APPLICATION NOTE

Tuning SANs for high availability

A total systems approach to optimizing data availability with a SAN

APRIL 2004

1 Executive summary

Today's business environment of scarce resources and growing expectations requires systems to be available 24 x 7. As always, time is money in the business environment. Today businesses operate 24 hours a day, seven days a week and try to fit 55 weeks in a year. In such a compelling environment, the loss of data could have a significant impact on the top and bottom lines of organizations.

In the past few years, organizations that have deployed Storage Area Networks (SANs) in their IT infrastructure have enjoyed a wide variety of technological and business benefits — increased data availability, ease and centralization of management, and reduced costs. When designing a SAN environment, data accessibility is as important a consideration as overall storage capacity. High availability solutions deployed by StorageTek® take into account not only availability of the storage solutions developed by StorageTek but the entire SAN infrastructure. High availability is not achieved through a single product, but rather through a comprehensive system design that includes all the components of a SAN.

In this paper, we examine some of the basic reasons for pursuing high availability SANs and designing different levels of high availability in your storage architecture. We describe a high availability SAN solution designed and implemented by StorageTek. The solution applies a total systems approach to create an adaptive storage environment that dynamically balances information's changing value with the cost of managing it over time. SAN high availability solutions are an important component to consider in managing information according to its value.

2 Complete SAN solutions

A SAN is a specialized high-speed network that enables fast, reliable access among servers and external or independent storage resources. SANs can be used to connect servers to storage, servers to each other, and storage to storage through hubs, switches, and routers. A SAN typically carries block level I/O traffic between servers and storage devices; it doesn't carry general-purpose traffic such as e-mail or other end-user applications. Thus, it avoids the difficult tradeoffs inherent in using a single network for all applications.

SANs allow the centralization of storage while providing a high level of flexibility in provisioning, distributing, scaling and reconfiguring storage resources. Storage on a SAN becomes a shared resource with centralized management, better utilization of disk and tape resources, and enhanced enterprise-wide data management and protection. In the end, a SAN solution is not a pure hardware-based solution; it includes software and services enabling users to reap the full potential of a SAN solution, as shown in Figure 1.

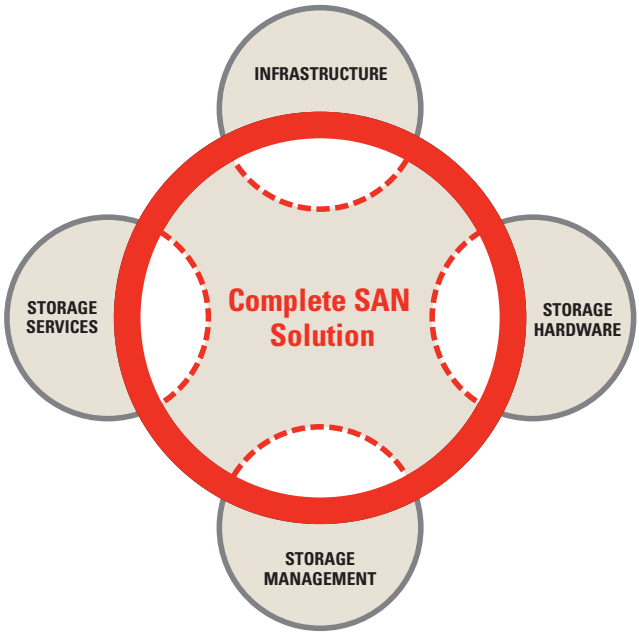

Figure 1. SAN solutions architecture.

These four critical components are integral to the creation of a complete and successful solution. Storage capacity can grow without affecting ongoing operations, its performance can scale along with the SAN topology, and its availability can increase by providing redundant systems and data paths. A switched Fibre Channel fabric can sustain an aggregate bandwidth in the multiple gigabits-per-second (Gb/s) range with very low latency. This translates to faster access to devices and high overall throughput. SANs have also been a key enabler in the process of storage and server consolidation that has helped organizations reap total cost of ownership benefits from resources, increased data protection and greater operational efficiency.

Data availability is a key factor in designing enterprise SANs. Each piece of stored information has a level of importance and criticality to the organization. By configuring a SAN with high availability in mind, constant access to your vital business information can be achieved.

3 High availability SANs

High availability SAN configurations are designed with the goal of achieving true 24 x 7 data access. There are many causes of downtime, as shown in Figure 2, but two stand out: planned downtime and software and hardware failures.

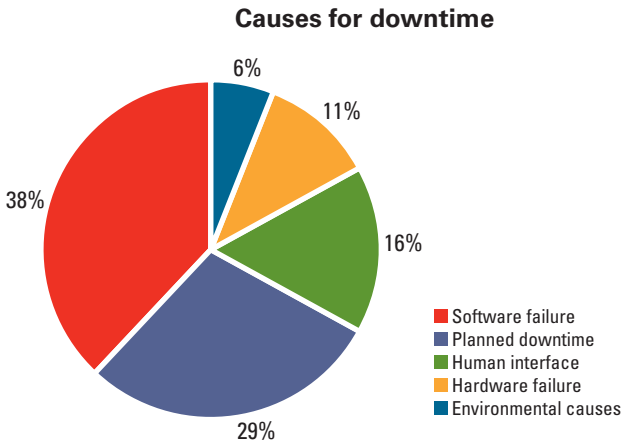

Figure 2. Source: StorageTek internal research.

High availability architectures allow you to eliminate or at least minimize some of these problems to help achieve the goal of true 24 x 7 availability. With the emergence of the Internet and proliferation of global e-business applications, more and more companies are implementing computing infrastructures specifically designed for continuous data and system availability. Faced with increased customer and internal user expectations and reduced budgets, companies are striving to optimize their IT infrastructure to achieve 99.999 percent ("five nines") availability. This level of availability equates to less than 5.3 minutes of downtime a year.

StorageTek uses proven technology and sound design practices to minimize the need for planned downtime, reduce hardware and software outages, and simplify human interfaces to raise the availability of SAN-based storage environments. Availability is much more than system reliability; hardware failure accounts for only 11 percent of overall system downtime. Pursuit of the goal of 99.999 percent availability requires a total systems approach to SAN design.

A true high availability system is not the result of a single product or set of products. Rather, it is a comprehensive system design process that factors in all the components of a SAN, servers, storage and applications, as well as a level of expertise in architecting the overall infrastructure and managing the interaction amongst the various hardware systems.

4 Levels of high availability

Developing high availability SANs starts with mapping the right combination of hardware, software and architecture to derive the appropriate level of service and availability required by an organization. One of the first tasks involves identification of specific availability requirements and predicting potential failures that could be disruptive to the company.

The purpose of a high availability configuration is to reinforce business continuance. Every organization has its own definition of critical data and its own list of applications that require 24 x 7 availability. The loss to an organization of a business-critical application(s) may be as much as five to ten times the cost of the infrastructure. The key is to know which data requires which level of protection. An optimized SAN can provide multiple levels of availability and protection, reducing waste without compromising service levels.

There are different levels of high availability that can be built into a solution to derive an appropriate level of service (see Figure 3). Any combination of the levels can be implemented to provide the level of data availability required for each application.

4.1 Level 1 — network level

Each network device should have built-in fault-tolerance, redundancy, and/or high availability features including power supplies, power cords and hot code load. Dual or redundant SAN fabrics provide network resiliency and diverse pathing. The design of the redundant fabrics may be as simple as two separate Fibre Channel switches or directors, or more complex such as the dual-redundant core-edge design depicted in Figure 3.

4.2 Level 2 — server level

Servers should be configured with at least two host bus adapters (HBAs) providing dual data paths, one connected to each fabric. Additional HBAs may need to be added in order to meet increased performance requirements. If practical, dual servers should be configured in failover or clustered mode to enable application failover. Tape and disk traffic should be separated with independent paths into the dual fabrics, or separate fabrics altogether. This can be accomplished with separate HBAs for tape and disk (HBAs from different HBA vendors are ideal) or at least separate Fibre Channel ports on a dual-ported HBA.

4.3 Level 3 — storage level

Disk storage devices should include the appropriate performance-based RAID data protection and have built-in fault-tolerance and redundancy in both controller design and disk storage connectivity. Redundant disk controllers should be independently linked to the dual fabrics, allowing for increased performance and multi-path access to the disk storage. Two disk arrays can also be linked together, allowing critical data to be duplicated on a second array for almost instantaneous data recovery. There are numerous methods of mirroring or duplicating data between storage arrays. Your StorageTek representative can assist you in determining the appropriate solution for your organization's needs. Tape devices require a different set of considerations. To build resiliency into the solution, provide enough tape devices to deal with simultaneous backups and restores, device failures, and/or increased workloads. Resiliency is based on the number of tape drives, not on the fault-tolerance or redundancy of the individual tape devices. In addition, tape devices can only connect to a single fabric and have no software capability for dual or failover pathing. This is true even for devices like the StorageTek T9840 and T9940 tape drives that have two Fibre Channel ports. The highest level of resiliency is found by distributing multiple tape drives across multiple fabrics.

4.4 Level 4 — total systems level

Designing a solution that combines all the recommendations discussed above can produce an architecture with no single point of failure across servers, data paths and storage.

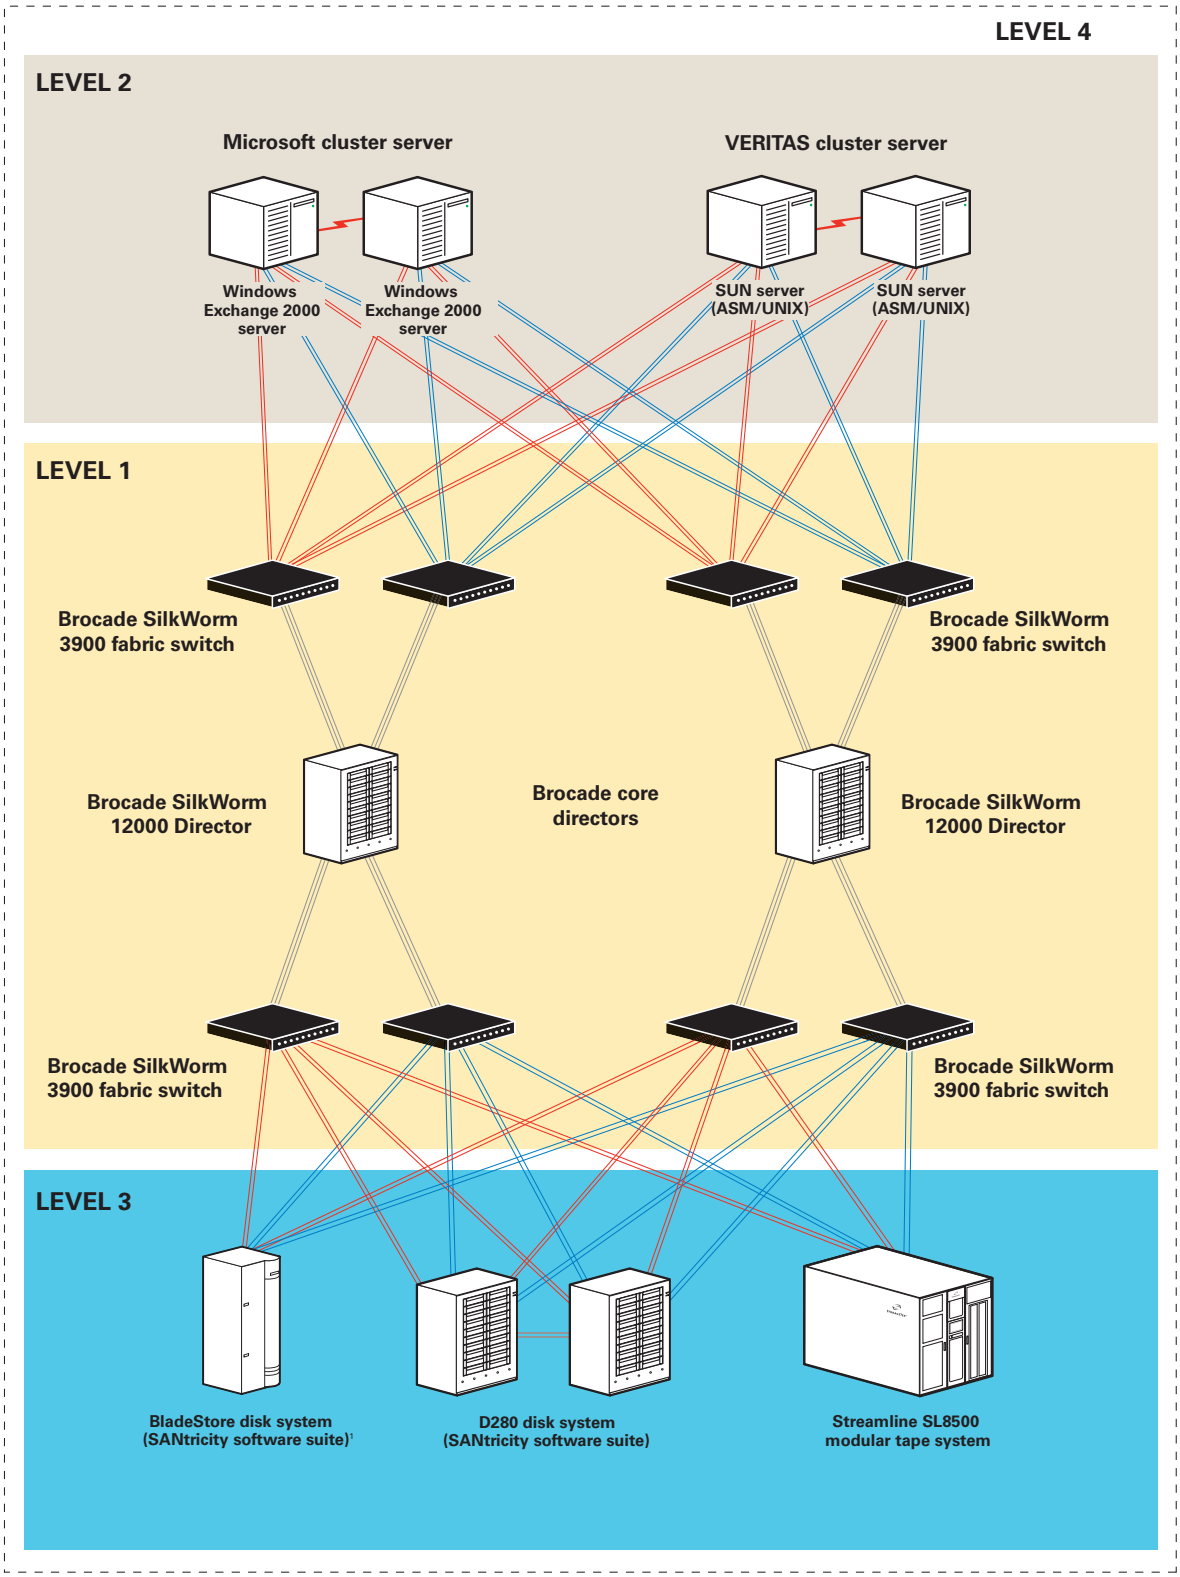

Figure 3. Different levels of high availability.

[1] SANtricity software suite, brought to you through a strategic alliance of StorageTek and LSI Logic Storage Systems, provides management of the disk systems.

The high availability SAN architecture not only provides for redundancy across the different components, but can also be tuned to allow load sharing that would significantly improve the performance of the entire architecture. Traditional high availability architectures are based on an active-passive configuration in which the failover devices do not contribute to the IT environment until the primary devices fail. Now it is possible to use an active-active (two active load sharing devices) configuration where all devices are active, sharing the load to maximize the performance of the environment.

The following example illustrates the power of a properly tuned high availability SAN configuration to enhance the overall value of your IT infrastructure.

5 Example: an e-mail SAN

As a leader in the storage industry over the past 30 years, StorageTek has built considerable expertise in designing and implementing high availability solutions targeted at specific customer problems. StorageTek's strategic partnerships with industry leaders such as VERITAS and Brocade help us design and deploy complete high availability solutions.

The following highly available archival solution for an e-mail application was designed and implemented for a customer using redundant fabrics. This configuration provides storage access for heterogeneous server platforms with a shared SAN environment. The customer is able to manage data across different levels of storage while maintaining the required levels of data availability as the data's value changes.

This solution, shown in Figure 4, utilizes the four levels referenced in Figure 1 to meet the customer's unique availability requirements for their e-mail system.

Level 3 (storage) availability was provided with the D280 disk system, the BladeStore (B-Series™) disk system and the L5500 tape library. These components provide the crucial data protection and multiple tiers of data storage required to support a bottomless e-mail configuration.

Level 1 (network) availability required a redundant fabric architecture using Brocade SilkWorm 3900 enterprise switches. The fabric not only provides high-speed connectivity but also provides a failover mechanism that eliminates single points of failure.

Level 4 (total system) availability is powered by a set of software applications that managed individual elements in the solution and the whole data flow in the architecture. Additional storage resource management (SRM) software could be added to the solution supplying overall management of the individual layers of the solution. The failover support between the primary storage devices is supported by Remote Volume Mirror software that performs non-disruptive replication of the mission-critical data.

StorageTek Global Services personnel assessed the e-mail environment, designed a SAN architecture to optimize the customer's e-mail infrastructure, and implemented the solution. StorageTek centrally manages ongoing maintenance through its TekCare™ Support Services and Remote Managed Storage (RMS) services.

6 The rationale for the design

The high availability e-mail solution was designed to eliminate all single points of failure. The redundant architecture (including clustered servers) allows for significant load sharing throughout the SAN, improving the performance of the whole solution.

The D280 disk system provides active primary storage. The disk systems are tuned for high performance and fast retrieval of data back to the users. Heterogeneous host support, high availability characteristics and online management functionality made D280 disk system an appropriate choice for this task.

The mirrored subsystem helps to achieve instantaneous recovery in case the primary storage is inaccessible for any reason. With non-disruptive mirroring, there is no need for expensive application server processing that would affect the performance of business applications.

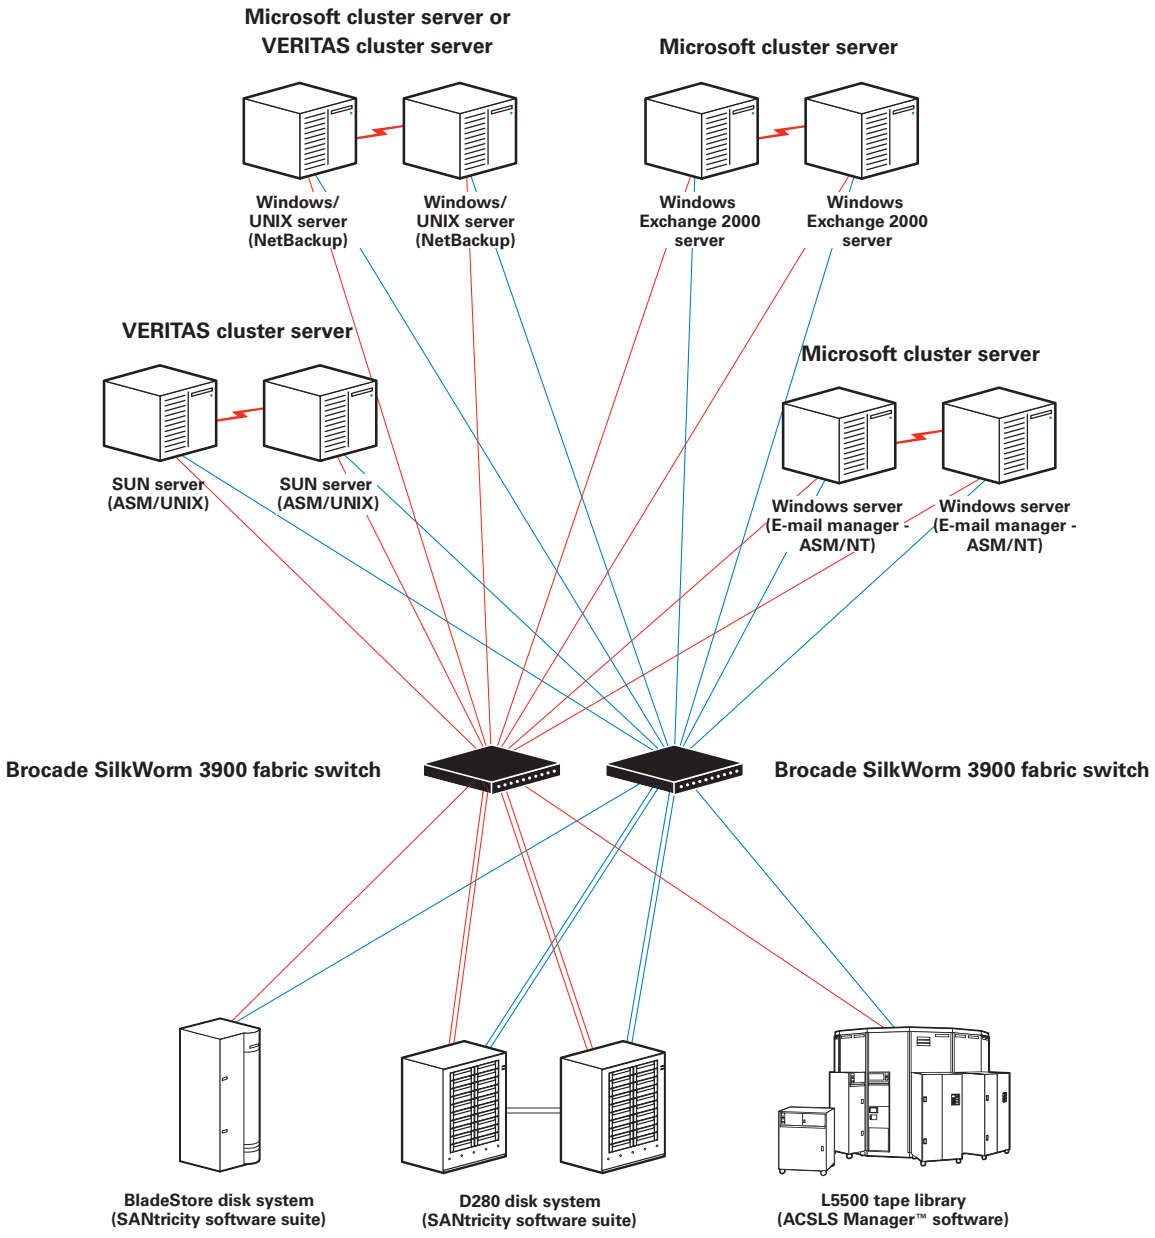

Figure 4. Customer architecture for a high availability e-mail solution.

Disk-to-disk data protection is provided by BladeStore disk systems. As e-mail ages and becomes less critical, it is automatically moved by StorageTek's Application Storage Manager® (ASM) software from the D280 systems to the less expensive ATA-based BladeStore systems. The data is seamlessly accessible by the users regardless of where it resides. Expensive primary storage is devoted to the most recent and frequently accessed e-mail data. The BladeStore system scales to 150 terabytes per system, making it an ideal disk-based data protection option.

Archival storage through the L5500 tape library provides near infinite capacity (13.2 petabytes native capacity) for archival data, while enabling the customer to share tape resources with backup and recovery operations. The L5500 tape library provides long-term high-capacity storage, while freeing up crucial disk resources.

The different levels of storage are seamlessly integrated through a combination of enterprise-class switches and ASM software. ASM software is implemented on a Sun server to manage user accounts to monitor access and perform policy-based automatic data migration. Application Storage Manager software implemented on the e-mail manager server enables monitoring of e-mail traffic and performs policy-based automatic data storage management. This allows the disk subsystems to be kept small and to operate at peak performance. When the data ages beyond a predetermined threshold, the ASM software automatically migrates data from D280 online disk systems to the BladeStore disk system.

The customer had been using VERITAS NetBackup for backup and restore operations. The new solution uses NetBackup to schedule backup jobs to move data from the BladeStore systems to the L5500 tape library. This reduces the need for additional software and training of personnel to implement the new system.

The new solution makes continuous data protection possible across all levels to minimize the risk of data loss while managing data availability based on the value of the data. Balancing data protection and availability needs in this way reduces costs and management complexity by eliminating waste in storage infrastructures and processes.

6.1 Bottom line impact

Based on customer-provided data, the estimated annual productivity cost savings from this solution was $2,013 per e-mail user. Total productivity cost savings of moving to the high availability SAN solution for e-mail was estimated to be $6.5 million annually.

Staffing is often overlooked when considering the cost of storage. Prior to the implementation of this solution, administrators were spending approximately 20 percent of their time managing the e-mail environment and retrieving lost e-mail for users. This solution helped to significantly lower the administrators' workload by automating backup and recovery and providing seamless access for end users with a bottomless mailbox. Administrator time required to manage the e-mail environment was reduced to 10 percent by automating e-mail management with a powerful search tool, improving system predictability and simplifying management processes.

7 Going forward

The solution presented above was designed for a specific customer. When adopting a similar solution, the following factors may affect your design:

- The server environments and the applications running on the servers
- The existing fabric design and scalability options
- The mix of storage devices currently in your environment and available free space on these devices
- Existing management platform and options to seamlessly plug into these environments

Designing a storage environment can be a challenge. Contact your StorageTek representative to find out how we can help customize a solution that would maximize benefits in your IT infrastructure.

ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK), a $2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151

MN 9125 A 04/04