WHITE PAPER

Breaking the Terabyte/Hour Boundary - Ultra-speed TMF Dump/Restore on NS Integrity with ETI-NET's BackHome/Ultra™

As databases grow in size, the criticality of being able to restore them in a very short time increases, since business downtime is both expensive and may even endanger the life of the company itself. While NonStop systems are very effective in protecting against equipment failures, software problems or user errors can result in database corruption. For this reason NonStop TMF permits reconstruction of files and databases to any point in time before the corruption occurred. With today's databases of many terabytes, the speed by which this can be accomplished when required is critical.

The Problems

Recovery to a point in time with TMF requires restoring a TMF dump of the files or databases and a roll-forward from that point by applying changes made to the data since the time of the dump. This involves reading a number of audit trail logs and sequentially applying the changes reflected in these logs. However:

- The audit logs must be applied sequentially. Thus the longer it has been since the prior backup/dump, the longer it will take to bring the database to the required point in time.
- Hence more frequent TMF dumps are desirable from a recovery time point of view. But since it is not known in advance when database corruption will occur, this implies that multiple generations of dump copies must be kept.
- Applying the changes from audit logs can't begin until the data is fully restored from the backup/dump, so the speed with which the restore can run is critical.
- For today's large databases, given the cost and capacities of NonStop disk drives, it can be prohibitively expensive to dump to NonStop disk, particularly where multiple dump copies are required. Hence tape is commonly used.
- If the TMF dump was partitioned appropriately, parallelism of restore streams can be used, but this assumes availability of sufficient backup devices, such as tape drives. But NonStop tape drives are expensive and use of many drives presents challenges in tape management.
- While more frequent TMF dumps can significantly reduce the amount of audit transactions that need to be applied, and hence the time to recover the database, the system load caused during business hours must be considered.
- Tape speeds can be problematic. While older model tape drives have slow transfer rates, newer models such as LTO-3 drives can have a hard time dealing with host transfer rates that are bursty, or slower than their optimal media write rate. This results in poor effective transfer rates with NonStop systems, despite high tape device speed capabilities.

Background

ETI-NET offers both hardware and software virtual tape products for use with NonStop systems. EZX/BackBox, the hardware solution, attaches to NonStop S-series systems via SCSI and to NonStop Integrity systems via Fibre Channel. It permits storage of virtual tape media on internal disks, SAN-based disk, file servers, or forwarding to enterprise storage managers such as Tivoli Storage Manager, Veritas NetBackup or Legato NetWorker. ETI-NET's BackHome family of products act as software virtual tape, and forward objects via SNA to an IBM mainframe (BackHome), or via TCP/IP to a Tivoli Storage Manager server (BackHome/TSM and BackHome/Ultra.)

In the past, with S-series systems and earlier, use of TCP/IP-based solutions was prohibitive where high transfer rates were required, due to the CPU overhead involved. Hence the evolution of SCSI, and subsequently Fibre Channel-based hardware solutions. These solutions permitted a great degree of parallelism of operation, emulating multiple tape drives.

Now, in a sense, the wheel has come full circle with the advent of NonStop Integrity systems. Their powerful Itanium processors, along with an efficient hardware and software implementation of TCP/IP over Gigabit Ethernet has enabled very high transfer rates. At the same time, the legacy code of the tape I/O Process (IOP) has become a bottleneck for transfers via Fibre Channel. Whereas FC is capable of up to 200 MB/sec, and LTO-3 tape drives have rated speeds of up to 80 MB/sec native (150 MB/sec with compression), actual transfer rates that can be achieved per drive via the NonStop tape IOP are typically in the range of 30 MB/sec.

A New Lease on Life for BackHome

While working on a special high-performance backup project for HP earlier this year, ETI-NET realized that the combination of Itanium CPU power and efficient gigabit TCP/IP in the NonStop Integrity series could be harnessed to provide ultra-high speed transfers of TMF backup data. The key was to compress the data before sending it out of the system via TCP/IP, thereby effectively increasing the effective throughput of each G4SA port, and reducing the CPU overhead of the TCP/IP processing in the NonStop and destination systems.

BackHome/TSM operates by intercepting transfers between TMF and the tape IOP, re-packaging and re-routing the data for transmission via TCP/IP and TSM's API to a Tivoli Storage Manager server, where it is stored in disk or tape pools according to designated policies. For the BackHome/Ultra product this processing was modified to compress the data immediately after interception from TMF, as well as to implement a number of other performance and management modifications.

Benchmark Configuration

The BackHome/Ultra product was benchmarked on a 16-processor Itanium-based system, with a 641 GB TMF-protected NonStop SQL database, partitioned such that 16 parallel TMF Dump jobs could be run. The NonStop system was connected with 2 GbE links via a switch to a Proliant DL385G4 server running Tivoli Storage Manager v5.3.2 client and server software. In turn, the Proliant used 2 HP StorageWorks MSA1000 disk arrays as JBODs, with 8 disk drives actively used in each as a TSM disk pool. The test configuration is shown in figure 1.

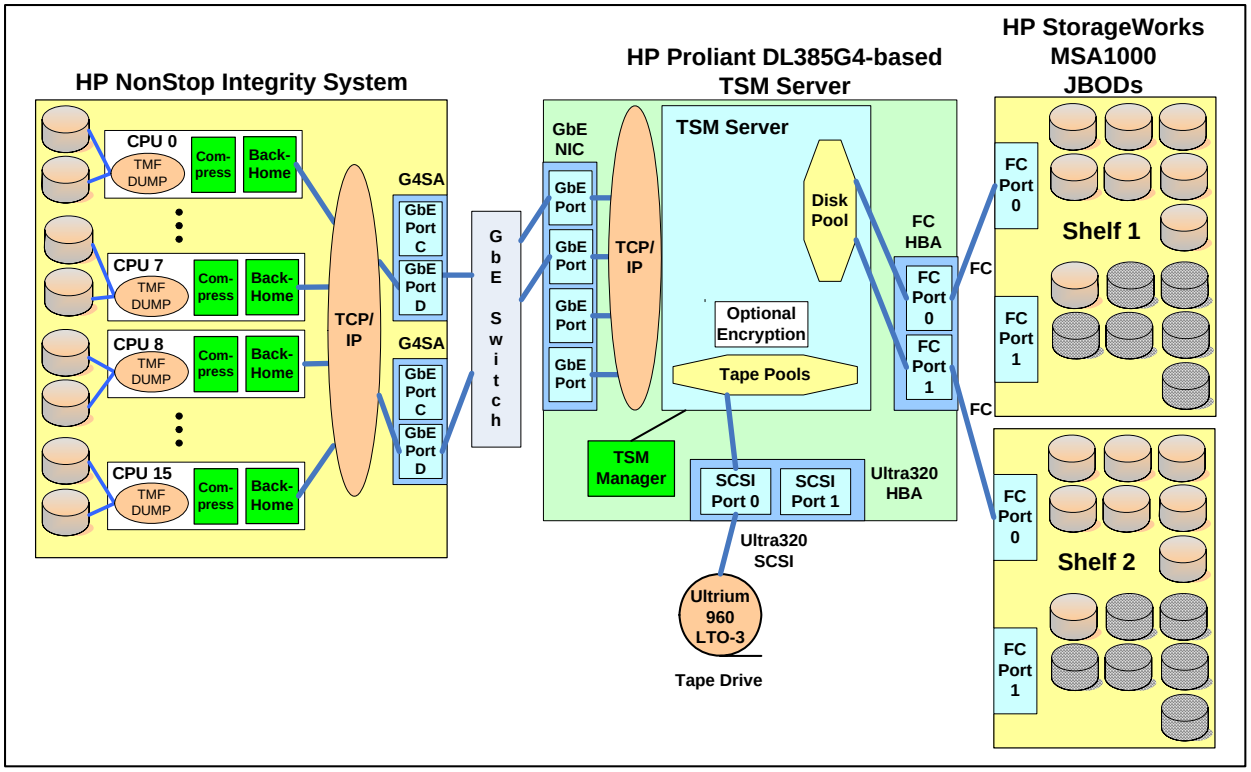

Figure 1: BackHome/Ultra Benchmark Configuration

TMF dump data from the 16 parallel streams was initially stored in the JBOD-based disk pool, and subsequently migrated from disk to TSM tape pool via a single LTO-3 tape drive.

Two different restores of the NonStop SQL database were performed – the first from the backup copies in the TSM disk pool, the second from the tape copy via staging in the TSM disk pool.

Benchmark Results

TMF dump to TSM disk pool

Time for TMF dump of 641 GB database:	~28 minutes
Instantaneous transfer rate to the TSM server:	441 MB/sec
Overall rate (*table sizes were not identical*):	390 MB/sec
	1.4 TB/hour

TMF restore from TSM disk pool

Time for TMF restore of 641 GB database:	~35 minutes
Instantaneous transfer rate to the TSM server:	352 MB/sec
Overall rate:	312 MB/sec
	1.1 TB/hour

TMF restore from single TSM tape drive

Time for TMF restore of 641 GB database:	~75 minutes
Overall rate:	146 MB/sec
	0.5 TB/hour

(*Note: rates are quoted in equivalent uncompressed source data bytes*)

Note that the speed of restoring from tape was limited by use of a single tape drive attached to the TSM server. Addition of a second tape drive is relatively inexpensive, as open system commodity tape drives can be use, and would almost double the end-to-end restore throughput, resulting in a throughput only slightly lower than that from the TSM disk pool.

NonStop System Overhead

BackHome/Ultra is intentionally designed to utilize available CPU cycles on the NonStop Integrity processors to maximize its compression and throughput. For example, in the above benchmark the CPU load was over 90% in all processors during TMF dumps. It is common for organizations to schedule backup windows during off-hours, so it is not unreasonable to run backups with high CPU loads and disk I/O rates during these periods. However, a major objective of the product was to enable more frequent TMF dumps to minimize the audit restore time required should TMF roll-forward ever be required.

To be a "team player" and co-exist with production database transaction loads, a "throttling" mechanism has been designed into BackHome/Ultra. The user can "dial-in" the rate at which it should run to limit interference with production traffic. And the process priority of BackHome/Ultra can be set to be lower than production applications, also limiting CPU load conflicts.

Summary

BackHome/Ultra is the perfect solution for ultra-speed backup and restore of large database on NonStop Integrity systems:

- It is the only backup product for NonStop that can achieve higher than 1 TB/hour.
- It can be run as frequently as needed to minimize recovery time, including TMF roll-forward.
- It uses NSK gigabit Ethernet ports (G4SA) very efficiently, achieving >80% utilization.
- It can be run full-speed during off-hours, or at a "co-operative" rate during production hours.
- The hardware platforms required for TSM are relatively inexpensive open-system products.
- Disk pools for storage of backup images can be implemented on different disk storage platforms, offering a range of availability, integrity, performance and cost tradeoffs.
- Restore speed can be increased by adding commodity tape drives.
- Tivoli Storage Manager, along with ETI-NET's TSM Manager product, is a powerful tool for automated management of backups and media, supporting tape libraries and drives.

While BackHome/Ultra is perfect for many NS Integrity database customers running TMF, other ETI-NET products, such as EZX/BackBox, may be better suited for other applications.

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.