IMPLEMENTING A FIBRE CHANNEL STRATEGY FOR TAPE BACKUP SYSTEMS

ATTO Technology, Inc.
The Power Behind The Storage

INTRODUCTION

The basic architecture for tape backup and archive systems has remained unchanged since the introduction and subsequent wide-scale implementation of the LAN nearly over 15 years ago. Fibre Channel technology allows for the implementation of new tape backup system architectures that will greatly benefit systems administrators, network engineers, and ultimately the customers that rely on these systems.

This article outlines suggested strategy for implementing a Storage Area Network (SAN) solution for tape libraries using Fibre Channel technology.

CURRENT TAPE ARCHITECTURE

The architecture for tape backup that is most often deployed for large systems uses tape backup servers and LAN resources. In this architecture, information that is contained on on-line storage and accessed though file servers or application servers (database servers, data warehouse servers, etc.) is pulled across a 100baseT Ethernet LAN by a tape backup server to which the tape library is attached.

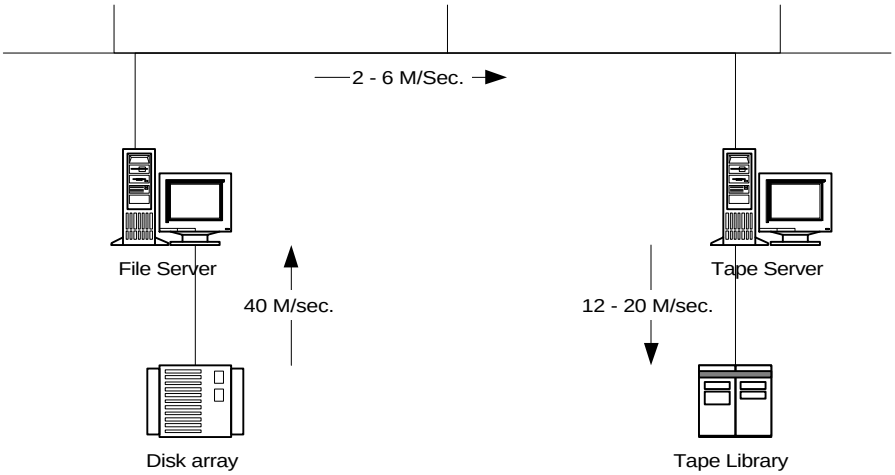

As the diagram points out there are several bottlenecks that act as a potential throttle on backup system performance. First and foremost is the LAN itself. Ethernet is, by storage standards, slow. LAN performance is one of the major reasons for less-than-desired backup performance. Backups also tend to drag down performance overall in the LAN owing to the large amount of data that is streaming through the LAN all at once. This reduces performance for everyone on the network segment that the tape system is attached to, reducing overall system efficiency.

The servers themselves represent a bottleneck as well, especially the file server. The collateral effect of reduced performance in the file or application server is slower overall performance for users attempting to access information through the server. In many cases, the file server cannot service

end-user requests at all during backups. The tape server is less affected, but still represents a choke point for the system. The relative (to disk storage) slowness of the tape units themselves also represent a small, even if less onerous throttle on total system performance.

Two new architectures, Fibre Channel-based LAN-free and Server-less backup, achieve large increases in overall, tape backup, system throughput by eliminating many of the bottlenecks.

STAGE 1: LAN-FREE BACKUP

LAN-free backup is the application of Fibre Channel technology to the tape storage sub-system to increase performance in the overall storage system by eliminating the need to pull data over the LAN and through a file or application server. Typically, it is deployed using a tape server, tape library, and disk-based storage all attached directly to Fibre Channel infrastructure as shown in the diagram below. The tape library is attached to the Fibre Channel network by using a bridge device, such as the ATTO FibreBridge, which also acts as a hardware buffer for incoming data.

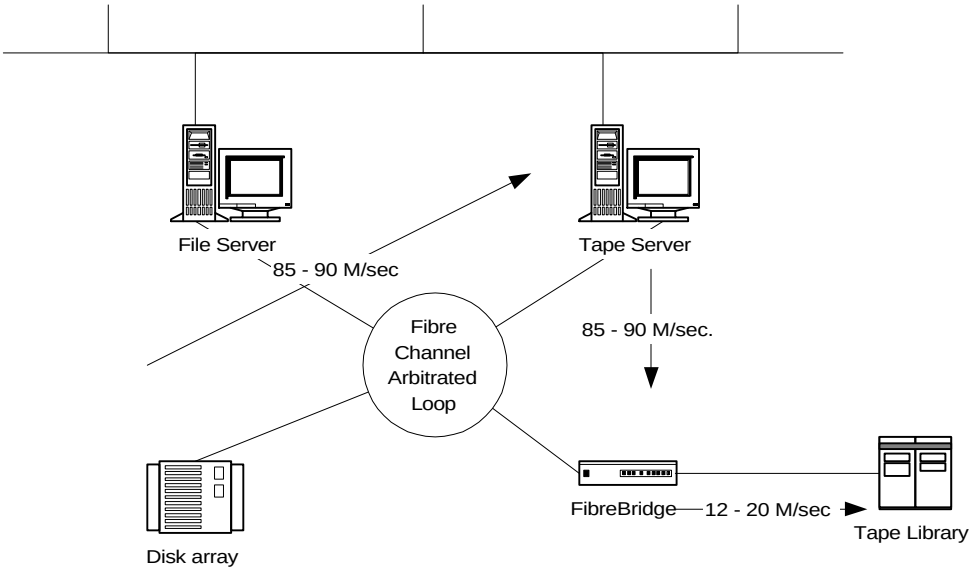

In this design, the tape server can stream data directly from the storage to the bridge device at 85 to 90 M/sec. The only bottleneck is the speed of the tape library itself, and the realized throughput of the tape server itself.

ADVANTAGES OF LAN-FREE APPROACH

The advantage of the LAN-free backup approach is increased throughput to the tape devices and, hence shorter backups. By removing the Ethernet bottleneck (2 - 8 M/sec.), the performance envelope is most affected by the throughput of the tape units themselves, usually between 12 - 20 M/sec. This improves performance by 2.5 to 10 times, an immediately realizable gain in efficiency. Use of tape RAID software can also be used to boost performance by aggregating tape device bandwidth in the same manner as disk RAID.

LAN-free backup can also leverage existing assets, keeping cost of entry to SAN architectures low. Thus LAN-free backup can be viewed as an upgrade to the existing tape storage sub-system, rather than as an entirely new installation.

By having the ability to deploy existing assets in a new manner, thereby extending the lifetime of those assets, a greater ROI is realized on these assets. Time to implement is also significantly shortened, which leads to more immediate benefit to the system and reducing the overall cost of implementation.

Besides the obvious financial benefits, LAN-free backup brings with it service increases that will positively effect end-users. Since the LAN is not involved there are no longer heavily loads placed on the LAN or on application servers during backups. This enables higher overall service levels in the LAN and better application response times.

LAN-free backup helps ensure that backups are complete without disruption to other systems. This in turn reduces adverse effects of backups on normal business operations.

STAGE 2: SERVER-LESS BACKUP

Server-less backup takes the LAN-free solution one step further. In this environment, the tape server is delegated to the role of system coordinator rather than data mover. A copy device such as the ATTO FibreBridge takes on the task of actually moving data from the disk storage to the tape

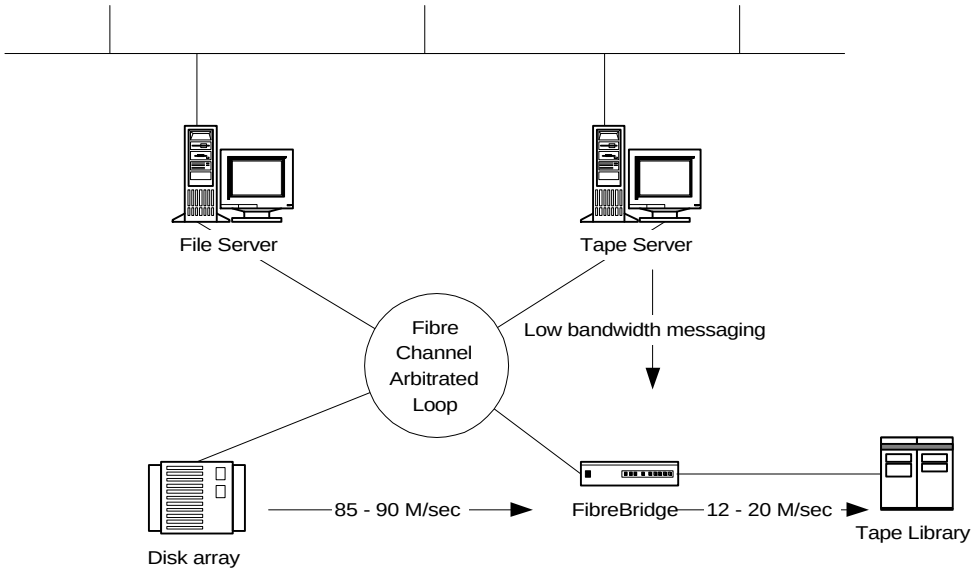

library.

Typically, there are several major elements to a server-less backup solution. First is the hardware infrastructure deployed for LAN-free backup. Second a bridge device such as the ATTO FibreBridge capable of acting as a copy device or independent data movement unit is needed to actual move the data. Finally, special control software such as Legato's Celestra, which issues commands to the copy device and insures smooth operation of the system. A tape server is still necessary, but acts as a place to house the control software more than as a system device dedicated to moving data.

The copy device follows a similar philosophy to network computing devices sometimes referred to as network appliances. It is a specialized device with sufficient and specialized resources to perform a specific rather than general activity within a network or SAN. In the case of a copy device for server-less backup the copy device needs to have enough compute power and memory to support the movement of large blocks of data. The copy device must also support connections to other device that may be involved in the movement of the data, in this case disk drives and tape libraries. Finally, the device must provide a software interface that allow it to interact with software applications that wish to control, manage, and track the movement of data in the SAN. Currently, the Extended Copy Command interface is the most popular interface for these type applications.

In general the market has looked to bridge devices since many of them, including the ATTO FibreBridge, have these attributes.

ADVANTAGES OF SERVER-LESS APPROACH TO END-USERS

There are a number of major advantages to this approach. While not as dramatic as the advantages of LAN-free backup, gains may be found from server-less backup that are not available in any other architecture.

Removing the remaining bottleneck in the system will create performance gains in an important area: The tape server. Even in LAN-free backup, the backup server's performance is directly related to the memory, I/O, and CPU performance of the backup server itself. This inhibitor to optimal performance is eliminated as the data moves through the high-performance copy device, optimized for data movement, rather than through a general-purpose computer bound by multiple needs and a non-specific architecture.

With server-less backup, cost savings may be realized by the elimination of expensive, high-end servers and their replacement with relatively inexpensive copy devices such as the ATTO FibreBridge and low-end control servers. In fact, since software such as Celestra can usually share space on another server, the dedicated tape server may be eliminated altogether for additional savings.

This architecture also makes possible the ability to stream the same data to several tape libraries at once even if geographically separated without the necessity of copying and moving the tapes themselves. This provides an effective element in a disaster recovery plan.

Finally, the system becomes simpler overall. A general-purpose server that requires significant administration and maintenance is replaced with a standalone device that need virtually no maintenance and can be replaced quickly in case of failure. This, coupled with the aforementioned cost savings, results in a **lower Total Cost of Ownership** for the backup sub-system and system infrastructure as well.

SUGGESTED DEPLOYMENT STRATEGY

This is a suggested deployment plan that migrates from a non-shared, parallel SCSI-based solution to a server-less backup solution for new customers

1. Upgrade Existing Systems to LAN-free Backup - by deploying a bridge device and other Fibre Channel components, current systems can be upgraded to a LAN-free backup architecture. This has the advantage of proving out the hardware infrastructure before moving to server-less backup. It is advantageous to start here because the step forward is not as radical and yet provides an immediate enhancement.

2. Upgrade Existing Systems to Server-less backup - upgrade the LAN-free systems to server-less where desired. Since this is likely done through software updates to the bridge device as opposed to hardware additions, this provides a method of achieving this functionality at a low incremental cost with little risk. Any risk is also mitigated by the ability to fall back on the existing LAN-free backup system.

3. Add Capacity - as need for more backup capacity grows, add more inexpensive copy devices and additional tape libraries. It is at this point that the cost effectiveness of this solution becomes apparent. Instead of having to add additional servers (that require extensive administration) and upgrades to the LAN, inexpensive copy devices such as the ATTO FibreBridge are added.

4. Add faster tape units - as the speeds of tape devices themselves increase, so will the overall efficacy and performance of the tape backup sub-system. Since the tape devices themselves are the bottleneck, increases in performance will be immediately realized when the performance of the tape units is increased. The current architecture places control of system performance with the LAN and server rather than the tape unit itself. LAN-free and Server-less backup architectures shift control of bandwidth and hence system performance to the high-speed Fibre Channel network, high-bandwidth copy device, and tape library.

By implementing a migration strategy away from the current server- and LAN-oriented backup systems toward a Fibre Channel based shared tape system, improvements in overall system performance and reductions in total cost of ownership of the system will be achieved. By extending the lifetime of existing assets while increasing overall system performance and reducing operating costs, a better return on investment will be realized.

Finally, end-users and companies will be served better by experiencing less disruption to their overall systems and hence regular business operations. For this reason alone, implementation of this strategy is worthwhile.

