

WHITE PAPER

Virtual Tape Backup Replication in NonStop System Disaster Recovery – The Last Mile

Background

Events of the last few years have caused many companies to re-examine their disaster recovery plans. Simply having the computer system reliability provided by HP NonStop systems is not enough, when floods, hurricanes or terrorist attacks can destroy the data centers themselves or make them unusable for an extended period of time. For NonStop-based applications, disaster recovery plans include designation of a remotely located computer center with another NonStop system onto which the primary data center's load can be transferred in case of a catastrophic event. The possible topologies are:

1. An active remote computer system, perhaps used for development or test, is designated to take over from a failed primary system ("Active-Passive" sites);
2. Computer systems at both sites are continuously active, sharing the production transaction load ("Active-Active" sites.) The load from a failed data center is picked up by the other;
3. Multiple data centers can designate a common DR site that could take over for any one experiencing catastrophic failure ("Many-to-One" sites);
4. An inactive system at a DR data center is designated to be activated in the case of a primary site catastrophic failure ("Cold Standby" site.) Often this is in a facility operated by a company offering business continuity services, such as SunGard.

The Challenges

Whichever the approach, there are a number of tough challenges in the operational management of a DR solution, including:

- Copies of backups from the primary data center must be transported to the DR site, traditionally on magnetic tape;
- Delivery must be timely to ensure that backups at the remote site are up to date;
- The tape transport method must be reliable and secure, to ensure that tapes aren't lost, stolen or their data exposed;
- Tapes at the remote site must be managed and returned to the primary site for re-use when they expire;
- A catalog of tapes and their content must be maintained current at the remote site so that, upon primary site loss it can be determined what is available to restore;
- A mechanism for restoring catalog information onto the NonStop DR system (e.g. DSM/TC catalogs or the TMF catalog) must be implemented.

And of course there is the age-old worry: will any given magnetic tape actually be readable when needed. Gartner estimates that one in 10 recovery images on tape is unrecoverable. For this reason, some IT organizations write duplicate sets of tapes, perform a read-verify pass after writing tapes, or both. Such measures consume resources and increase backup time windows on the production systems, and still don't ultimately guarantee that all tapes will be readable when needed.



Security worries over losing tape media when moved off-site (based on recent high profile instances that received a lot of press) have led IT organizations to explore tape encryption, with cost and operational complexity implications.

All of this can be a concern even with Active-Active data centers, as there can be significant quantities of data local to each data center that would only be needed at the other after a catastrophic failure. So, backup tapes are exchanged between sites in many of these configurations as well.

Elimination of magnetic tape use in DR has been a dream of IT organizations. But until recently there were no solutions for NonStop host systems that were cost-effective, reliable and operationally simple. Communications-based remote backup data transfer solutions required significant resources, both in NonStop host processing and in bandwidth between the sites, not to mention software license fees.

NonStop DR solutions have usually imposed significant limitations in handling Guardian tape catalogs. To backup the DSM/TC catalog to tape that will be transferred to the DR site involves a complicated procedure of quiescing all tape activities and ensuring that Mediacom is inaccessible during the catalog backup. This can have significant operational impact, discouraging frequent DR tape catalog updates, and creating a window of un-cataloged backups for DR.

A Better Way

Now two technologies have come together to provide a cost-effective, efficient and easy to manage DR solution: ETI-NET's EZX/BackBox® virtual tape for NonStop systems, and Data Domain's data de-duplication products with replication capability. BackBox emulates tape drives SCSI-attached to NonStop S-Series systems and Fibre Channel-attached to NonStop Integrity systems. It is highly integrated with Guardian, storing its metadata and catalogs on the NonStop system and linking to DSM/TC and TMF tape catalogs.

While BackBox can store its virtual tape image files on a variety of back-end storage subsystems, its integration with Data Domain's "Restorers" enables it to take advantage of their data de-duplication, achieving effective compression of backup data by 20x or more. This leverages disk storage to permit many more backup iterations to be retained, or to greatly reduce the cost of the disk storage required for a given number of iterations.

Another benefit of data de-duplication is a corresponding reduction in the quantity of data needed to be transmitted between a local Data Domain unit and its corresponding unit at the DR site. Required WAN bandwidth can be reduced by 10x or more. This reduction also reduces the required throughput and cost of WAN encryption equipment, such as VPN servers, needed to ensure the security of the data being transferred between data centers.

But simply getting copies of backup virtual tapes to a DR site does not solve the problem of how to know what has been transferred and how to use it in restoring the primary system data onto the DR system. ETI-NET has filled in this "last mile" of NonStop DR operations with an innovative approach to catalog restoration.

On the primary system the DSM/TC and TMF catalog changes are captured and are automatically and transparently replicated to the remote Data Domain Restorer, along with backup data virtual tapes, where they can remain in that site's Data Domain Restorer until needed for DR restoration. When activation of the DR site configuration is initiated, BackBox's DR Domain volume catalog is restored, followed by automated updating of DSM/TC's and TMF's catalogs. Once completed, the DR system has all of the catalog information for backups and dumps that were present at the time of BackBox's last checkpoint on the primary NonStop system. Operations personnel can then quickly determine the latest backups and their content, and initiate their restoration.




Figure 1: Bi-Directional and Selective Remote Vaulting

How Does It Work?

EZX/BackBox can be configured with active and passive “Domains”. Active Domains could be designated for production, development or test use, for instance. Passive Domains are remotely populated with DR content, but are not activated for virtual tape access on that system until needed. Within a BackBox Domain, multiple “Data Stores” for storage of virtual tape images can be defined for different purposes, and can optionally be designated to be replicated. BackBox Data Stores that are mapped to Data Domain Restorers have their virtual tape content de-duplicated. This typically results in 1.5-3x effective compression of the first instance of a backup or TMF online dump, and 10-100x compression of subsequent instances.

Writing a NonStop backup or TMF dump onto virtual tapes in a given BackBox Data Store that is Data Domain-based and designated for replication triggers a copy of the de-duplicated differences in file content to be transmitted via WAN to the remote DR Data Domain Restorer. The WAN bandwidth used for replication can be throttled by time of day to limit its impact on other production traffic during peak hours.

Catalog Content Capture

Upon close of each virtual tape volume written by NonStop Backup, BackBox software in the NonStop system captures the TAPEVOLUME and TAPEFILE information that was updated in the DSM/TC catalogs. On a user-selectable schedule, BackBox software also extracts the DISKFILE information associated with these tapes from DSM/TC. A similar process is used to capture TMF catalog updates. The catalog information, along with BackBox’s DR Domain volume catalog, is written to a “special” virtual tape on the local Data Domain Restorer and is thus also replicated to the DR Restorer.

Upon a decision by operations personnel to activate DR, a restore of the “special” virtual tape populates the “passive” (DR) BackBox Domain catalog, and then applies the captured updates to the DSM/TC VOLCAT and FILECAT that are designated for DR use. This permits the DR site’s normal VOLCAT and FILECAT to be retained so that backups and restores for that system can continue as normal.



DR Topologies

For Active-Active, Active-Passive or Many-to-One site topologies replication can be bi-directional. So any system can serve as a DR system for the other. In the case of Many-to-one, multiple DR "Passive Domains" can be defined at the central DR site, with any of them able to be activated in the case of a site's catastrophic loss.

"Active-Standby" site topology is particularly interesting, as no live NonStop system is required at the DR site. A BackBox Virtual Tape Controller (VTC) and a Data Domain Restorer are used to accumulate the virtual tape images and catalog data. Upon loss of the active primary site, the VTC can be connected to an available NonStop host system and the NSK software uploaded from the VTC. After a few simple configuration steps, the catalog information can be restored onto the NSK host, and operations staff can use Mediacom, Restore or TMF to initiate restoring the appropriate backups from virtual tape.

Summary

Use of ETI-NET's EZX/BackBox virtual tape product on NonStop systems at the primary data center enables high-speed reliable backups and TMF dumps. The need for tape-verify passes and writing duplicate tapes is eliminated, and backup windows shortened. Storing the virtual tape images on Data Domain Restorers reduces their "footprint", using disk storage more efficiently by up to 20x or more, and reduces the WAN bandwidth required to replicate it to a DR site by 10x or more.

Now it is possible to entirely eliminate use of magnetic tape – the last holdout in many data centers has been DR. And cost-effective transmission of DR backup data via WAN connections enables use of common, extremely secure WAN communications security equipment to protect the content during transmission.

Elimination of physical tape results in:

- Drastic operations cost reductions;
- Elimination of tape recovery risk;
- No tape transport security issues;
- Reduction in backup windows.

But perhaps most important, BackBox's regular and reliable transmission of NonStop catalog information about backup content, along with its automated restoration as part of a DR takeover, removes the unknowns and minimizes the human error potential in the crisis atmosphere accompanying a catastrophic loss of a primary data center.

New NonStop DR site architectures can now be considered: Active-Active, Active-Passive, Many-to-One, or Cold Standby. And DR can now be implemented easily in a symmetrical configuration, where two sites back up each other.

Since physical tape is no longer needed for DR, the DR paradigm can be reconsidered and re-architected around business objectives and cost-effective use of resources.

Company Information

ETI-NET develops products that help customers manage the complexities of multi-vendor computer systems. ETI-NET products are designed to easily integrate storage resources from dissimilar computers and provide cost-effective consolidation and management of backup and archiving operations. With a product development center in Montreal, Canada and field operations centers in Boca Raton, Florida and San Mateo, California, ETI-NET has been shipping products for HP NonStop systems since 1987. ETI-NET supports customers worldwide and can be contacted by phone at 1-800-546-9101 or by email at information@etinet.com. To learn more about ETI-NET, visit www.etinet.com.