplanned workload switches using automated switching policy and includes a 60-second failure detection interval. Both of these results would not be possible employing disk replication techniques. In the two-site switching scenarios[10], a planned site switch was achieved in 20 seconds and an unplanned site switch was achieved in 150 seconds using automated policies and a 60-second failure detection interval[11] and these times are expected to improve across the board as development continues. Going from hours to seconds with virtually no data loss seems an enviable achievement for this development effort.

Hardware and Software Support

All processors supported by *z/OS V1 Release 11* and above, which includes the *z9*, *z10*, and *zEnterprise* families, are supported by GDPS

[8] This may be considered a worst-case scenario, where data is not recoverable. If the data is just temporarily "stranded" it will be fully recovered and hence an RPO=0.

[9] A workload consisting of CICS applications fronting a DB2 database was used for this test.

[10] A combination of 9 CICS-DB2 plus 1 IMS workload was used for this test.

[11] Prior to this time, GDPS solutions that include disk replication would require 1 to 2 hours for a full site recovery.

Active/Active. On the production side, *InfoSphere IMS Replication for z/OS V10.1* was announced to support that database product, in addition to the existing replication product for *DB2*. For the Controller, *GDPS/Active-Active V1.1* and *IBM Multisite Workload Lifeline V1.1* were announced. Existing products *NetView for z/OS*, System Automation for z/OS, and IBM Tivoli Monitoring, also are required, and have been updated with new releases.

The Future of GDPS – and Some Speculation

As stated earlier, the announcement of GDPS Active/Active establishes a roadmap for future developments. The Active/Standby configuration is the first to be announced and delivered. At the same time, IBM also announced a Statement of Direction (SOD)[12] for the *Active/Query* configuration. This configuration option will provide the capability to distribute connections (each connection consisting of one or more read-only transactions) to both sites based on workload balancing policy with essentially the same objective criteria as the Active/Standby configuration.[13] Beyond that, GDPS architects are unwilling to discuss explicit developments.

With the recent second round of zEnterprise announcements[14], it begs the question how hybrid workloads will fit into the GDPS picture, because the focus of the solutions heretofore has been exclusively on z/OS workloads. This is an issue with which the GDPS development team is well aware and appears to have in sharp focus. The complexity of an *end-to-end* solution that includes all workloads and data both inside and outside the z/OS domain would seem to be enormous. Consequently, I would not expect to see this problem begin to be addressed for a couple of years.

More immediately, data outside the DB2 and IMS domains, in particular VSAM files, need to be included within the GDPS Active/Active architecture, as this data still constitutes a substantial portion of most System z customers' inventory. Currently, recovery and continuity of VSAM data is only within the scope of existing disk replication technologies.

[12] SODs are usually promulgated by IBM to inform its customers of a product or technology enhancement that is beyond an announcement horizon, typically about 1 year, and thus is not committed.

[13] Active/Standby and Active/Query configurations may coexist within the same two-site DR/CA complex.

[14] For more detailed information on the recent zEnterprise System announcement (z114), see **The Clipper Group Navigator** entitled *IBM zEnterprise in the Midmarket - Revolution or Evolution*, dated July 12, 2011, and available at http://www.clipper.com/research/TCG2011024.pdf.

Conclusion

GDPS Active/Active is IBM's newest and most innovative System z solution for customers that must have a disaster recovery and continuous availability solution that allows them the freedom to spread and protect their resources over significant distances. Frankly though, it is complex, requires experienced z/OS skills, and is expensive (which is a natural by-product of the necessary duplication of equipment and software at the remote site(s)). For all intents and purposes, every element, whether hardware, software, or network, must be redundant – any single-point-failure dooms the solution and surely Murphy's Law will be demonstrated quickly and tragically. These facts notwithstanding, nearly 600 installs (as shown in Exhibit 3 on page 4) confirm the conclusion that protecting enterprise data, business processes, and reputation with a GDPS solution is far less costly than the consequences of a prolonged outage.

Four features of the GDPS Active/Active solution are most compelling and innovative:

1. Unlimited distance between centers,
2. The granularity of the solution down to the workload level
3. The high degree of automation that may be enabled, and
4. The elimination of customized (RYO) recovery and continuity code in every application.

Each by itself would be something to crow about, but lack of ambitious goals has never been one of IBM's shortcomings.

Look for much more to come in the GDPS Active/Active solution for disaster recovery and continuous availability, an arena in which IBM's System z certainly seems to have the upper hand. The GDPS Active/Standby configuration is available now so that better protection and outcomes can be had by your enterprise, in ways that previously were impossible. Check it out!

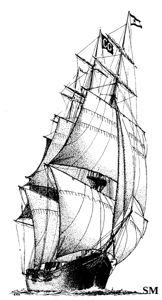

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Stephen D. (Steve) Bartlett **is a Senior Contributing Analyst for The Clipper Group.** Mr. Bartlett's interests include enterprise solutions including servers, system software, middle-ware, their underlying technologies, and their optimal deployment for responsiveness to emerging business requirements. In 2010, he joined the Clipper team after over 42 years with the IBM Corporation as an account and program manager in large system sales, product development, strategy, marketing, market research, and finance. During that time, he received several awards for his contributions in innovative market research and contributions to the business. He has a B.S. from Rensselaer Polytechnic Institute, and an M.S. from Union College.

- ***Reach Steve Bartlett via e-mail at steve.bartlett@clipper.com or at 845-452-4111.***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.