SAN over 10G *ip* ™

Executive Summary

Data is growing at explosive rates in today's businesses. A typical data record has tripled if not quadrupled in size in just the last five years. Multiply this times thousands of records of customers or other information bases and it becomes glaringly apparent that the size of the databases is growing exponentially. Aside from a company's human resources, data has become the most valuable corporate asset both tangibly and intangibly. How to effectively store, access, protect and manage critical data is a new challenge facing IT departments. A Storage Area Network (SAN) applies a networking model to storage environments in the data center. The SANs operate behind the servers to provide a common path between servers and storage devices. Unlike server-based Directly Attached Storage (DAS) and file-oriented Network Attached Storage (NAS) solutions, SANs provide block-level or file level access to data that is shared among computing and personnel resources. The predominant SAN technology is implemented in a Fibre Channel (FC) configuration, although new configurations are becoming popular.

With the growth of SANs and the worldwide domination of Internet Protocol (IP), storage using IP networks to transport storage traffic is in the forefront of technical development. IP networks provide increasing levels of manageability, interoperability and cost-effectiveness. By converging the storage with the existing IP networks (LANs/MANs/WANs) immediate benefits are seen through storage consolidation, virtualization, mirroring, backup, and management. The convergence also provides increased capacities, flexibility, expandability and scalability. The two main standards utilizing the IP protocol are FCIP (Fibre Channel over IP), also known as iFCP in hybrid form, and iSCSI (ip Small Computer System Interface). Both carry either Fibre Channel or SCSI commands incorporated into an IP datagram. The difference between the two is that SCSI can work with existing Ethernet devices, while the FCIP or iFCP, alternatively defined as Fibre Channel Tunneling, can only work with Fibre Channel Components. Both were developed by the IETF (Internet Engineering Task Force). Tunneling is the encapsulation of Fibre Channel commands within an IP packet for transmission over an IP network.

Today, 10Gigabit Ethernet is becoming increasingly popular in the backbones of corporate data centers. Gaining a competitive edge from deploying 10 Gigabit Ethernet in the enterprise requires a robust IT infrastructure. Based on Siemon's commitment to reliability and quality, 10G *ip* ™ is a reliable foundation for data centers, backbone networking components and SAN networking. With 10G *ip* ™, users are provided with an open and standard based structured cabling system that can support multiple converged applications on one cabling infrastructure.

This document provides some insight into new storage technologies. It addresses both currently adopted standards as well as a brief look at standards that are in development and expected to be published.

Evolution of Storage Technology

With the advent of the Internet, corporate intranets, e-mail, e-commerce, business-to-business (B2B), ERP (Enterprise Resource Planning), Customer Resource Management (CRM), data warehousing, CAD/CAM, rich media streaming, voice/video/data convergence, and many other real time applications, the demands on the enterprise storage capacity has grown by leaps and bounds. The data itself is as important to a business's successful operation as its personnel and systems. The need to protect this strategic asset has far exceeded the capabilities of a tape backup. Tape access and capacities have simply not addressed all of the issues at hand. Growing data stores meant having to implement tape libraries. Even then, however, there are inherent issues with tape media that could only be addressed with either supplemental storage or replacement of the media altogether.

Downtime is one critical factor in today's businesses. Based on recently published study by Contingency Planning Research and Internetweek (4/3/2000), 77% of the responses indicated an average down time cost of $20,333 per hour with 4% indicating that downtime costs are in excess of $50,000 per hour. These costs alone have pushed the storage industry to provide redundancy and high-availability. Further, Federal mandates for the medical and financial industry have created yet another mandate for security and high availability. The greatest advantage of utilizing the IP protocols is that security and encryption can be built into the communications, providing a distinct advantage over older methods. Storage network technology has developed in the following three main configurations: Direct Attached Storage (DAS), Network Attached Storage (NAS), and Storage Area Networks (SAN). DAS is the traditional method of locally attaching storage devices to servers via a direct communication path between the server and storage devices.

As shown in Figure 1, the connectivity between the server and the storage devices are on a separate dedicated path than the connectivity to the network. Access is provided via an intelligent controller. The storage can only be accessed through the directly attached server. This method was developed primarily to address shortcomings in drive-bays on the host computer systems. When a server needed more drive space, a storage unit was attached. This method also allowed for one server to mirror another. This functionality is also was accomplished via directly attached server to server interfaces.

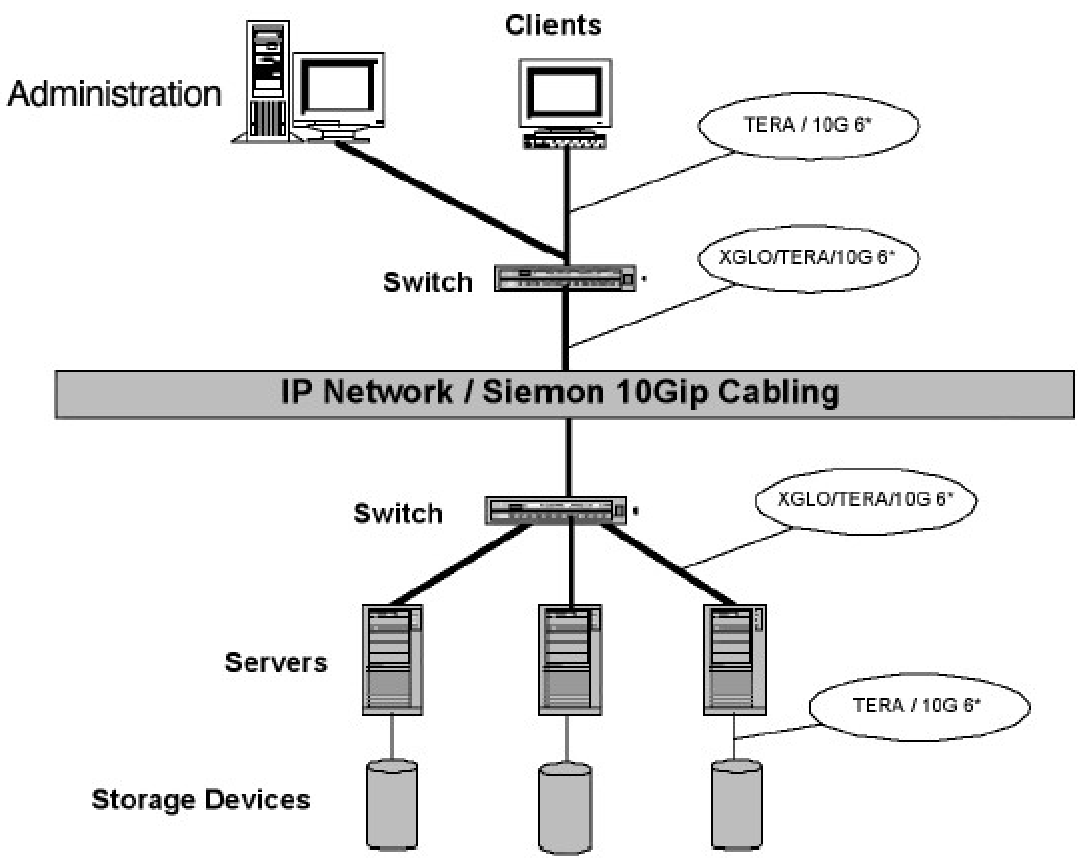

Figure 1: A simple DAS Diagram

Footnote: *10G 6 implementations are based on theoretical channel capacity of unshielded category 6/class E cabling. 10Gb/s transmission over UTP is under evaluation by the 10G-BASE-T study group and may be limited in length or require some mitigation of ANEXT depending on capabilities within the active electronics.

Network Attached Storage (NAS)

NAS is a file-level access storage architecture with storage elements attached directly to a LAN. It provides file access to heterogeneous computer systems. Unlike other storage systems the storage is accessed directly via the network (see Figure 2). An additional layer is added to address the shared storage files. This system typically uses NFS (Network File System) or CIFS (Common Internet File System) both of which are IP applications. A separate computer usually acts as the "filer" which is basically a traffic and security access controller for the storage. The advantage to this method is that several servers can share storage on a separate unit. Unlike DAS, each server does not need its own dedicated storage which enables more efficient utilization of available storage capacity. The servers can be different platforms as long as they all use the IP protocol.

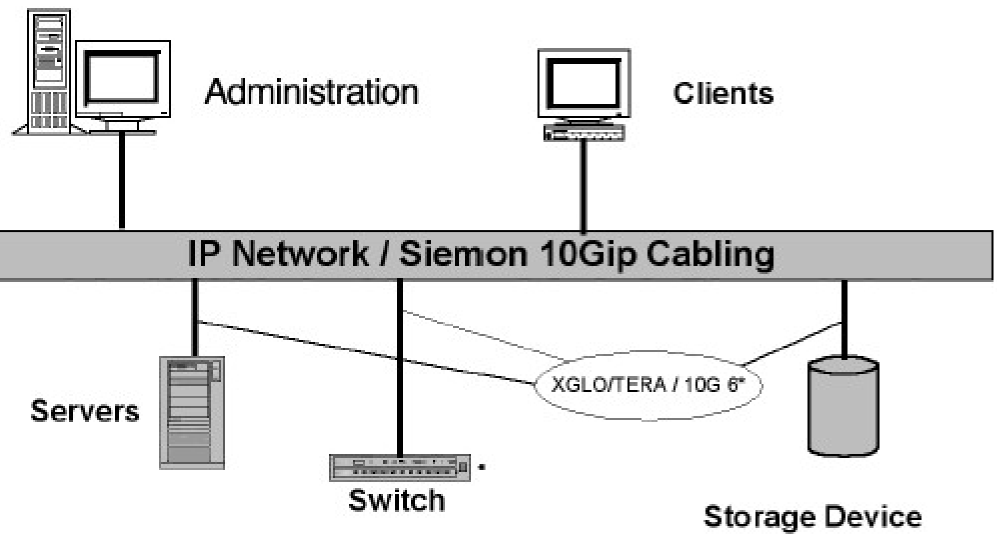

Figure 2: A simple NAS Diagram

Storage Area Networks (SANs)

Like DAS, a SAN is connected behind the servers. SANs provide block-level access to shared data storage. Block level access refers to the specific blocks of data on a storage device as opposed to file level access which refers to access at the file level. One file will contain several blocks. SANs provide high availability and robust business continuity for critical data environments. SANs are typically switched fabric architectures using Fibre Channel (FC) for connectivity. If you refer to Figure 3 on page 6, the term switched fabric refers to each storage unit being connected to each server via multiple switches which provides redundancy within the paths to the storage units. This provides additional paths for communications and eliminates the switch as a single point of failure.

According to IDC's Worldwide Disk Storage Systems report, worldwide storage systems factory revenues were $4.8 billion in the first quarter of 2003. The overall SAN market has grown nearly 14% annually in terms of factory revenue, and the investment in storage systems will reach $120 billion worldwide by 2005. IP based SANs represent a significant portion of the storage solutions today. In fact, surveys conducted by CompTIA in 2002 indicate that nearly 70% of mid-tier companies will consider deployment of IP based SANs in whole or in part within the next a few years.

Ethernet has many advantages similar to Fibre Channel for supporting SANs. Some of these include high speed, support of a switched fabric topology, widespread interoperability, and a large set of management tools. In a storage network application, the switch is the key element. With Ethernet ports showing a significant increase in numbers shipped, especially for Gigabit and to 10-Gigabit. Leveraging IP and Ethernet for storage is a natural progression.

SAN over IP

IP was developed as an open standard with complete interoperability of components. Two new IP storage network technologies are FC over IP (FCIP or iFCP a hybrid form) and SCSI over IP (iSCSI). FCIP supports Fibre Channel communication across a standard IP network via Fibre Channel Tunneling or storage tunneling and has the benefit of utilizing storage in locations that may exceed the directly attached limit of nearly 10 km if fiber is used as the transport medium. Internal to the data center, legacy Fibre Channel can also be run over coaxial cable or twisted pair cabling, but at significantly shorter distances. iFCP, the hybrid form carries iSCSI commands over Fibre Channel. iSCSI maps SCSI commands, data and status over Ethernet networks. It has the advantage of operating over geographically disparate networks without needing to be attached to the Fibre Channel fabric as it uses existing Ethernet connectivity.

The incorporation of the IP standard into these storage systems offers performance benefits through speed, greater availability, fault tolerance and scalability. These solutions, properly implemented, can almost guaranty 100% availability of data. The IP based management protocols also provide network managers with a new set of tools, warnings and triggers that were proprietary in previous generations of storage technology. Security and encryption solutions are also greatly enhanced. With 10G gaining popularity and the availability of new faster WAN links, these solutions can offer true storage on demand.

Fibre Channel over IP (FCIP)

Legacy Fibre Channel (FC) Overview

Legacy FC is a standards-based SAN interconnection technology within and between data centers limited by geography. It is an open, high-speed serial interface for interconnecting servers to storage devices (discs, tape libraries or CD jukeboxes) or servers to servers. FC has large addressing capabilities. Similar to SCSI, each device receives a number on the channel. It is the dominant storage-networking interface today. The Fibre Channel can be fully meshed providing excellent redundancy. FC can operate at the following speeds: 133Mb/s, 266Mb/s, 530Mb/s, 1Gb/s, 2 Gb/s, with 1Gb/s and 2Gb/s being predominant. The transmission distances vary with the speed and media. New 10G Fibre Channel switches are already being shipped. According to the IEEE 802.3ae standard, when using single mode optical fiber cables, the distance supported is 10 kilometers, and up to 300m when using new 50 Micron multimode fiber. Laser optimized fiber is an important consideration in fiber selection for 10Gb/s transmission.

Fibre Channel supports three different connection topologies: point-to-point, arbitrated loop, and switched fabric. Switched fabric, as the name implies, is the better solution as it allows for a mesh within the Fibre Channel. It may also be configured in what is known as fabric islands. Fabric islands connect geographically diverse Fibre Channel fabrics. These fabrics may be anywhere within the range of the medium without IP. With IP, the fabric can reach greater distances as it is extended by routers and links outside of the fabric. They may also comprise different topologies (cascaded, ring, mesh, or core-to-edge), but may require additional connectivity for shared data access, resource consolidation, data backup, remote mirroring, or disaster recovery.

FCIP Topology

A basic Fibre Channel IP configuration is shown in Figure 3. In order to overcome distance limitations of legacy FC, the tunneling protocol Fibre Channel over IP (FCIP) was developed. FCIP supports FC communication across a standard IP network, enabling the interconnection of FC SANs, which can be dispersed over an entire enterprise, regardless of geography. For replication between storage systems over a wide area network, FCIP provides a mechanism to interconnect islands of FC SANs over the IP infrastructure (LANs/MANs/WANs) to form a single, unified FC SAN fabric.

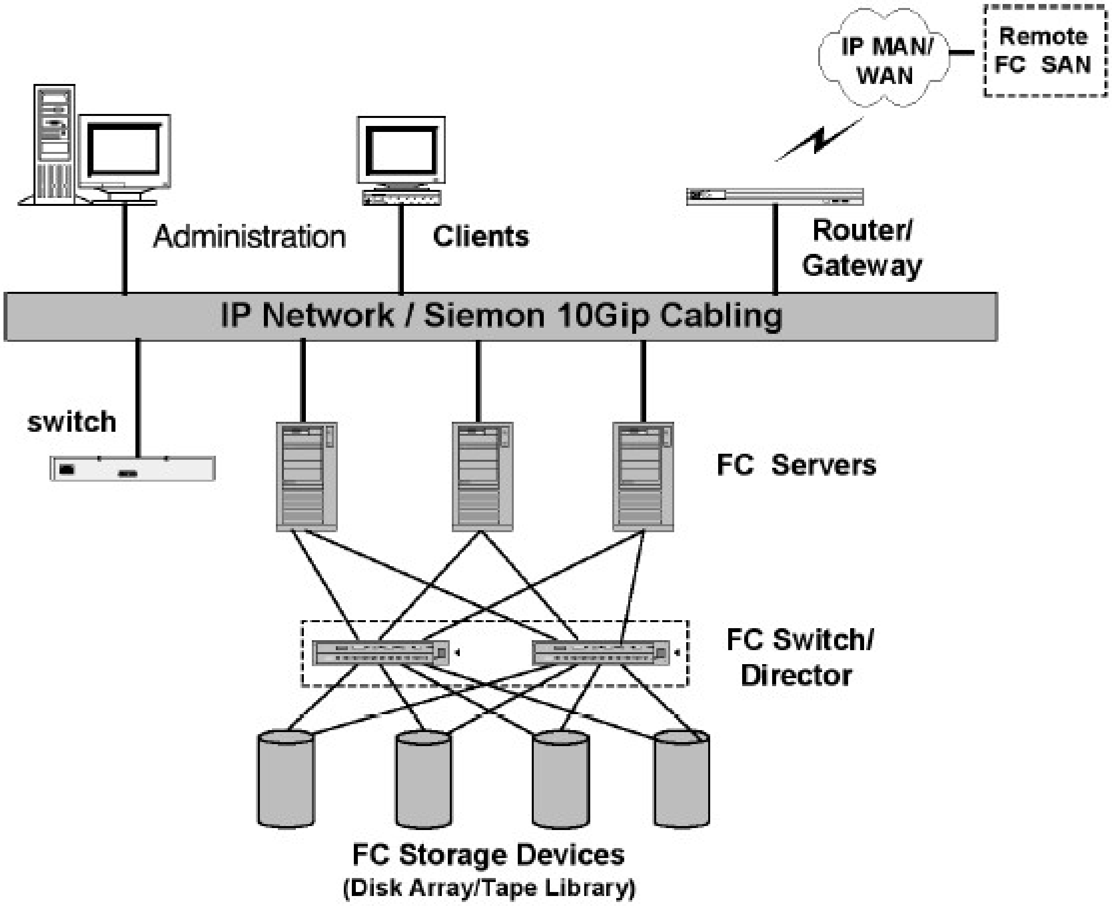

Figure 3: A Simple FCIP SAN Diagram

Typical Component and Elements

Fibre Channel hardware interconnects storage devices with servers and forms the Fibre Channel fabric through the connection of the following:

- Interconnect device: switches, directors
- Translation devices: Host bus adapters (HBAs installed in server), adapters, bridges, routers, and gateways
- Storage devices: non-RAID or RAID (Redundant Array of Independent Disks) disk arrays, tape libraries
- Servers: The server is the initiator in the Fibre Channel SAN and provides the interface to an IP network. Servers interact with the Fibre Channel fabric through the HBA.
- Physical layer/media: Coax, twisted-pair and/or fiber-optic cables

The FC SAN switches are classified into the fabric as either switches or directors. A fabric switch contains a low to medium port count, while a director is a high port count switch (generally above 64 ports). Fibre Channel switches can be networked together to build larger storage networks. The HBA is more complex than a traditional Ethernet card. It connects the Fibre Channel network to the IP network. A bridge may be used to connect legacy SCSI or ESCON (Enterprise System Connection) storage devices to the Fibre Channel network.

By combining the advantages of both SCSI and networking, iFCP, FCIP's hybrid successor, allows FC devices to communicate over networks using the iSCSI commands and traditional Ethernet switches.

Small Computer Systems Interface (SCSI) over IP (iSCSI)

The iSCSI protocol unites storage and IP networking. iSCSI uses existing Ethernet devices and the IP protocol to carry and manage data stored in a SCSI SAN. It is a simple, high speed, low-cost, longdistance storage solution. One problem with traditional SCSI attached devices was the distance limitation. By using existing network components and exploiting the advantages of IP networking such as network management and other tools for LANs, MANs and WANs, iSCSI is expanding in the storage market and extending SAN connectivity without distance limitations. It is more cost effective due to its use of existing equipment and infrastructure. With a 10x increase from existing 1Gigabit to 10Gigabit Ethernet, it will become a major force in the SAN market. Using 10Gigabit Ethernet, SANs are reaching the highest storage transportation speeds ever.

iSCSI Typical Component/Elements

- iSCSI Host Bus Adapter (HBA) or NIC (installed in server)
- Storage devices disk arrays or tape libraries
- Servers
- Standard IP Ethernet Switches and Routers
- Storage Switches and Routers
- Gateways
- Physical layer media - Fiber, twisted-pair

Generally, to deploy an iSCSI storage network in a data center, connectivity is provided via an iSCSI Host Bus Adapters (HBAs) or storage NIC which connects the storage resources to existing Ethernet via IP Ethernet switches or IP Storage switches and routers. Specified storage IP routers and switches have a combination of iSCSI interfaces and other storage interfaces such as SCSI or Fibre Channel, they provide multi-protocol connectivity not available in conventional IP and Ethernet switches.

When connecting to FC SANs, an IP storage switch or router is needed to convert the FC protocol to iSCSI. IP storage routers and switches extend the reach of the FC SAN and bridge FC SANs to iSCSI SANs. For example, an IP storage switch allows users to perform FC-to-FC switching, FC-to-iSCSI switching, or FC-to- Ethernet switching in addition to Ethernet to Ethernet switching.

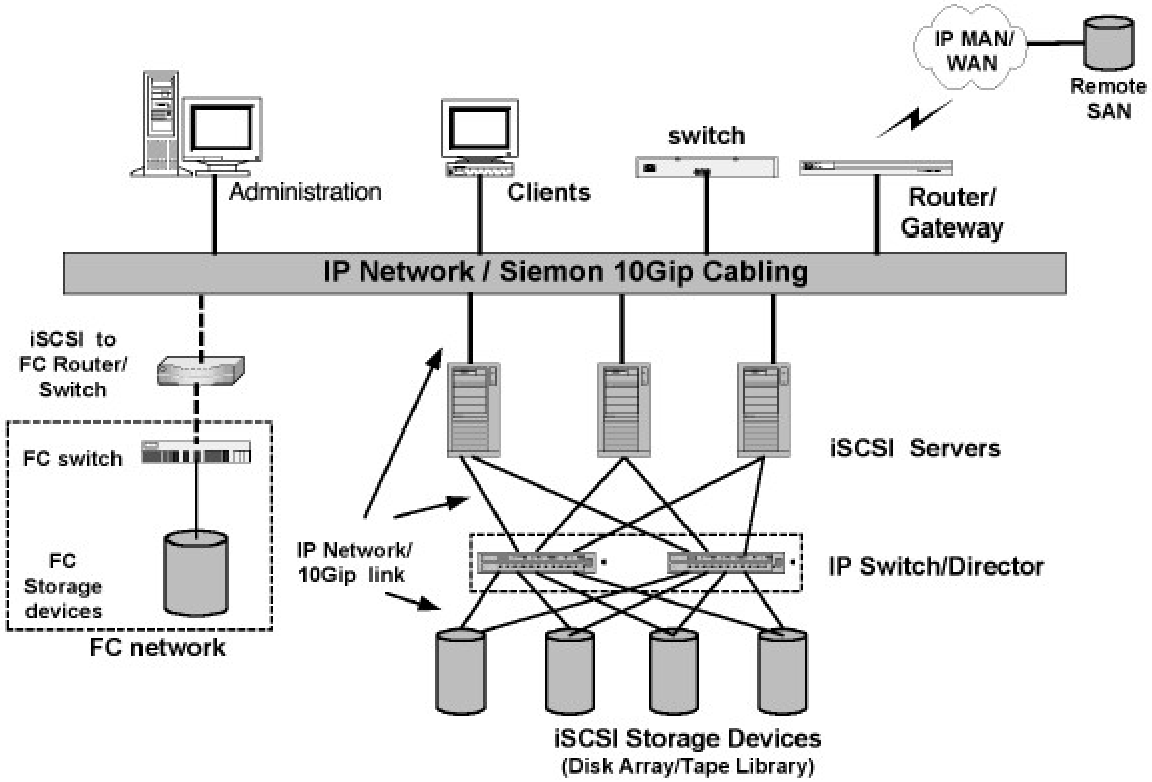

Figure 4: iSCSI SAN Diagram

Figure 4: iSCSI SAN Diagram

Mixed Architectures Storage Networks

Flexibility and low cost are the important driving factors for implementing an iSCSI approach, especially for long distance storage. In addition, as Ethernet speeds are continually increasing, it is believed that the 10 Gigabit Ethernet based iSCSI will be widely used for SANs in data centers today. A number of devices and even iFCP have been developed to address the large installed base of legacy FC storage solutions in place today. In order to protect an organization's current investment in storage technology, SAN installations will likely evolve from a single specific storage network to a mix of Fibre Channel and SCSI products.

Furthermore, a convergence or integration of NAS and SANs is expected and multilingual (combination) Fibre Channel and Ethernet switches are expected to evolve. The integrated SAN and NAS network will be scaleable and cost-effective, it will support multiple protocols and interfaces. According to Dataquest of Gartner, more than 20% of all units shipped will have a combination of iSCSI and Fibre Channel HBAs in 2006. This integration will enable customers to optimize their legacy Fibre Channel SANs by providing reliable connections over long distances using existing electronics by providing a convergence between Fibre Channel and iSCSI protocols.

Evolving Standards for SANs

FC standards are developed by the technical subcommittee NCITS/T11 of the National Committee for Information Technology Standards (NCITS). The original FC standards were approved by the ANSI X.3230 in 1994. The first SCSI standard was ratified by ANSI in 1986. Since then, there have been multiple amendments mirroring changes within the industry.

The Internet Engineering Task Force (IETF) is expanding on these standards through IP protocol enhancements to the existing interface and operational standards above. In February, 2003, the iSCSI specification was officially approved as a "proposed standard" by the IETF. Additionally, the Storage Networking Industry Association (SNIA), the Fibre Channel Industry Association (FCIA), and other industry groups are also working on the SAN standard's implementation and development. The data center is the critical infrastructure hub of an organization. Besides the SANs /NASs components, a typical data center includes a variety of other components and connectivity. To address the evolutions of data centers, the TIA TR-42.1.1 group is working on the "Telecommunications Infrastructure Standard for Data Centers to be published as ANSI/TIA/EIA-942." The standard will cover the cabling system design, pathway, and spaces.

Cabling Considerations and Design Factors for SANs in Data Centers SANS are most prevalent in data centers, but they also include video, voice, and other converged applications. A robust network foundation is essential. In a data center environment the basic requirements for the cabling system are:

- Standards based open system
- High performance & high bandwidth, and quality
- Support for 10Gigabit
- Support for multiple types of SANs / NASs and protocols
- Support for cumulative bandwidth demands for converged applications
- High Reliability
- Redundancy
- Flexible, scaleable and provides mechanisms for easy deployment of MACs

To meet all above requirement, 10G *ip* ™ is the first choice. It is the robust infrastructure for enterprise communications in data centers worldwide and is fully aligned with the draft TIA data center standard. In order to improve the reliability of the communications infrastructure, redundancy is a principal design consideration in a data center. The redundancy can be achieved by providing physically separated services, cross-connected areas and pathways, or by providing redundant electronic devices in fabric topologies.

Note: Siemon 10G *ip* ™ XGLO ™ fiber-optic cabling system can be configured in either a distributed or centralized topology. For detailed product and design information on Siemon 10G *ip* ™ cabling solution, please visit the Siemon website (www.siemon.com) or contact The Siemon Company.

Conclusion

Storage Area Networks are but one component of converged applications that traverse today's networks. The benefits of these systems are not only numerous, but completely essential to a business. Providing the bandwidth necessary for all networked applications with a state of the art infrastructure will assure their functionality for years to come. Adjusting your infrastructure reactively is costly. Industry experts agree that your cabling infrastructure should be planned to carry data for at least 10 years and Siemon's 10G *ip* ™ system can provide a company with the bandwidth needed and protection of their investment for this full 10 years. This includes not only your cabling, but its related connectivity as well.

The Siemon Company dedicates its efforts to assure that our end-users have the best solution for their investment dollars. Backed by one of the industry's best warranties, a global system of certified installers and state of the art components, a company can rest assured that their cabling system will add value and help protect the ROI of their electronics and the applications that run on them. Whether you chose XGLO ™, TERA® or 10G 6 ™, your investment in today will last through tomorrow. Cabling systems are as different as electronics systems, it makes sense to protect your critical assets with the best infrastructure.

Bibliography

Per Hour Downtime Costs, Contingency Planning Research and Internetweek (4/3/2000)

- Worldwide Disk Storage Systems Report, IDC, www.idc.com
- SAN for the Masses, Computing Technology Industry Association, http://www.comptia.org/research/
- Storage Network Infrastructure, 2003 Forecast (Executive Summary), Dataquest of Gartner, www.gartner.com
- ANSI, American National Standards Institute, www.ansi.org
- TIA, Telecommunications Industry Association, www.tiaonline.org
- EIA, Electronics Industry Alliance, www.eia.org
- IETF, Internet Engineering Task Force, www.ietf.org
- SNIA, Storage Networking Industry Association, www.snia.org
- FCIA, Fibre Channel Industry Association, www.fibrechannel.org

---

The Siemon Company
Siemon Business Park
101 Siemon Company Drive
Watertown, CT 06795-0400 USA
Phone:(1) 860.945.4200
Fax: (1) 860.945.4225