



Data Center Storage Evolution Executive Summary

Data is growing at explosive rates in today's businesses. Big Data is increasing storage demands in a way that could only be imagined just a few short years ago. A typical data record has tripled if not quadrupled in size in just the last five years, however this data now has many forms including structured, semi-structured and non-structured. In fact, according to a recent IBM® study, 2.5 quintillion bytes of data are written every day and 90% of global data has been created in the last two years alone. It is glaringly apparent that the size of 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 in the data center. The SANs operate behind the servers to provide a common path between servers and storage devices. Unlike server-based Direct 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 including iSCSI and Fibre Channel over Ethernet (FCoE). The media on which the data is stored is also changing.

With the growth of SANs and the worldwide domination of Internet Protocol (IP), 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 FCoE (Fibre Channel over Ethernet), and iSCSI (ip Small Computer System Interface). Both carry either Fibre Channel or SCSI commands incorporated into an IP datagram. FCoE is different in that Fibre Channel commands are encapsulated into IP traffic, but this requires a converged network adapter (CNA) that is capable of speaking both Fibre Channel and Ethernet for encapsulation. iSCSI operates over standard Ethernet networks and standard Ethernet adapters at the edge device called the initiator.

Today, 10Gigabit Ethernet is becoming increasingly popular as the horizontal application of choice in corporate data centers. Gaining a competitive edge from deploying 10 Gigabit Ethernet in the enterprise requires a robust IT infrastructure. Increasingly, 10GBASE-T and 10Gb SFP+ applications provide a reliable foundation for data centers' networking components and SAN networking. With a structured cabling

system capable of 10GBASE-T, users are provided with an open and industry standards-based infrastructure that can support multiple converged applications.

This paper provides some useful insight into both existing and new storage.

Storage Technologies

With the advent of the Internet, Big Data, 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 can simply not address the growing demands. Growing data stores meant having to implement tape libraries. Even then, 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 a recently published study by Dun & Bradstreet, 59% of Fortune 500 companies experience a minimum of 1.6 hours of downtime per week. Wages alone levy a downtime cost of \$896,000 per week or just over \$46 million per year. A recent conservative Gartner study lists downtime costs at \$42,000 per hour. A USA today survey of 200 data center managers found that over 80% reported that their downtime costs exceed \$50,000 per hour, and another 20% said they exceed \$500,000 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 due to compliance requirements.

Storage network technology has developed in the following three main configurations: Direct Attached Storage (DAS), Network Attached Storage (NAS), and Storage Area Networks (SAN).

Direct Attached Storage (DAS)

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 dedicated path separate from the network cabling. 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. The mirroring functionality may also be accomplished via directly attached server to server interfaces.




Figure 1: A Simple DAS Diagram

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 as shown in 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 which may be incorporated into the unit itself. 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: Simple NAS Architecture




Figure 3: Meshed SAN Architecture

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. 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. As shown in Figure 3 the term switched fabric refers to each storage unit being connected to each server via multiple SAN switches also called SAN directors which provide redundancy within the paths to the storage units. This provides additional paths for communications and eliminates one central switch as a single point of failure.

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 the significant number of Gigabit and 10 Gigabit Ethernet ports shipped, leveraging IP and Ethernet for storage is a natural progression for some environments.

SAN over IP

IP was developed as an open standard with complete interoperability of components. Two new IP storage network technologies are Fibre Channel over Ethernet (FCoE) and SCSI over IP (iSCSI). IP communication across a standard IP network via Fibre Channel Tunneling or storage tunneling has the benefit of utilizing storage in locations that may exceed the directly attached limit of nearly 10 km when using fiber 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. 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 (FC) and Fibre Channel over Ethernet (FCoE)

Native 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: 1, 2, 4, 8, 16 and 32 Gb/s with 8Gb/s to 16 Gb/s currently being predominant. The transmission distances vary with the speed and media. With FCoE, the packets are processed with the lengths and distances afforded by an Ethernet Network and again, vary according to speed and media. According to the IEEE 802.3ae standard for 10Gigabit Ethernet over fiber, when using singlemode optical fiber cables, the distance supported is 10 kilometers, up to 300m when using laser optimized 50 micron OM3 multimode fiber and up to 400m with OM4 as compared to native Fibre Channel with a distance of only 130m. Laser optimized OM3 and OM4 fiber is an important consideration in fiber selection for 10Gb/s transmission.

FC Topology

Native 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.

FCoE Topology

Fibre Channel is accomplished on a separate network than the Ethernet network. With Fibre Channel over Ethernet, Converged Network Adapters are used in place of Ethernet adapters and allow a single channel to pass both Ethernet and Fibre Channel encapsulated packets across a standard IP network extending distance over an entire enterprise, regardless of geography via Ethernet routers and bridges. For replication between storage systems over a wide area network, FCoE provides a mechanism to interconnect islands of FC SAN or FCoE SANs over the IP infrastructure (LANs/MANs/WANs) to form a single, unified FC SAN fabric.

Native Fibre Channel SAN 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, however fiber is the most predominant.

The FC SAN switches are classified as either switches or directors. A SAN 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 via the networking cabling subsystem. A bridge may be used to connect legacy SCSI or ESCON (Enterprise System Connection) storage devices to the Fibre Channel network. The bridge will serve to translate and/or encapsulate the various protocols allowing communication with legacy storage devices via the SAN.

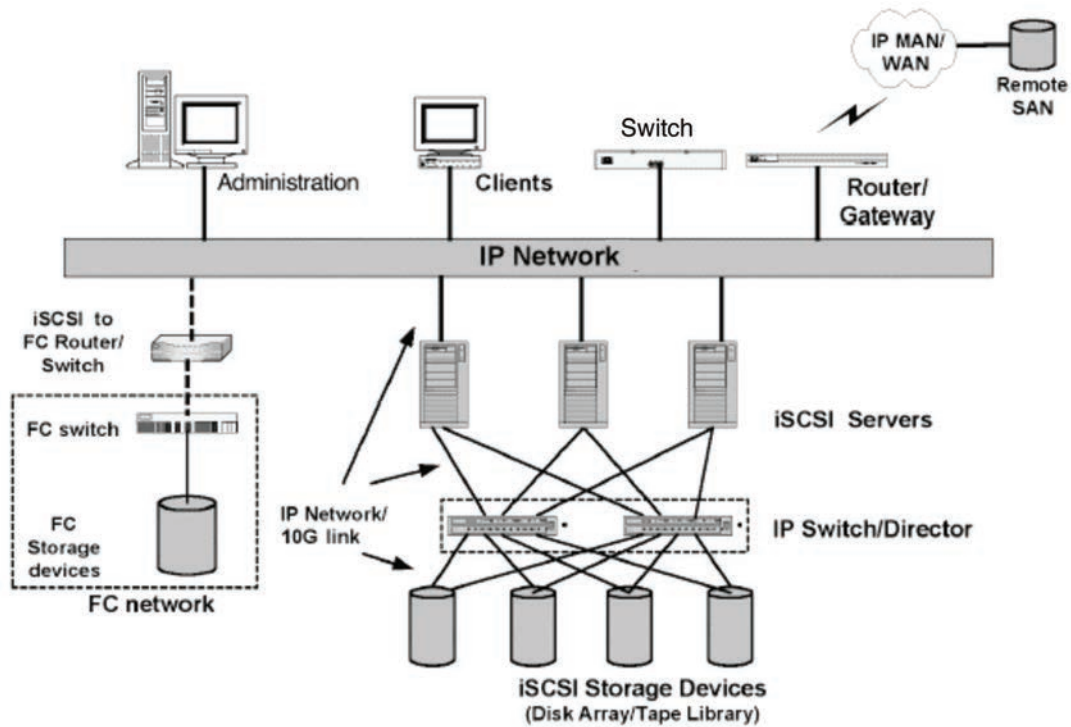


Figure 4: iSCSI SAN Diagram

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, long distance 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 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.

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. A number of devices have been developed to address the large installed base of native FC storage solutions in place today. In order to protect an organization's current investment in storage technology, SAN installations may evolve from a single specific storage network to a mix of Fibre Channel and iSCSI 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. This integration will enable customers to optimize their native Fibre Channel SANs by providing reliable connections over long distances using existing electronics by providing a convergence between Ethernet, 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 SAN /NAS 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 developed the "Telecommunications Infrastructure Standard for Data Centers" published as ANSI/TIA/EIA-942 and later amended and published as TIA 942-A. The standard covers the cabling system design, pathway, and spaces. Likewise, ISO developed ISO 24764 international cabling standard for data centers.

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

- Standards-based open system
- Support for 10GbE, 8, 16 and 32Gb/s FC
- Support for multiple types of SAN / NAS and protocols
- Support for cumulative bandwidth demands for converged applications
- High Reliability
- Redundancy
- Flexible, scaleable and provides mechanisms for easy deployment of MACs
- It is highly desirable to use the highest performing fiber with low loss connectors to allow reconfigurations without running new fiber.

To meet all above requirement, 10GbE copper and laser optimized multimode fiber are the first choices. TIA recommends category 6A as a minimum copper cabling standard and now OM4 as the minimum fiber standard. ISO 24764 recommends 6A as a minimum for copper and OM3 for fiber. A 10GbE capable infrastructure is predominant in data centers today, with 40 and 100GbE fast approaching for backbone applications. 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.

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 using a high performance structured cabling infrastructure will ensure their functionality for years to come. Upgrading or replacing your infrastructure reactively is costly. Industry experts agree that cabling infrastructure should be planned to carry data for at least 10 years.

Storage solutions are plentiful and there is no one size fits all for today's data centers. In fact some data centers utilize a variety of storage architectures depending on the application requirements. While Fibre Channel in native form is the predominant architecture for storage, iSCSI and FCoE are gaining some momentum. When fibre channel SANs complement Ethernet networks, dual paths for moving data are provided. Converging fibre channel over Ethernet decreases the number of connections required, but doubles the traffic over the used channels. Increasing bandwidth from gigabit to 10GbE provides more bandwidth for these applications. When increasing the horizontal server to switch speed, uplink ports also need to increase in speed, generally using multiple 10GbE links or newer 40/100GbE speeds. Siemon's data center design assistance experts can help design a storage and network architecture to support your business needs.

The Siemon Company is a global market leader specializing in high performance, high quality cabling systems. Siemon offers a broad range of copper and fiber cable, connectivity and cable management systems for Data Centers including Storage Area Networks and beyond. Siemon's LightStack™ Fiber Plug and Play system combines superior performance with ultra high density (144 LC and 864 MTP fibers in 1U) and best in class accessibility. Siemon cabling systems are backed by an extended warranty covering product quality, performance headroom and applications assurance for up to 20 years. For more information on Siemon Data Center solutions please visit: www.siemon.com/datacenter.

Bibliography

- 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

Worldwide Headquarters North America

Watertown, CT USA
Phone (1) 860 945 4200 US
Phone (1) 888 425 6165

Regional Headquarters EMEA

Europe/Middle East/Africa
Surrey, England
Phone (44) 0 1932 571771

Regional Headquarters Asia/Pacific

Shanghai, P.R. China Phone
(86) 21 5385 0303

Regional Headquarters Latin America

Bogota, Colombia
Phone (571) 657 1950