

Storage & Data Management: Three-Way Disaster Recovery Solutions for System z

by Ray Lucchesi
September 1, 2010

** Read this article online at <http://www.mainframezone.com/it-management/storage-data-management-three-way-disaster-recovery-solutions-for-system-z>

After the events of 9/11, governments everywhere began to reconsider their Disaster Recovery (DR) requirements for “critical” organizations. Prior to 9/11, most companies employed dual-site DR planning, where IT operations could continue when a single data center went down by transferring activity to another site located nearby. After 9/11, critical organizations were asked to guard against “region-” or 350-mile-wide disasters. Regionwide DR Solutions

To sustain operations in the event of such vast disasters with minimal data loss, a company would need three data centers—two located close to one another, with the third outside the region defined by the other two. In one scenario, the primary and secondary data centers within region sites synchronously replicate data between themselves, while the primary asynchronously replicates data to the remote or tertiary out-of-region data center. If disaster struck the primary site, activity could continue at the secondary data center, which would also take over asynchronous replication to the tertiary data center.

Another alternative is to use cascaded replication where the secondary site asynchronously replicates data to the tertiary site from the start. Since more often than not primary storage is down while processing elements remain operational, with cascaded replication, a company may only need storage at the secondary location. In this case, storage access swaps to the secondary site with minimal downtime. If the primary site fails, work moves to the tertiary site. Also, with a fully redundant, three-way DR solution, costs are significant, but having only storage at the secondary site can be more economical.

The IBM System z presents multiple asynchronous replication alternatives,



such as z/OS XRC, as well as proprietary storage subsystem-based solutions similar to those from EMC and HDS. Unlike proprietary vendor storage solutions, when using XRC, you can replicate to IBM, EMC or HDS storage but XRC consumes CPU processing resources, only supports CKD disk, and doesn't support cascaded replication.

Similarly, IBM, EMC and HDS also support their own proprietary synchronous replication solutions. In addition, EMC and HDS license IBM's proprietary synchronous replication facilities known as IBM Metro Mirror and, as such, can supply IBM-compatible services from EMC-to-EMC or HDS-to-HDS storage.

IBM's GDPS Three-Way DR Solutions

To facilitate the newly expanded government DR mandates, IBM released its three-way DR solutions with GDPS version 3.3 in 2006. IBM supports Metro Mirror and either z/OS XRC or its subsystem-based Global Mirror asynchronous replication capabilities in two three-site configurations called GDPS MzGM and GDPS MGM, respectively. With more than 500 two- and three-site installations (based on stats from IBM), GDPS also provides many features to ease business continuity and DR, specifically:

- **HyperSwap Manager** swaps primary site processing to use secondary site storage.
- **Consistency groups** join volumes and/or Logical Unit Numbers (LUNs) to conserve update sequences across sites and to initiate recovery for any volume or LUN failure.
- **Run-book automation** provides a script repository and automates script execution to restart operations.

EMC Three-Way DR solutions

EMC created GDDR, its replacement for GDPS, to take advantage of the company's proprietary SRDF replication services. In contrast, EMC can also operate in a fully GDPS MzGM-compatible mode with Metro Mirror synchronous and XRC asynchronous replication.

Also, GDDR provides AutoSwap and ConGroup features similar to IBM's GDPS HyperSwap and consistency groups. In addition, GDDR supports both cascaded and non-cascaded asynchronous replication to the tertiary data center.

Aside from SRDF, the other major benefit to using EMC's GDDR is its run-book expert system. GDDR's expert system makes defining and maintaining run-book scripts considerably easier than using GDPS.



HDS Three-Way DR Solutions

Like EMC, HDS supports both GDPS MzGM-compatible operations as well as its own proprietary replication solutions. However, HDS' proprietary solution only operates with GDPS, using the HDS Universal Replicator (HUR) and Business Continuity Manager (BCM). Using these facilities, HDS supports both cascaded and non-cascaded asynchronous replication to the tertiary site.

Summary

It sometimes takes a disaster to show the flaws in one's recovery plans. The events of 9/11, although tragic, have revealed that some catastrophes can impact multiple localities. As a result, IBM, EMC and HDS have all responded with three-way solutions that provide automated recovery for region-spanning disasters. Such capabilities can help organizations sustain operations whenever the next wide-ranging calamity occurs.

About the Author

Ray Lucchesi

Email: Ray@SilvertonConsulting.com

Website: <http://www.SilvertonConsulting.com>

Blog: <http://RayOnStorage.com>

Twitter: <http://twitter.com/RayLucchesi>

Ray Lucchesi is president of Silverton Consulting and has worked in data storage for more than 30 years. He helps both Fortune 500 and start-up storage vendors develop and market their products, and data centers better manage and optimize their data storage systems. Previously, he worked for StorageTek, where he focused on marketing and product strategy as well as Information Lifecycle Management (ILM) architecture and directed microcode development for disk subsystems.