Implementing Tape Storage on Apple Systems

Apple's Xserve and Xserve RAID comprise an exceptional platform to support new applications within the traditional Apple customer base and those currently running on non-Apple systems.

Xserve's increasing encroachment into "corporate-centric" computing is driving the need to implement the use of tape; a widely used technology that stores and protects critical information should the primary system, or system storage fail. This information can be in the form of traditional text-based content or art in the form of very large audio or graphics files - and the loss of either can severely impact the well being of an organization.

This paper addresses the use of tape devices in Mac OS X environments.

Managing Tape Technology Under Mac OS X

Mac OS X itself contains no code to manipulate tape devices. TOLIS Group has developed tapectl™ and libctl™ tools that seamlessly manage the operation of tape devices and tape libraries under Mac OS X:

- tapectl - manages tape drives (rewind, erase, append, setblk, etc)
- libctl - manages tape libraries (load, unload, inventory, move tapes, etc.)

These tools are integrated into each TOLIS Mac OS X backup solution and supports all tape formats, including LTO, AIT and VXA as well as most other tape technologies (all generations).

Tape Device Performance

The key performance characteristics of tape devices are throughput and storage capacity - the major selling points touted by manufacturers. The throughput (speed) of tape technology is measured in MB/sec (megabytes per second), and capacity is measured in GB (gigabytes), reflecting the amount of information that can be stored on each discrete tape cartridge.

The throughput and tape capacity performance realized when using a particular device is directly affected by "data compression." Virtually every modern tape device today provides data compression at the device level. When the data is compressible, the end result is a reduction in the overall physical size of the file. The size reduction of a file is directly related to its compressibility. When the data is highly compressible, a greater reduction in size is achieved.

For example, a 100 MB file of compressible data can be reduced in size significantly. If a compression ratio of 2:1 is achieved, the 100 MB file will be reduced in size to 50 MB. If only a portion of the 100 MB file contains compressible data, the resulting compressed file may only

become 85 MB in size, for example. Not all data is compressible, and some file formats even automatically compress the data (MPEG 3&4, TIFF & JPEG graphics) and therefore cannot be compressed further.

Not all tape device manufacturers base their performance specifications on the same criteria. Some tape device manufacturers will quote "native" performance figures without data compression (ratio of 1:1) while others quote performance based on a specific compression ratio (i.e. 2:1 or 2.6:1). Some manufacturers will use both native rates and data compressed rates to document the performance of their different tape products.

When viewing product documentation, one product can appear to be significantly faster than another, yet each may perform the same in a real world application. The "fine print" must be considered when evaluating the performance capabilities of tape devices.

BRU solutions offer software compression, but data compression at the software level should only be used when hardware data compression is not supported by the tape drive. Employing both hardware and software data compression in parallel does not increase data compression efficiency, and can, in fact, result in a decrease in throughput performance.

Tape Drive Throughput

Quite often a discrepancy arises between a user's perception and reality regarding tape drive performance. For example, a manufacturer lists the sustained throughput performance of their drive at 60 MB/sec. When the specification is based on compressed data, it will usually be asterisked or footnoted in the product documentation to relate; "assumes 2:1 compression." The number of 60 MB/sec is "burned" into the user's consciousness. That rate will be achieved given the correct conditions - the ability to write data that can be compressed at a 2:1 ratio. That same device backing up non-compressible data at a sustained rate of 30 MB/sec will be operating at full-bore efficiency.

In addition to the data compression concern, the optimal block transfer size of each discrete tape drive model varies. Some drives work best with 64K transfers, while some deliver optimal performance at block sizes of 32K, or some other size. Counter intuitively, larger block sizes are do necessarily result in improved throughput. The architectural foundation of each particular drive model dictates the transfer size "sweet spot." BRU solutions for Mac OS X systems provide the ability to set the I/O buffer size, allowing the user to set the optimal block transfer size.

In the past, the network and tape technology were the constraining factors to backup system throughput. Advances in tape drive technology have improved performance to the point that inefficient backup software can now actually be the gate to backup system efficiency. The BRU technology used in the Mac OS X backup products can stream (fill the voracious appetite of) the most advanced tape device.

Tape Cartridge Capacity

Data compression not only affects the throughput capability of tape drives, but also how much information can be contained on a tape. For instance, a Sony AIT-3 tape cartridge rated at 260 GB (2.6:1 compression ratio) will only be able to contain a native capacity of 100 GB of non-compressible data.

Understanding the compressibility characteristics of you data will allow you to configure a backup system that will deliver the performance and capacity necessary to meet your business requirements.

For best results, always start your backup design plans assuming NO compression.

Tape Performance

On an Xserve/Xserve RAID system with the storage subsystem configured to run at 2 gigabit, BRU is able to pull data off the RAID subsystem at 123 MB/sec. At a rate of 123 MB/Sec, BRU can not only stream the most modern tape technology available today, but also provides investment protection because it will feed the forthcoming generations of tape technology.

Customer Experience #1

Industry:
Entertainment/Production
Data Mix:
Mostly non-compressible
Hardware:
Xserve, Xserve RAID, SCSI LTO-2 tape drive
Software:
OS X Server 10.2, BRU

In a production environment, BRU backed up 122.5GB of data, comprised of nearly 75% non-compressible data (video/music), in 69 minutes. BRU sustained a single stream average throughput in excess of 29.5MB/sec to the LTO-2 tape drive. This performance was achieved out of the box without any system tuning. The customer was so impressed that he posted the results on the Web at http://forums.osxfaq.com/viewtopic.php?p=24735#24735

Customer Experience #2

Industry:
Children's Apparel
Data Mix:
Non-compressible
Hardware:
Xserve, Xserve RAID, Spectralogic LTO-2 tape library
Software:
OS X Server 10.2, BRU

In this production environment, the customer is writing a sustained, single stream of non-compressible data at 29.6 MB/sec through the LTO-2 SCSI drive in the Spectralogic tape library, resulting in the backing-up of 102 GB/hr. of data.