*i*ntelligent Bridging Architecture™

Speed Write

White Paper

Intelligent Bridging Architecture™

Speed Write

Storage Area Networks (SANs) are quickly becoming the method of choice for consolidating enterprise storage. As SAN deployment grows, so does the need to be able to back up these large pools of data. As the amount of data increases, so does the need to include high performance, scalable tape-based storage systems in the SAN.

Traditional methods for backing up data on a server rely upon a separate tape drive on each server. This typically requires manual efforts by the IT Department to kick off backups, change out tapes and manage the library of archived information. Once the backup process gets too large to affordably manage, one could consolidate by bringing the data over the network to a dedicated backup server connected to an automated tape library. Due to the amount of traffic across the LAN, this architecture can dictate that administrators limit backups to off hours. With the amount of data increasing at a rate of 100% a year, the off hours backup window is continually shrinking. It is at this point that IT departments need to consider more modern solutions.

Many companies are turning to SANs as an answer to their backup problems. The Storage Area Network, (see Figure 1) most often implemented using Fibre Channel technology, allows multiple servers direct connections to a shared pool of storage. Tape backup systems can easily be integrated into this shared storage pool.

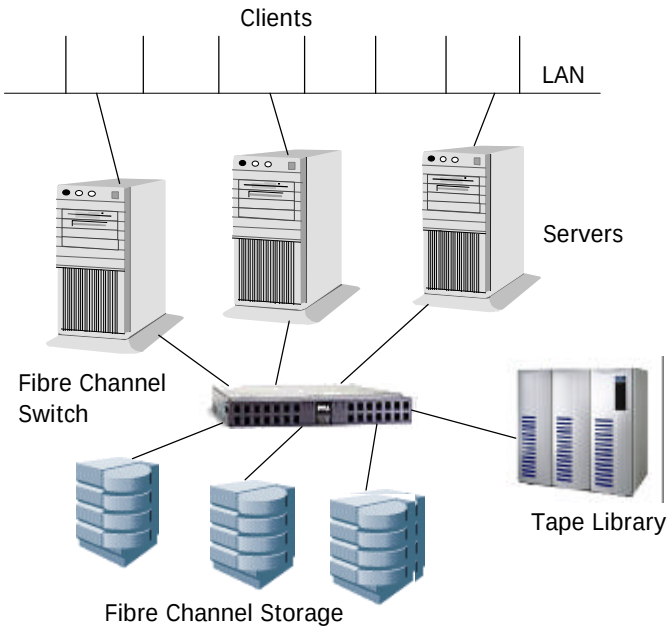

Figure 1: Fibre Channel SAN

The SAN provides for LAN free backup operations as the data moves from the storage pool into the backup server and back out to the Tape Library. This eliminates any backup traffic from moving across the LAN and allows one server to manage the backup process for all of the disk storage.

When it comes to implementing a tape backup solution into a SAN, there are two basic ways to go:

(1) Native Fibre Channel tape drives
(2) SCSI based tape drives with a Fibre Channel to SCSI bridge in front of them

Using native tape drives eliminates the need for a Fibre Channel to SCSI bridge or converter device, which adds complexity, but brings out other challenges. These issues are primarily a function of the Fibre Channel Arbitrated Loop protocol itself.

The amount of data to be backed up on a SAN most often requires a tape library, consisting of a robot controller and multiple tape drives. If native Fibre Channel devices are used, each device requires a direct connection into the Fibre Channel SAN. With switch ports costing between $1000 - $2000, an eight-drive library can prove to be a costly addition. Because of this, most manufacturers and integrators recommend the use of an arbitrated loop hub to connect the tape and robot devices into the SAN. Each drive connects to a port on the hub, with one hub port connected to the switch. Hub ports are typically one-fifth the cost of switch ports, making the solution much more affordable.

The trade off of using a hub forces you to deal with an issue in Arbitrated Loop called a LIP (Loop Initialization Primitive). A LIP can be simply explained as the process a Fibre Channel port must go through to initialize, determine a loop master, obtain an AL_PA (loop ID) and basically introduce itself to all of the other nodes on the loop. LIPs typically occur when devices on the loop are powered on or off, or if a device loses sync with the rest of the loop. If a LIP occurs while tape I/O is occurring, traffic on the loop could be delayed long enough to cause the backup to fail or abort, requiring the backup job to be manually restarted. Unfortunately, tape simply is not very tolerant of any kind of interruption.

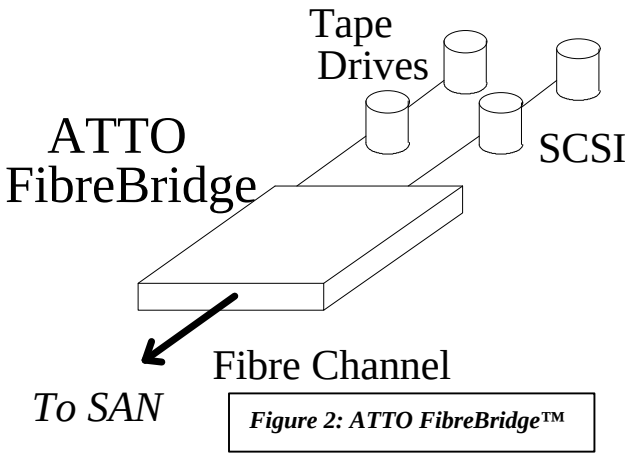

Figure 2: ATTO FibreBridge™

The alternative to native Fibre Channel tape libraries is to use SCSI tape drives and robot controllers with a Fibre Channel to SCSI bridge in front of it (see Figure 2). As Storage Area Networks emerge as the leading method to transition from server-based to storage-based architectures, companies need to focus on a total systems approach to designing and managing the storage. Adding intelligence to the storage is the key.

The ATTO FibreBridge™ not only converts the Fibre Channel protocol into SCSI, but also adds value to the SAN because it is an intelligent storage management device. The technology applied to move data efficiently and safely is a collection of knowledge ATTO Technology has developed over the last fourteen years. Some of which will be discussed below.

There have been many advances in the last year to improve upon the backup process and the time it takes to complete. One of the more profound improvements is the increase in throughput that the newer drives can sustain. The emergence of SDLT and LTO tape

drives that are capable of moving data at rates near 40 MBytes/sec. today (80 MBytes/sec. in the near future), require that the SAN infrastructure be able to deliver the data accordingly.

This is not an issue when native Fibre Channel tape drives connected directly to a switch are used because the Fibre Channel pipeline can easily support the bandwidth. As mentioned above, this can be a very costly implementation. Using a hub poses an additional problem beyond the LIP issue. It has limited bandwidth. With a 1-Gig Fibre Channel hub, the available bandwidth to be shared among all connected devices is 100 MBytes/sec. With tape drives capable of 40 MBytes/sec., only two tape drives can be used without creating a bottleneck. Standard Fibre Channel to SCSI bridges or routers can also pose a problem because they use a "store and forward" design. This means that data comes into the bridge, is temporarily buffered, and then forwarded. With a sequential device like tape, only one command can be processed at a time.

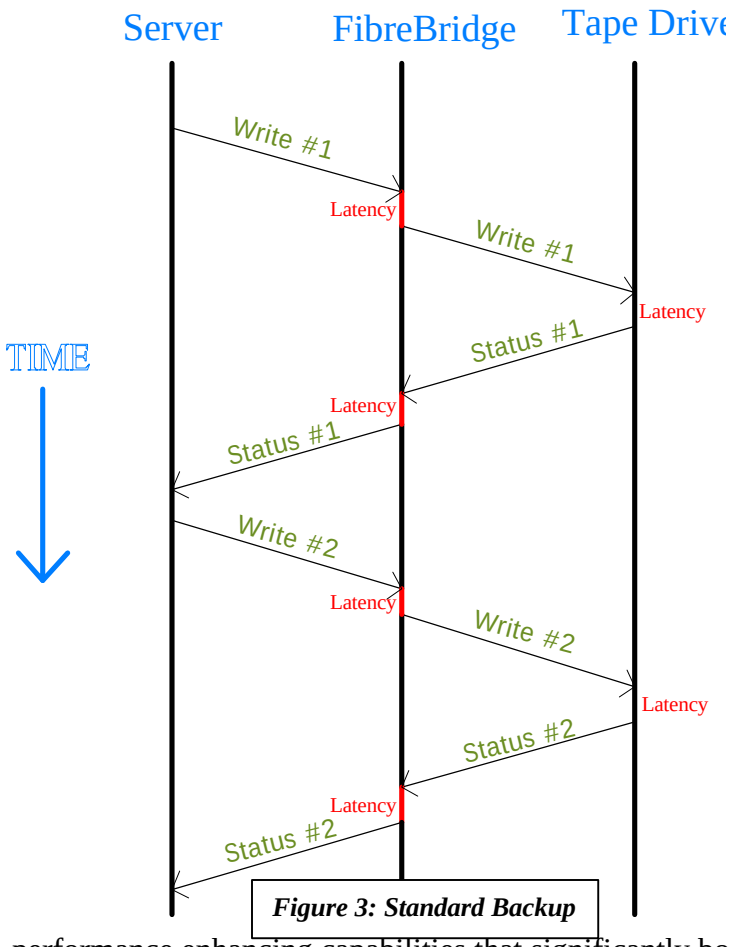

Figure 3: Standard Backup

Figure 3 shows the flow of data during a backup as it passes from the server through a bridge, to a tape drive. A Write command must complete and the drive must respond with status before the next Write can begin. There is latency associated with the bridge as the command is processed as well as a time delay to actually write the data to the medium. The inherent sequential nature of tape drives together with a bridge reduces the overall bandwidth available to the backup process.

One of the proprietary intelligent Bridging Architecture™ features that ATTO Technology has added to the line of FibreBridge products is Speed Write. This feature is the latest in a series of performance enhancing capabilities that significantly boosts the write performance of the FibreBridge by efficiently managing the Write commands between the host and tape drives. IT Managers, System Integrators, OEM's etc. will achieve faster backups when using ATTO's Speed Write with any of the industry-leading backup applications.

Speed Write is a mode of operation the FibreBridge runs in where SCSI Write commands are processed using 'Write-Behind' and 'Deferred Error' handling, as described in the SCSI specification, to return completion status back to the host prior to actual command completion. The result is that a host can issue new Write commands to a FibreBridge SCSI target in faster succession than is possible with a standard architecture.

Referring to Figure 4, when Speed Write is enabled, the server will send a Write command to the FibreBridge and the bridge will immediately respond with status. At the same time, the bridge actually writes the data to the tape drive. Using 'Write Behind' the bridge is ready to accept the next sequential Write command from the server much sooner than a standard bridge could handle. The overall result is that the throughput between the host and the tape drive is significantly increased.

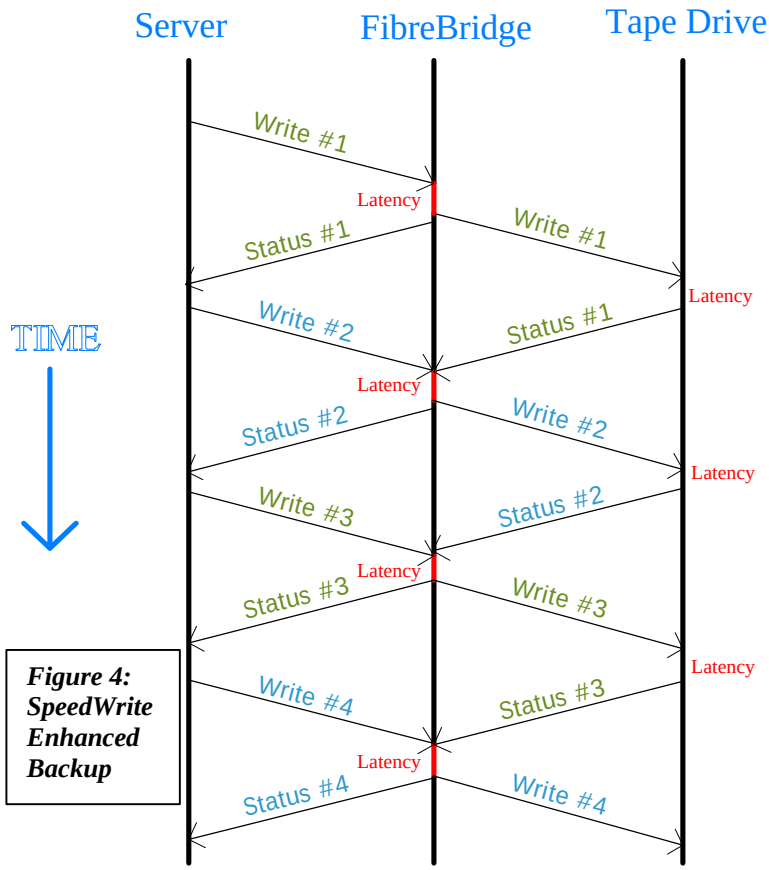

Figure 4: SpeedWrite Enhanced Backup

Using Write Behind requires added intelligence in the FibreBridge to handle errors. If an error were to occur during Write #2, the FibreBridge would respond with a Deferred Error Check Condition in the status frame for Write command #3 or #4, depending on the timing involved. The host would then be required to rewind and re-send the appropriate data to correct the issue. This is known as deferred error handling. While most backup applications support this function today, it is a requirement for the Speed Write feature to work properly in the FibreBridge.

Speed Write operates most effectively under the assumption that a large stream of successive Write commands is issued from a FC host to a SCSI target connected to the FibreBridge. The idea is to keep the path between the host and the bridge filled with data. Every time there is a pause, additional latency is introduced resulting in slower performance.

The Speed Write performance enhancement operates transparent to the user. When the feature is enabled for a device, the FibreBridge maintains a write performance equal to or greater than its current write performance. The amount of the increase depends upon the particular Fibre Channel host adapter, the type of tape drive and compression setting, the content of the data sent from the host, and how much additional traffic is flowing out of the backup server or through the FibreBridge.

Figure 5 shows the results of two tests run to measure the ability of the FibreBridge to sustain the high throughput rates achievable by high performance tape drives. An *ATTO ExpressPCI UL3S* Ultra 3 SCSI host adapter was used to establish the baseline (expected)

performance levels of the tape drive using highly compressible data. Tests were then run with the FibreBridge *Speed Write* (SW) function disabled and then enabled on a Quantum SDLT and an IBM Ultrium LTO tape drive.

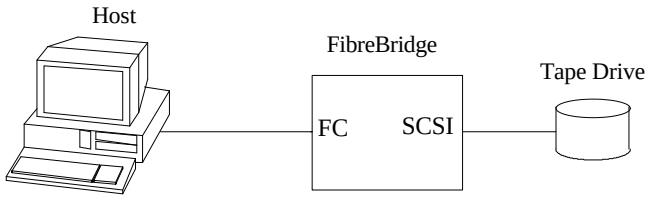

Figure 5

→ Quantum SDLT Tape Drive Test

Transfer Size	*Throughput (Mbytes/sec)*		
	SCSI Baseline	***FibreBridge Speed Write OFF***	***FibreBridge Speed Write ON***
	Writes	**Writes**	**Writes**
64 K	62.5	31.3	62.5
128 K	62.5	34.7	62.5
256K	65.8	36.8	65.8
512 K	67.6	38.5	65.8

→ IBM Ultrium LTO Tape Drive Test

Transfer Size	*Throughput (Mbytes/sec)*		
	SCSI Baseline	***FibreBridge Speed Write OFF***	***FibreBridge Speed Write ON***
	Writes	**Writes**	**Writes**
64 K	62.5	28.4	62.5
128 K	41.7	27.7	41.7
256K	36.8	25.5	36.8
512 K	35.2	25.3	34.7

There are operating procedures to be taken into account when using the Speed Write feature. Speed Write is established between a FC host/SCSI target pair. If Speed Write is enabled and a Write operation is occurring between a host (host 1) and a SCSI target, and another host (host 2) attempts a command to that same SCSI target, the FibreBridge 'stalls' and waits for outstanding SCSI Write commands from host 1 to complete before issuing the command from host 2. Also, if host 1 sends a non-Write command during its stream of Write commands, the FibreBridge again stalls and waits for the outstanding SCSI Write commands to complete before issuing the non-Write command to the SCSI target. This behavior ensures that data integrity is preserved on the SCSI target device. These concerns are performance related only and should be taken into account when designing the backup system.

Configuring the FibreBridge for the Speed Write option is very simple. All FibreBridges ship from the factory with the feature disabled. A command line interface can be used to

enable Speed Write through an RS-232 serial or Ethernet connection into the FibreBridge using the *Set SpeedWrite* command.

An additional benefit with the way Speed Write is implemented in the FibreBridge is the fact that you can selectively enable/disable on a LUN-by-LUN basis. The advantage is when tape drives and other types of storage devices (hard drives, MO drives, CD/DVD) are connected to the same bridge. Speed Write only adds value when using high-performance tape drives. It is recommended to leave it disabled for all other types of storage devices.

Speed Write is just the beginning of ATTO's effort to enhance customer applications. With an ATTO FibreBridge, limited backup windows are a thing of the past.

The ATTO FibreBridge family of products is the premier choice for taking Storage Area Networks to levels beyond simple connectivity. The embedded 'intelligent Bridging Architecture' is the engine that powers the platform for adding value.

ATTO Technology, Inc.

155 CrossPoint Parkway
Amherst, NY 14068
716-691-1999
www.attotech.com