IBM's Global Mirror for ESS hits the mark

Designing an IT infrastructure that ensures business continuity is no easy task. It involves designing, implementing, and testing failover procedures for servers, applications, and storage—processes that become increasingly complex as various flavors of hardware and software are added.

IBM has taken some of the complexity out of the process with TotalStorage Resiliency Family, an integrated series of hardware and software products to enhance business continuity. The family is a mix of new and existing IBM hardware and software. Its roadmap promises to automate key business-continuity processes, involving servers, storage, and everything in-between.

> "Up to eight ESS subsystems can be supported within a Global Mirror environment. Up to seven of these subsystems can be local with the remainder at remote locations. This compares favorably to many competitive offerings, which only support one local to one remote controller."

Business Continuity is...

Business continuity is not just about disaster recovery but about high availability. Disaster-recovery processes detail the necessary steps to recover from an unplanned event, such as a power outage, a flood or fire, or even a more common hardware or software failure. It is one subset of business continuity.

Proper business-continuity planning goes beyond disaster recovery, ensuring continuous, high-availability application processing and operation of servers, applications, and storage during planned outages, such as scheduled maintenance and backups.

An effective business-continuity plan involves several different hardware and software layers. The foundation of a solid business-continuity plan is fault-tolerant server and storage hardware. Data protection hardware and software (e.g., backup and restore) comprise the second layer, and the integration of second layer components with high-availability software for servers makes up the third layer. The last layer is composed of application and database automation products ensuring that applications are always kept up and running.

Within the first layer, IBM has several storage products, including ESS, FAStT, SAN switches and VTS, 3494, LTO, 3590, and 3592 tape products (see Figure 1). The second layer includes data protection and replication software such as Synchronous PPRC (now called Metro Mirror), Asynchronous PPRC (now Global Mirror) and FlashCopy point in time software. It also includes key recovery and replication management capabilities that are part of the IBM TotalStorage Open Software Family such as the IBM TotalStorage Productivity Center— including the Multiple Device Manager and Tivoli Storage Manager products that provide automated backup/restore and D/R. The third layer has z/OS, UNIX, and Windows automation software. The application layer includes automation software for SAP, Oracle, Siebel, SQL, and DB2.

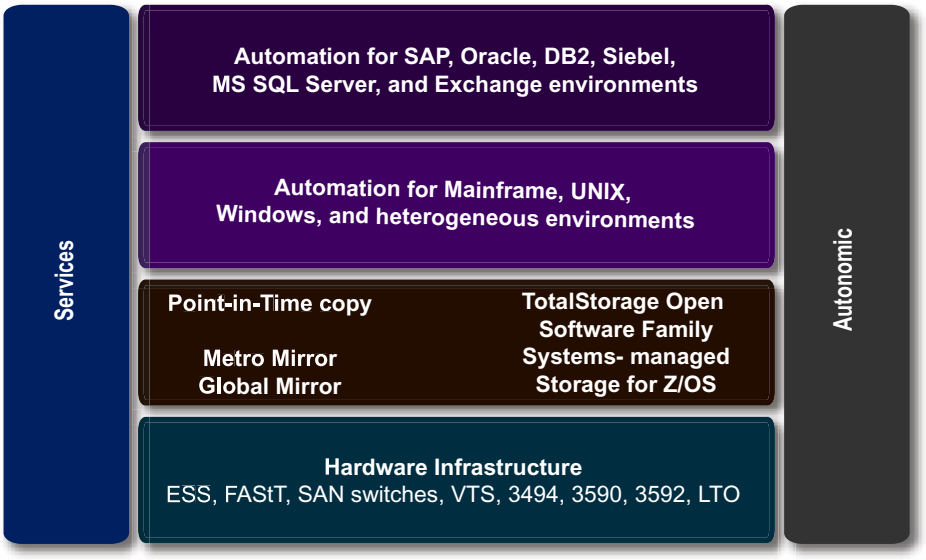

Figure 1: IBM TotalStorage Resiliency Family

IBM plans to tightly integrate these products across all layers and provide additional products and services that automate key business-continuity functions. A key component of this TotalStorage Resiliency family is Global Mirror.

Extending Mirrors Globally

IBM offers several products that replicate volumes over distances. For example, Peer-to-Peer Remote Copy (PPRC) allows users to replicate volumes from one IBM Enterprise Storage Server (ESS) disk array to a second like system.

For users looking to replicate data across short distances, IBM offers Synchronous PPRC. The software, which has been renamed Metro Mirror, writes updates to both the local and the remote controller before the operation is considered complete. Metro Mirror software is available for other storage products, such as IBM TotalStorage FAStT and SAN Volume Controller. In fact, all synchronous replication products within IBM will be renamed Metro Mirror.

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.881.9031 Fax: 603.881.9035

Similarly, IBM has renamed its Asynchronous PPRC software to better reflect its distance capabilities. Global Mirror, as it now called, is designed to replicate volumes over greater distances without affecting overall performance. The newly renamed software also features a variety of new capabilities, including automating consistency groups, which should be of particular interest to customers with large databases and other application volumes that span multiple ESS systems.

The Consistency Problem with Asynchronous Replication

Asynchronous replication, by definition, first updates the local volume and then at a later time updates the secondary volume. This means that the second volume can be several updates behind the primary volume. Should a primary volume fail, the secondary volume can be used for recovery purposes without the missing updates. However, recovery can become sticky if application volumes span multiple disk subsystems.

For example, suppose volume A on controller A, volume B on controller B, and volume C on controller C in the local data center are replicated to volumes A1, B1, and C1 on remote controllers A1, B1, and C1. Now, a power failure occurs in the local data center.

At the time of the outage, volume A1 is missing the last five updates from volume A , volume B is missing three updates, and volume C1 is missing two updates. The remote volumes are not consistent with each other—that is, they contain updates from different points in time.

In this scenario, database recovery may not be possible until all updates to all volumes are rolled forward or backward to a point in time where the secondary volumes are consistent with one another at a particular point in time. One way to get around this problem is to use IBM Global Mirror.

The Consistency "Fix"

Global Mirror addresses the consistency problem by creating a set of remote volumes every few seconds. Using the above example, one of the three local controllers would serve as the master controller, which would communicate with all the other controllers in the Global Mirror configuration without requiring additional software on the servers.

To ensure consistency, Global Mirror performs the following steps:

1. Volumes A, B and C are asynchronously replicated to volumes A1, B1, and C1 (see Figure 2).

2. It then suspends the asynchronous replication process, which transmits all remaining updates (up until the time the suspend command was issued) to the remote volumes. Volumes A1, B1, and C1 are now consistent images of the database at that specific point in time.

3. It then creates a copy (volumes A2, B2, and C2) of each of the three remote volumes, which will serve as the recovery volumes.

4. The replication process is then restarted and only the changes that have occurred since the process was suspended are transmitted.

5. Steps 2 through 4 are repeated.

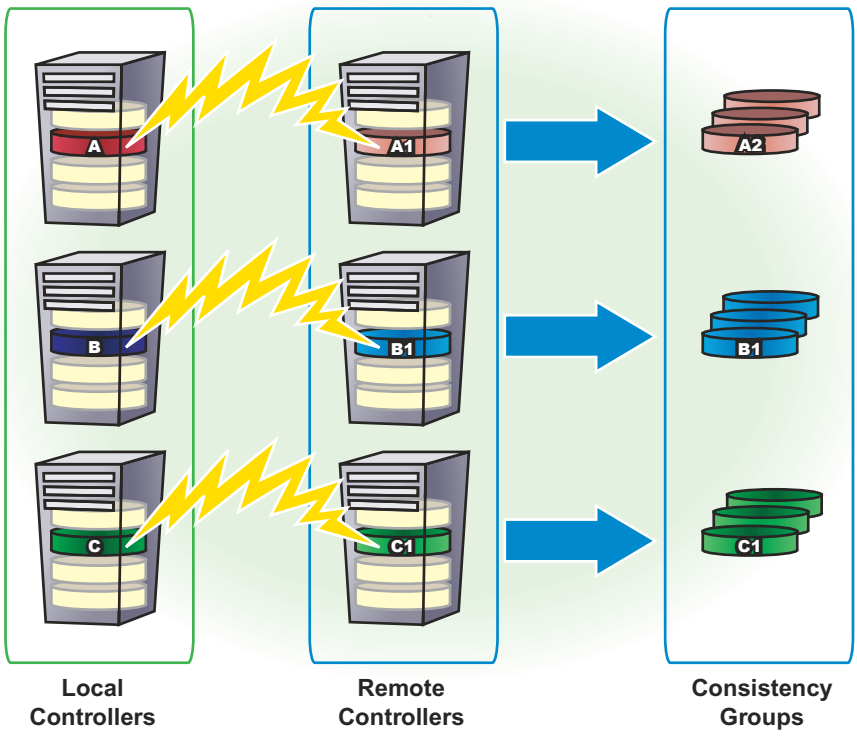

Figure 2: Creating Consistency Groups

IBM's design objective is to create consistency groups every three to five seconds. This objective can be met as long as sufficient transit time between remote sites and sufficient network bandwidth are available.

Asynchronous replication, by nature, does not guarantee that all updates are captured at the remote site. Creating consistency groups every three to five seconds rather than every 30 to 60 seconds, for example, increases the currency of the data that is captured in the consistency groups.

The Benefits of Global Mirrors

IBM Global Mirror includes a number of new features that will appeal to enterprise customers with large files and databases spanning multiple subsystems, including:

Automation of asynchronous replication process
Automation eliminates the need to write and test time consuming scripts.

Creation of consistency groups of volumes at the remote location
These volumes can be used to recover files and databases to a consistent point in time after an outage.

Support for u p to eight ESS subsystems
Up to eight ESS subsystems can be supported within a Global Mirror environment. Up to seven of these subsystems can be local with the remainder at remote locations. This compares favorably to many competitive offerings, which only support one local to one remote controller. IBM is expected to extend its support to more than eight subsystems by the end of the year.

Two Fibre Channel links per ESS
In most environments, Global Mirror only requires two Fibre Channel links per ESS subsystem, resulting in overall lower telecommunications costs compared to competitive products.

Support for z/OS and open systems volumes
Global Mirror supports not only open systems, but it also supports z/OS volumes. The consistency group can span z/OS and open systems, and can span multiple open systems.

It is free
A microcode upgrade is free to customers running PPRC version 2 on local and remote controllers and FlashCopy version 2 on remote controllers.

The one caveat is that it only works with ESS Models 750, 800, and 800T.

Conclusion

Synchronous replication may be less complicated but it also requires a lot more bandwidth than asynchronous approaches (since all updates to local volumes must be immediately written to remote volumes). And, synchronous replication may require installation of server software to maintain consistency groups. Finally, synchronous replication does impose distance limitations on remote locations.

So, for IT organizations that need to replicate over long distances, asynchronous replication is the only option. With the release of Global Mirror, IBM has removed much of the complexity of its former Aysnchronous PPRC offering. It has automated the process, provided support for numerous server platforms, including z/OS and various flavors of UNIX, and extended hardware support to up to eight ESS subsystems in a single configuration.

And the best part is—it is free for customers running PPRC and Flashcopy Versions 2!

Implementing Business Continuity processes is not easy—but it has now been simplified with the integrated solutions available within the IBM TotalStorage Resiliency Family.

76 Northeastern Blvd. Suite 29A Nashua, NH 03062 Phone: 603.881.9031 Fax: 603.881.9035