ATTO Technology, Inc.

***i*ntelligent Bridging Architecture™**

White Paper

Increasing the Backup Window using the ATTO FibreBridge™ for LAN-free and Serverless Backups

White Paper

*i*ntelligent Bridging Architecture™

Increasing the Backup Window using the ATTO FibreBridge™ for LAN-free and Serverless Backups

ATTO Technology, Inc.

Storage Area Networks (SANs) are quickly becoming the method of choice for consolidating Enterprise storage. As SAN deployment grows, so does the need 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. This paper will explore those options and provide a cost-effective solution.

I. Traditional Backup Methods

Traditional methods for backing up data on a server rely upon a separate tape drive(s) directly attached to 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.

Figure 1 depicts a typical backup architecture using Ethernet as the means of passing data. The backup server controls the data flow. During a backup, the backup server will request the data from the fileserver. This request runs through the LAN/WAN and is processed by the fileserver. The fileserver then requests the data from the disk storage subsystem. Once received, the fileserver must process the data and send it through the LAN/WAN back to the backup server. Once the backup server has the data, it still has to Write the data to the tape storage sub-system. During a restore operation, the flow of commands and data is reversed. This model has two major bottlenecks. First is the obtainable bandwidth through the LAN/WAN. The second is the number of total CPU cycles, I/O Bus and memory resource constraints for both servers being used in the tape backup operation.

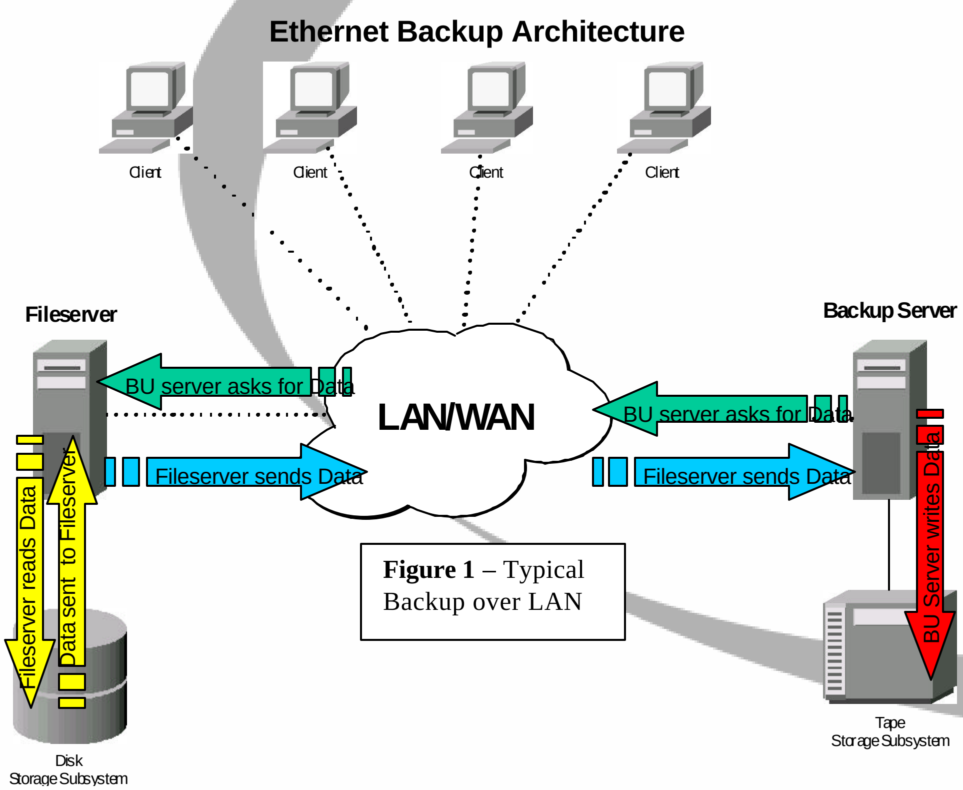

Figure 1 – Typical Backup over LAN

The first strain on the LAN/WAN is the bandwidth requirements during backups and restores. In an Ethernet-based backup model, the LAN/WAN is usually either 100bT or 1000bT, which means either 10 MB/sec. or 100 MB/sec. theoretical speeds. In practice, these speeds are actually between 1-4 MB/sec. and 20-30MB/sec., respectively, during off-peak times. During peak times these numbers drop even lower. With today's tape drives capable of up to 15 MB/sec. native, even one tape drive will strain a 100bT network. For 1000bT, the available bandwidth is strained by only two tape drives. 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, often becoming almost impossible to complete in a given night. It is at this point that IT departments need to consider more modern solutions.

II. SAN-based Backup Methods

Many companies are turning to SANs as an answer to their backup problems. The Storage Area Network, 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.

When it comes to implementing a tape backup solution into a SAN, there are two basic models being followed today:

- Native Fibre Channel Tape Drives
- Fibre Channel Bridge Converting SCSI-based Tape Drives to Fibre Channel

Native Fibre Channel Tape Drives:

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:

- Native Fibre Channel tapes – A new technology
- Switch Port Costs
- Arbitrated Loop Limitations

The technology behind native Fibre Channel tape drives is relatively new. SCSI tape drives have been around for more than a decade and are far more interoperable. Native Fibre tape drives use mature Fibre drives, not tape drives that have internal bridges integrated.

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 $500 - $2000, an eight-drive library can prove to be a costly addition.

Finally, because of this switch port cost, 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 in using a hub is that it 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.

FibreBridge™ to Convert SCSI to Fibre Channel:

The alternative to native Fibre Channel tape libraries is to use SCSI tape drives and robot controllers with a Fibre Channel bridge (see **Figure 2**). As SANs 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. ATTO achieves that through its proprietary *i*ntelligent Bridging Architecture™. It is a key differentiator designed into all ATTO bridge products because its flexibility allows ATTO to match price and performance requirements to each level of the market; it provides a common software platform throughout the line; and it is the platform for adding value into the bridge. This means ATTO customers have a path for adding value into their products via the physical connectivity, the additional functionality of industry-required applications like extended copy or device virtualization, and through customer-specific applications.

Figure 2 – FibreBridge to convert Fibre Channel to SCSI

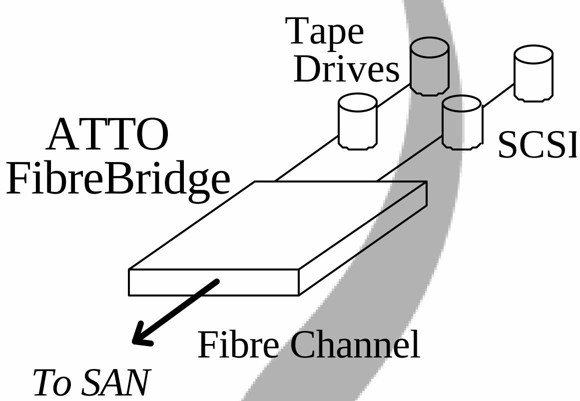

The ATTO FibreBridge™ not only converts the Fibre Channel protocol into SCSI, but also adds value to the SAN because it is an intelligent storage-management device. *i*ntelligent Bridging Architecture is the technology applied to move data efficiently and safely and is a collection of knowledge ATTO Technology has developed over the last 14 years.

The SAN provides two ways of approaching backup: LAN-free and serverless. In LAN-free backup operations, the data move from the storage pool into the backup server and back out to the tape library, as shown in **Figure 3**. This eliminates the need for backup traffic to move across the LAN and allows one server to manage the backup process for all of the disk storage.

LAN-free Backup Advantages:

- Increased Performance
- Centralized and Consolidated
- Storage Pool
- Scalability
- Better Fault Tolerance
- Lower Total Cost of Ownership

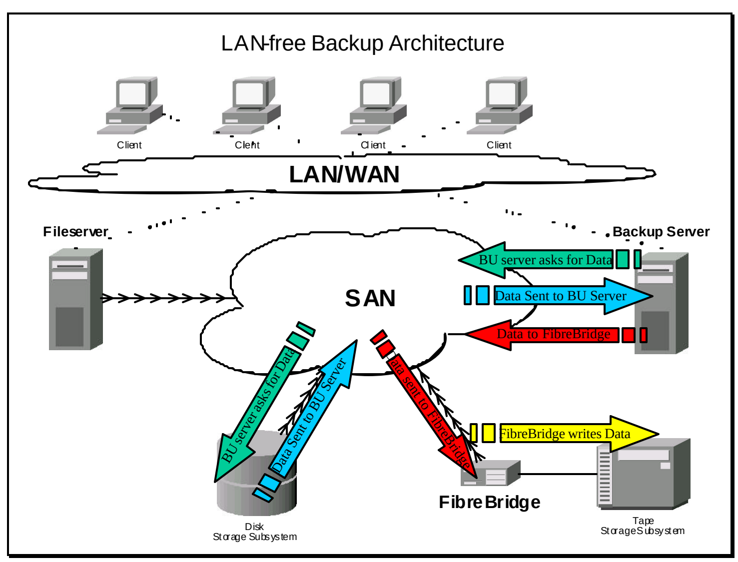

Figure 3 – LAN-free Backup

While the backup server still controls the data flow in a LAN-free architecture, the data path no longer goes through the LAN/WAN. The backup server reads the data from the disk storage subsystem. It then processes the data and Writes it to the tape storage subsystem. Because Fibre Channel protocol does not have packet collisions, the achievable data rates are almost the theoretical data rates of 100 MB/sec. for 1-Gigabit or 200 MB/sec. for 2-Gigabit Fibre Channel. The tape storage subsystem data rate is not restricted like it used to be going across a 30 MB/sec. LAN. Additionally, the LAN/WAN is no longer hindered by backup or restore processes. The servers used to serve files or other applications are able to process more efficiently because they do not have to serve files to a backup server. This reduces the need for additional servers, Ethernet switches, and Ethernet routers.

In a SAN environment, the storage is centralized where all of the storage subsystems, both disk and tape, can be pooled together for all servers to share. In the Ethernet model, every storage subsystem requires its own server, making management difficult. A centralized and consolidated storage pool is more easily managed. In the long run, this saves on staffing resources.

A Fibre Channel SAN is able to scale more effectively in terms of performance as well as storage capacity when compared to the direct-attached model. If more tape performance or capacity is needed, more Fibre Channel switch connections and drives can be added. With the scalability and increased performance over a LAN-free architecture, backup times become drastically shorter, giving the IT Department more time to perform other functions.

If clustered servers are used to manage the backup process, a single cluster can now manage the entire process. Reducing the number of servers required reduces the probability of a failure. Because the storage subsystems are no longer directly attached to a server, the possibility of that server being a point of failure is removed.

Lower Total Cost of Ownership:

All of the advantages of using LAN-free backup combine to considerably decrease the Total Cost of Ownership.

Reduced Backup Times + Fewer Servers + Fewer IT Personnel + Increased Scalabilty + Increased Up Time = Lower TCO

However, even the LAN-free backup architecture can begin to experience bottlenecks in larger installations where many backup or restore operations need to occur simultaneously. Because the data flow is controlled by the backup server, this server can become a bottleneck. By its nature, a backup is very CPU intensive. Running several backups from one server will degrade the performance of the backups significantly by placing a heavy CPU load on the server.

One solution is to increase the number of backup servers on the SAN to distribute the load. While this solution will address the problem, users will begin to lose the advantages of using a LAN-free backup model.

A better alternative would be to implement a serverless backup model.

Serverless Backup:

With serverless backup, also known as Extended Copy, XCopy, or Third-Party Copy, the actual control of the data movement is off-loaded from the backup server to a third party device called a *Data Mover*. ATTO FibreBridges all have the built-in ability to act as the Data

Mover. To implement serverless backups, the backup application (Copy Agent) uses a SCSI command called Extended Copy. This instruction tells the ATTO FibreBridge where the data are and where they are going. The ATTO FibreBridge processes the Extended Copy command, moves the data, and reports back to the Copy Agent with the status. A server is still required to run the backup application, but its only responsibility is to initiate commands and monitor status. Actual data no longer move through it.

Serverless Backup Advantages:

- All of the Advantages of LAN-free Backup, plus
- Increased Performance
- Lower Total Cost of Ownership (compared to LAN-free Backup)

Figure 4 – Serverless Backup

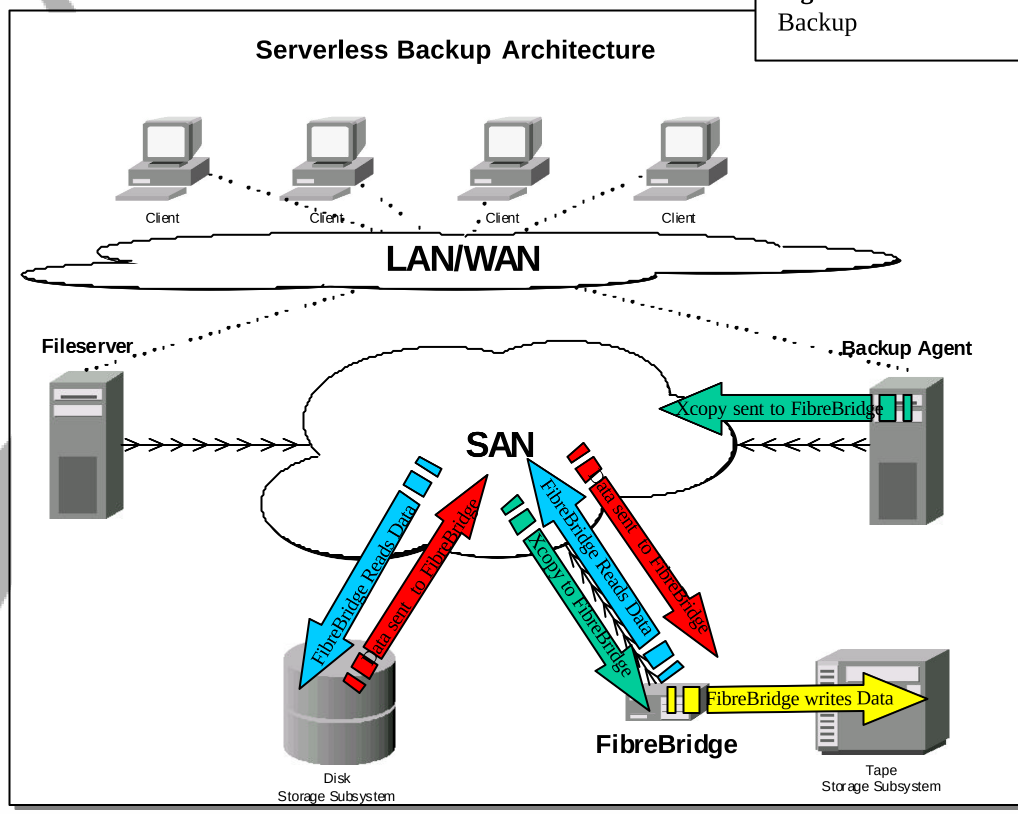

Figure 4 shows a typical Fibre Channel SAN. Notice that the model is exactly the same as the LAN-free backup model, with the only exception being the data path. The data no longer need to pass through the backup server. Data simply move from the disk storage subsystem to the FibreBridge, then to the tape storage sub-system.

In this model, the server is just sending control data to the FibreBridge, using very little server resources and removing the server as the bottleneck. Because the server is not actually moving any data, multiple con-current backups run faster than a standard LAN-free backup, as shown in **Figure 5**. The server itself will also witness additional performance increases since its CPU is not being used by the backup. **See Figure 6**.

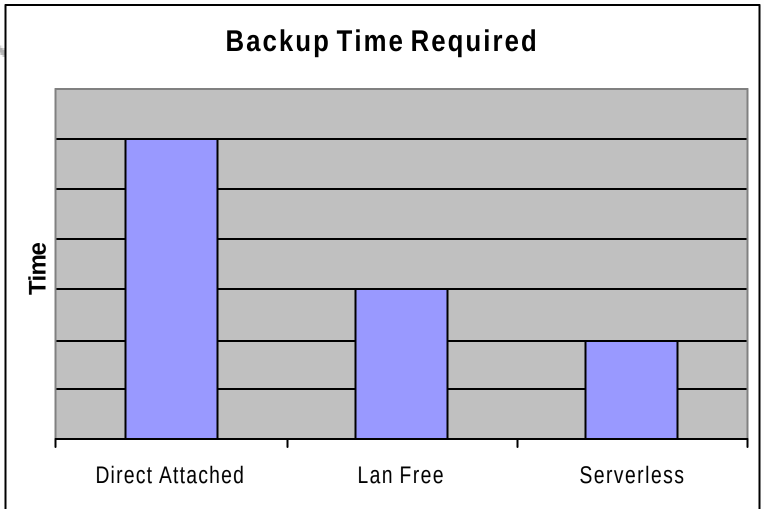

Figure 5 – CPU Utilization Comparison

Lower Total Cost of Ownership:

Once the server bottleneck is removed by implementing a serverless backup environment, it becomes more economical to scale the SAN as capacity demand increases. The FibreBridge is controlling data flow and will now become the bottleneck. This is easily remedied by adding additional bridges to the SAN and distributing the tape drives among them. FibreBridges are less expensive in initial investment, maintenance, and support costs as compared to adding servers. **Figure 7** compares the TCO for a serverless backup architecture to a LAN-free and a direct-attach architecture. The initial investment for a SAN is high for smaller capacity installations, but as the amount of data grows, a SAN is definitely more cost effective.

As the amount of data grows to where a LAN-free backup server becomes the bottleneck, the serverless backup option becomes the most economical.

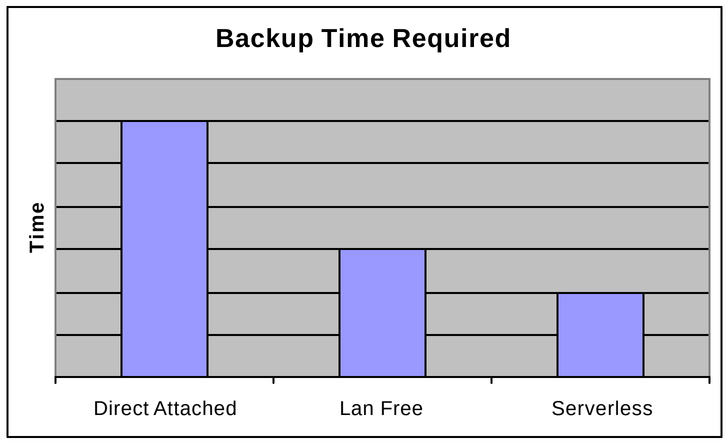

Figure 6 – Required Backup Time

How Does Extended Copy Work:

Extended Copy enables host application software to manage backups performed entirely on the SAN with minimal server interaction. The backup application software on the host creates an Extended Copy command that is sent to a data mover, such as the ATTO FibreBridge, which then performs the individual Read and Write commands on various devices. **Figure 8** provides a breakdown of the Extended Copy command.

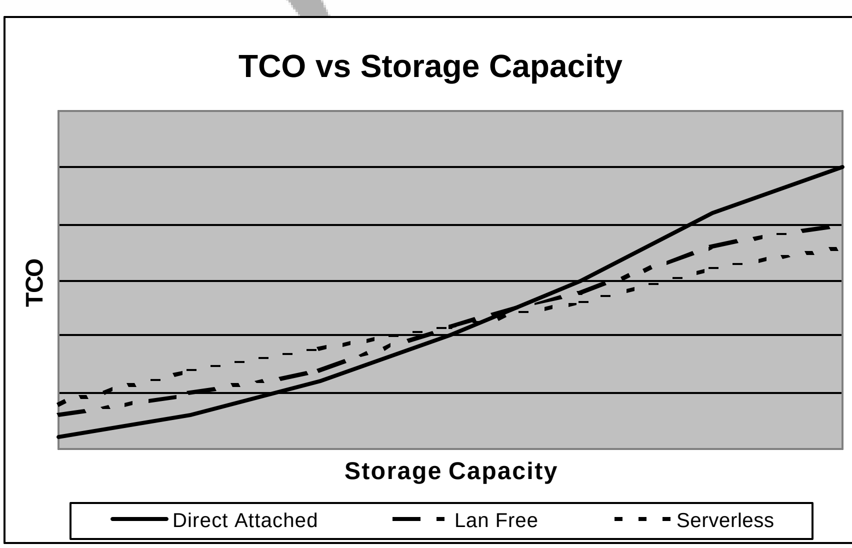

Figure 7 – TCO as Capacity Increases

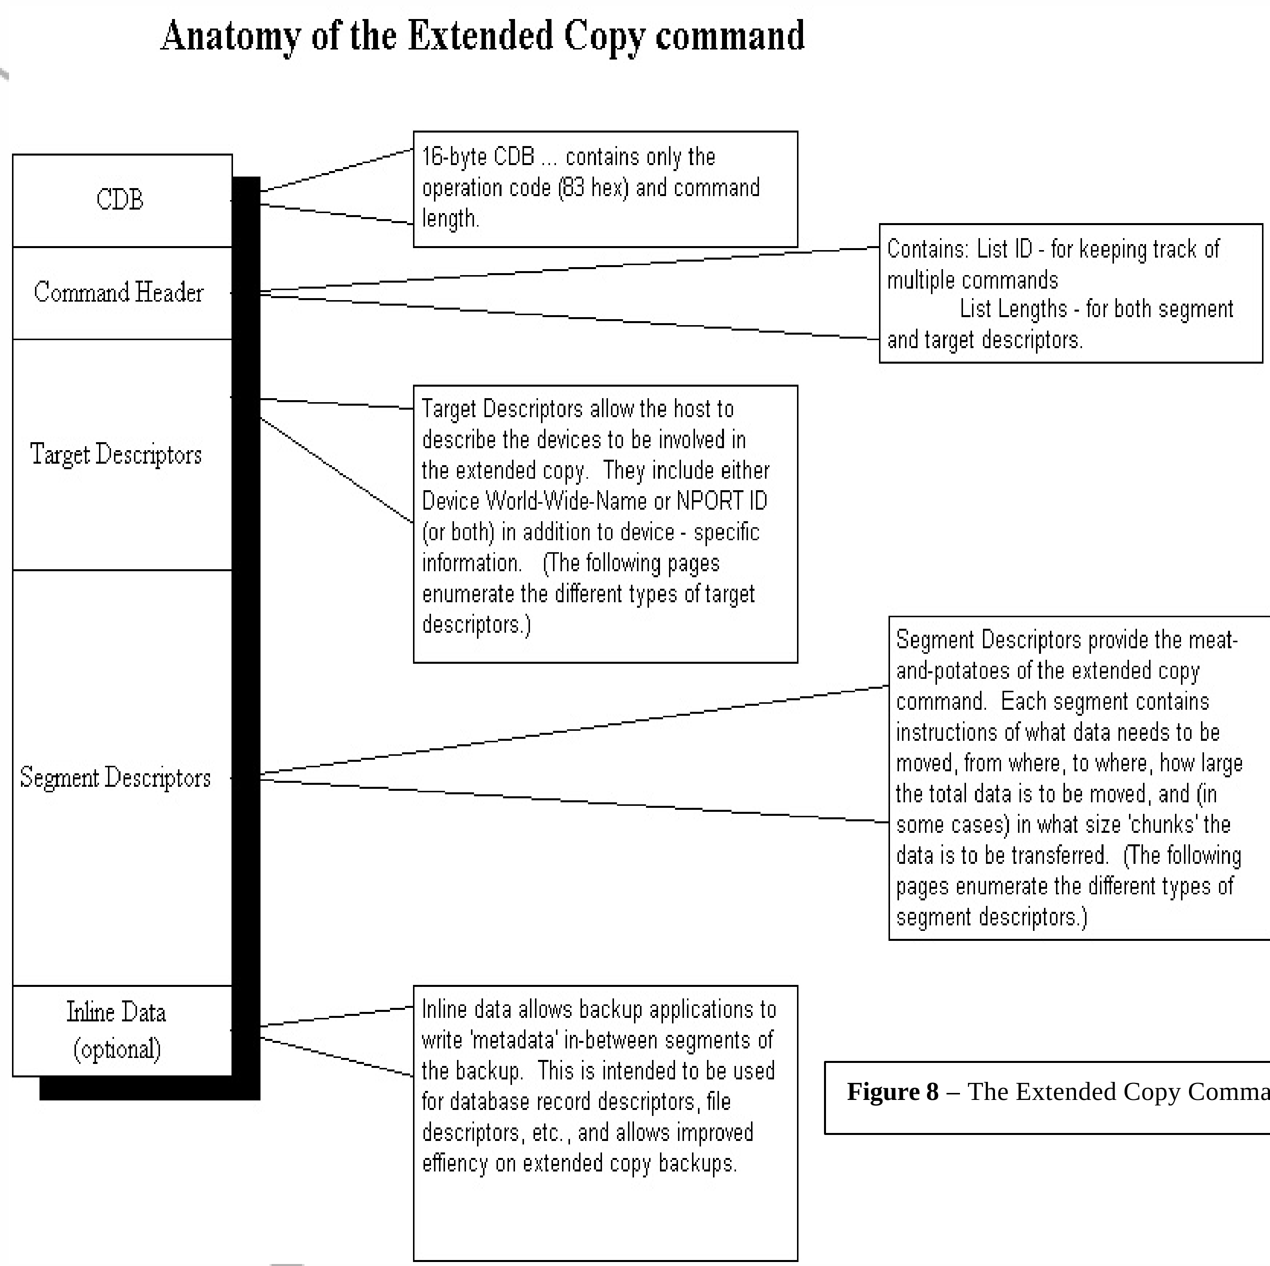

Figure 8 – The Extended Copy Command

Each Extended Copy command contains a list of target devices, called Target Descriptors, and a list of instructions, called Segment Descriptors. Target Descriptors allow the host to describe the devices to be involved in the Extended Copy. These descriptors include a device's World Wide Name (WWN) and/or NPort_ID as well device-specific information. There must be a target descriptor for each device referenced by the Extended Copy Segment Descriptors. Ordering of the target descriptors is important, as the segment descriptor uses the index of the target descriptor to point to the device involved in the copy. **Figure 9** gives a breakdown of both the WWN Target Descriptor and NPort_ID Target Descriptor.

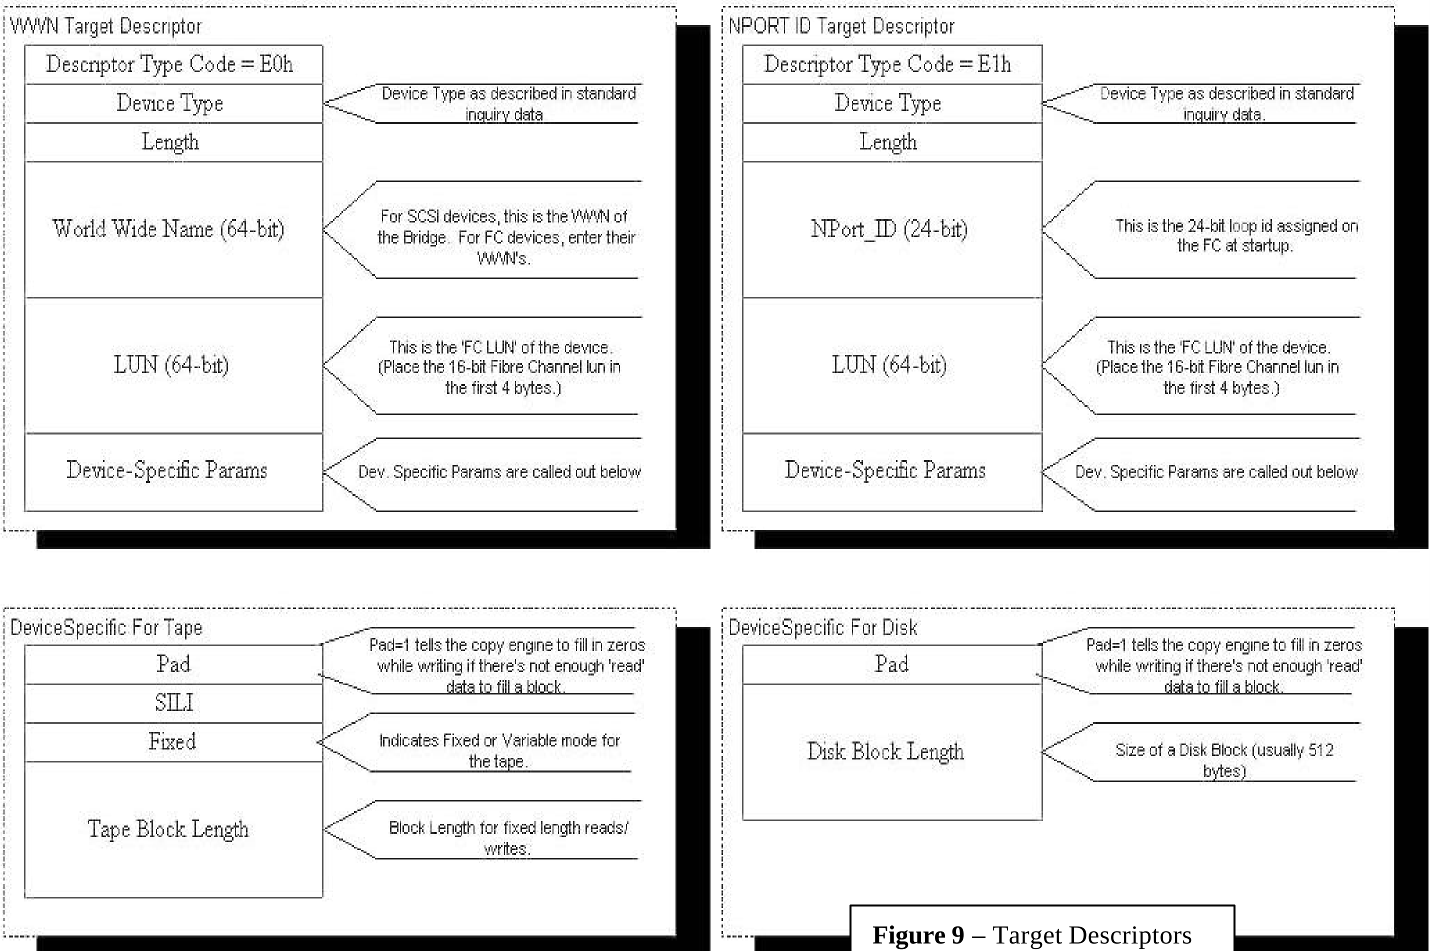

Figure 9 – Target Descriptors

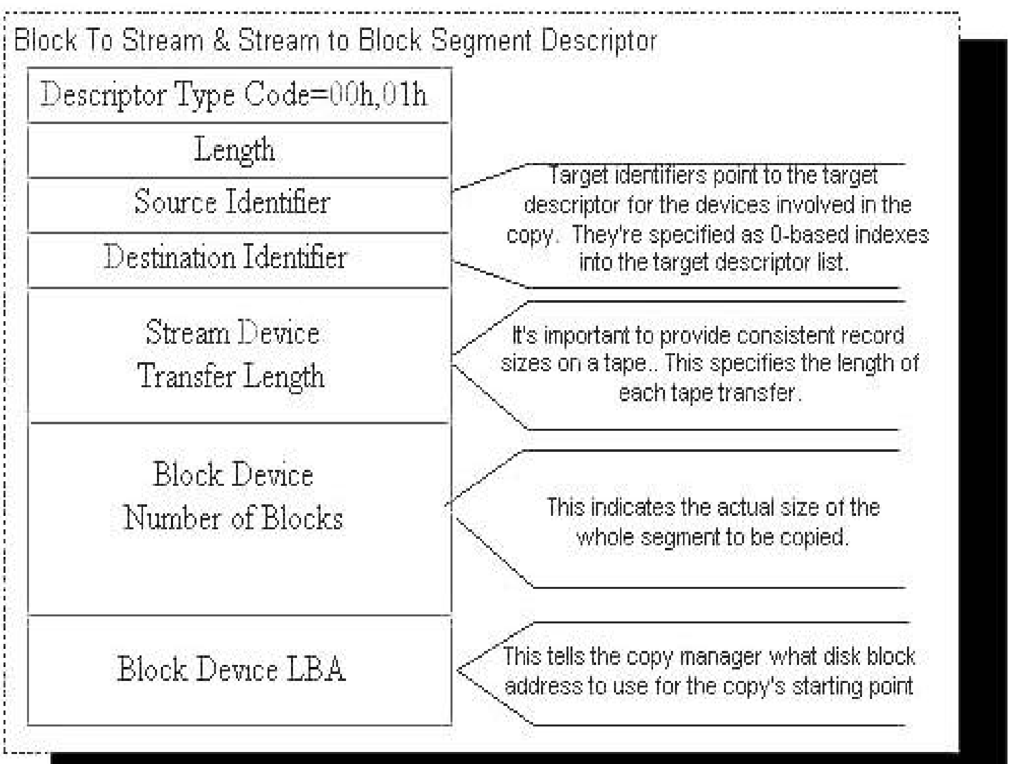

Figure 10 – Segment Descriptors

Segment Descriptors describe the data to copy, where the data are coming from and going to, the length of the data to copy, and size of the data to copy. Two of the most important Segment Descriptors are *Block to Stream* and *Stream to Block*. The *Block to Stream* descriptor instructs the data mover to perform a Read from a block device, such as a disk, and Write to a stream device, such as a tape. The *Stream to Block* descriptor contains instructions for the opposite operation, allowing restoration of files from a tape to a disk. **Figure 10** is a diagram describing these descriptors.

When an Extended Copy command is received by the FibreBridge, it executes each Segment Descriptor it receives. The FibreBridge has a buffer that is able to store data to be copied. Reads are issued to get the data and place them into the buffer. Writes are issued to move the data from the buffer to the destination device. Because the source and destination devices typically have different block sizes, there is a possibility that data can be left over at the end of the Segment Descriptor. This data may need to be placed at the front of the data from the next Segment Descriptor. When this is required, concatenation bits, or CAT bits, in the command indicate what to do about the data. If the CAT bit is set, the data are held over for the next Segment Descriptor. If the CAT bit is not set, the data are either padded or discarded. Once each segment is complete, a status message goes back to the host and the FibreBridge awaits the next command.

Configuring the FibreBridge for Serverless Operation:

All FibreBridges ship from the factory with the feature disabled. There are a variety of ways to enable serverless operation. The simplest is to use the ATTO BridgeTools graphical interface configuration application. It is a Java-based utility that provides an easy way to flash firmware, configure and monitor certain functions of the FibreBridge from any operating system. BridgeTools provides for three methods of connecting into the bridge:

- In-band, directly over the Fibre Channel interface
- 10/100 BaseT Ethernet
- RS-232 Serial interface

Simply connect a host to the bridge and launch the BridgeTools application. Click the button to enable FcInitiator, save the changes and reboot the bridge.

A command line interface can also be used to enable Serverless operation through an RS-232 serial or Ethernet connection into the FibreBridge using the *Set Fcinitiator* command.

The ATTO FibreBridge family of products is the premier choice for taking Storage Area Networks to levels beyond simple connectivity. The proprietary embedded '*i*ntelligent Bridging Architecture' is the engine that powers the platform for adding value and enables ATTO to easily customize the bridge for any application. By implementing a serverless backup application complete with ATTO FibreBridge, users will realize significant reductions in CPU utilization, thereby freeing up the server for more mission-critical functions.