Internetworking the Storage Area Networks

STANISLAV MILANOVIC [*], NIKOS E. MASTORAKIS [#]

[*] WSEAS, Highest Institute of Education, Science and Technology
Haghiou I. Theologou 17
15773, Zographou,
Athens, GREECE

[#] MILITARY INSTITUTIONS OF UNIVERSITY EDUCATION
HELLENIC NAVAL ACADEMY
Terma Hatzikyriakou, 18539, Piraeus, GREECE.

Abstract: - A SAN internetworking has been highlighted in this paper describing new technologies available for building an enterprise-wide SAN and connecting Fibre Channel SANs over Wide Area Network. The new storage-centric infrastructure includes an open, modular and scalable storage network, not tied to any one server or application.

Key-Words: - Storage Area Network, Fibre Channel, e-business, WAN, Internet, ATM, Gigabit Ethernet

1 Introduction

Companies across all industries are launching new business-critical applications such as enterprise resource planning, business intelligence, customer relationship management and supply chain reengineering — turning information into a strategic corporate asset. With information taking center stage as a competitive differentiator, storage has moved from supporting player to a leading role in a company's IT strategy and budget.

To remain successful in such a dynamic marketplace, organizations need reliable storage systems that can effectively manage and protect critical business information. These systems must be able to scale quickly to manage anticipated data growth — a difficult problem for many traditional storage approaches. As a result, organizations are now accessing and managing the ever-increasing amount of enterprise data through innovative Storage Area Networks (SANs).

2 SAN Technology Overview

SAN is a separate, centrally managed (but functionally distributed) networked environment that provides a scalable, reliable IT infrastructure to meet the high-availability, high-performance requirements of today's most demanding e-business applications. SAN is focused on the single task of managing storage resources and removing that task from the LAN or servers. SANs accomplish this by providing any-server-to-any-storage connectivity through the use of Fibre Channel switch fabric technology. From a client network perspective, the SAN environment complements the ongoing advancements in LAN (Local Area Network) and WAN (Wide Area Network) technologies [1, 2, 3, 4, 5, 6, 7, 8, 9] by extending the benefits of improved performance and capabilities all the way from the client and backbone through to servers and storage. SANs are currently used to connect shared storage arrays, cluster servers for failover, interconnect mainframe disk or tape resources to distributed network servers and clients, and create parallel or alternate data paths for high performance computing environments [10]. Most SAN solution vendors also mandate SNMP (Simple Network Management Protocol) support for all the SAN networking components, which is then used and exploited by SNMP based network management applications to report on the various disciplines of SAN management [11].

The central foundation of the SAN is Fibre Channel technology (FC). The first generation SAN provides 1Gb/s (2Gb/s full-duplex), while the second generation SAN supports multiple data rates up to 2Gb/s (4 Gb/s full-duplex) [12]. The FC specification is a set of standards being developed by ANSI (American National Standards Institute) and is ideal for storage, video, graphic and mass data transfer applications [13].

FC is a layer 2 technology which operates over copper and fiber optic cabling with maximum distances appropriate to the media (30m for copper, 500m for short-wave laser over multimode fiber, 10Km for long-wave laser over single mode fiber) [14]. FC supports protocols such as SCSI (SCSI over FC is called FCP), ESCON (Enterprise Systems Connection), FICON (Fiber Connectivity), SSA (Serial Storage Architecture), IPI (Intelligent Peripheral Interface), HiPPI (High Performance Parallel Interface), ATM, IP [15]. FC is designed to transport large blocks of data with greater efficiency and reliability than IP-based networks, which can significantly improve backup and recovery performance. FC supports three topologies: point-to-point, arbitrated loop (FC-AL), and switched — FC-SW (1Gb/s, i.e. 2Gb/s full duplex) or FC-SW2 (2Gb/s, i.e. 4Gb/s full duplex); all three topologies are fully interoperable, so the topology is transparent to the attached devices. The principal change in this new architecture is the externalization of server storage onto the SAN as a shared resource attached to multiple servers [16] (Figure 1). It does not matter whether the attached storage resources are mainframe DASD (direct access storage device), open systems disk arrays, or even remote storage used for a variety of reasons. They are all connected by a SAN utilizing high-speed storage interconnect technologies.

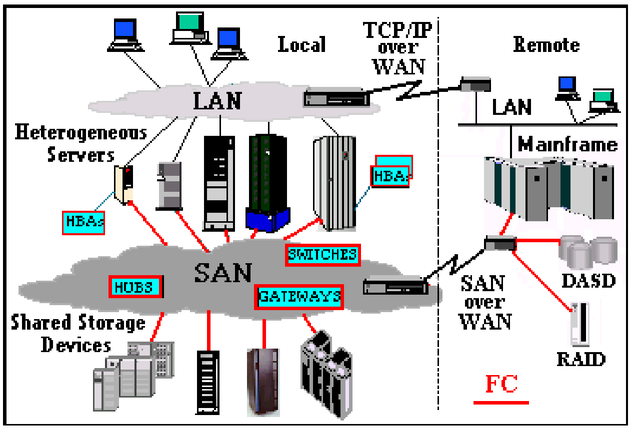

Figure 1. The SAN

The LAN meets the SAN at the server. The server is provisioned with both a LAN interface card and a FC Host Bus Adapter (HBA). As a user requests data from the server over the LAN, the server retrieves the data from storage over the SAN [17].

3 Objective

The commitment was to providing an open, vendor-neutral SAN solution to meet the following objectives:

- Enhanced performance

Take storage traffic off the enterprise LAN and enterprise servers. The SAN enables bulk data transfer from each server to shared SAN storage, but the LAN is used only for communication (not data) traffic between the servers. Sophisticated backup and recovery software applications still control the process, tracking the backup and recovery data. The result is a faster, more scalable, and more reliable backup and recovery solution — with more effective utilization of storage, server, and LAN resources [18].

- Mission continuity
  - Consolidate backups and archives
  - Disk mirroring to disaster recovery sites
- Applications
  - High availability mission critical databases
  - Distributed server clustering
  - Disk virtualisation

4 Requirements

Since standards from SNIA (Storage Network Industry Alliance) and the FCIA (Fibre Channel Industry Alliance) are just now coming into being, the major requirement was to deploy FC switches whose firmwares were upgradeable to newer versions as new standards become available.

Specific requirements that were driving towards storage area networking solution include:

- Scalability,
- High Speed Storage Access,
- Heterogeneous connectivity,
- Flexibility in Server and Storage Placement,
- Secure transactions and data transfer,
- 24x7 response and availability.

5 SAN Internetworking

Following are network components that had been employed to build and interconnect an enterprise-wide SAN infrastructure:

- SilkWorm 3800 switch (SW3800) provides 16 ports with auto-sensing 1Gbit/s (2Gb/s full duplex) and 2Gbit/s (4Gb/s full duplex) interfaces for seamless integration with existing FC fabrics. SW3800 includes Brocade Advanced Fabric Services that increase security through hardware-enforced World Wide Name zoning [19]. SW3800 switches were deployed with fiber optic cabling to support all FC-related topologies providing reliable, high-performance data transfer that is critical to efficient SAN applications, such as

LAN-free backup, server-free backup, storage consolidation, remote mirroring and high-availability clustering configurations. Inter-switched links (ISLs) between two central SW3800 switches (Figure 2) create a single logical high speed trunk running up to 4Gb/s (8Gb/s full duplex). The trunked ISL is fault-tolerant in that it will withstand failure of individual links. This feature will improve core fabric throughput and performance [20].

- The ATTO FibreBridge 4500C/R/D Fibre Channel-to-SCSI bridge is configured with three independent Fibre Channel ports and four independent Ultra2 SCSI buses [21].
- Compaq ProLiant (Windows NT/2000, Linux), IBM RS6000 (AIX), HP9000 (HP-UX), Sun Enterprise (Solaris) servers; IBM Enterprise Storage Server 2105 F20 RAID (Redundant Array of Independent Disks), Nextor 18F (FC target JBOD - Just a Bunch of Disks); Terasystem DataSTORE LTO TLS 8000 (FC or SCSI tape library); Terasystem OptiNET CD-ROM/DVD Jukebox and CD-ROM/DVD Jukebox Controller; Adaptek 2940U2W (Ultra2 SCSI adapters); FC 6227, Qlogic 2100 and 2200 series Fibre Channel HBAs.
- Tivoli Storage Manager (TSM) server running on AIX. TSM delivers data protection for file and application data, record retention, space management, and disaster recovery. The client software (Storage Agent) running on different systems (PCs, workstations or application servers), in conjunction with the server software (TSM) enables the LAN-free data transfer exploiting SAN infrastructure [22].
- The ADVA DiskLink SAN Interconnect Gateway enables the distribution and mirroring of data across multiple sites. DiskLink is a high-performance storage networking device that provides 2Gb/s Fibre Channel (4Gb/s full duplex) and Ultra 2 SCSI networking over unlimited distances across ATM networks at 155/622 Mbps, or over next-generation Internet services (Virtual Private Networks) with Gigabit Ethernet. Configurable levels of intelligent storage command handling allow DiskLink to reduce the impact of transmission delays and eliminate distance barriers [23].

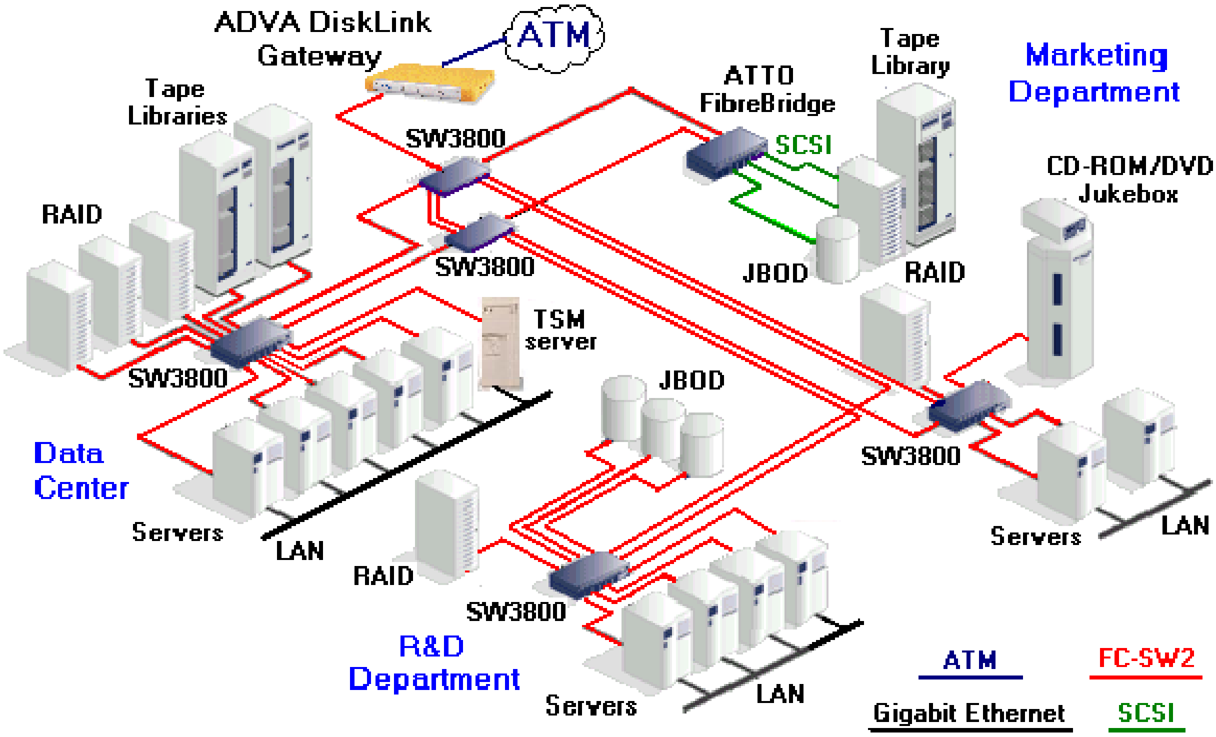

Figure 2. Local SAN island deployment at Greek corporate headquarter

Native FC technology provides the 10 km or greater link extended distance connectivity required to maintain geographically separate disaster recovery facilities or mirroring operations. Nevertheless, remote switches can provide virtually unlimited distance connections through the ATM protocol over existing WAN [24]. This SAN interconnectivity allows a prominent communications provider to implement remote backup and disaster recovery plans while utilizing the existing storage facilities and network infrastructure [25]. Figure 3 depicts deployed enterprise-wide SAN solution spanning European sites of a well-known international organization. In this scenario, the storage devices appear logically as if they were on the same bus or loop, but gain the advantages of a switched infrastructure. The ATM network sets up direct paths between devices and becomes a Virtual SAN switch — a Virtual SAN (V-SAN). The V-SAN is distance independent while storage and servers can be located anywhere [26]. The V-SAN facilitates remote archiving, disaster recovery, web caching and data warehousing (multi-site, central storage).

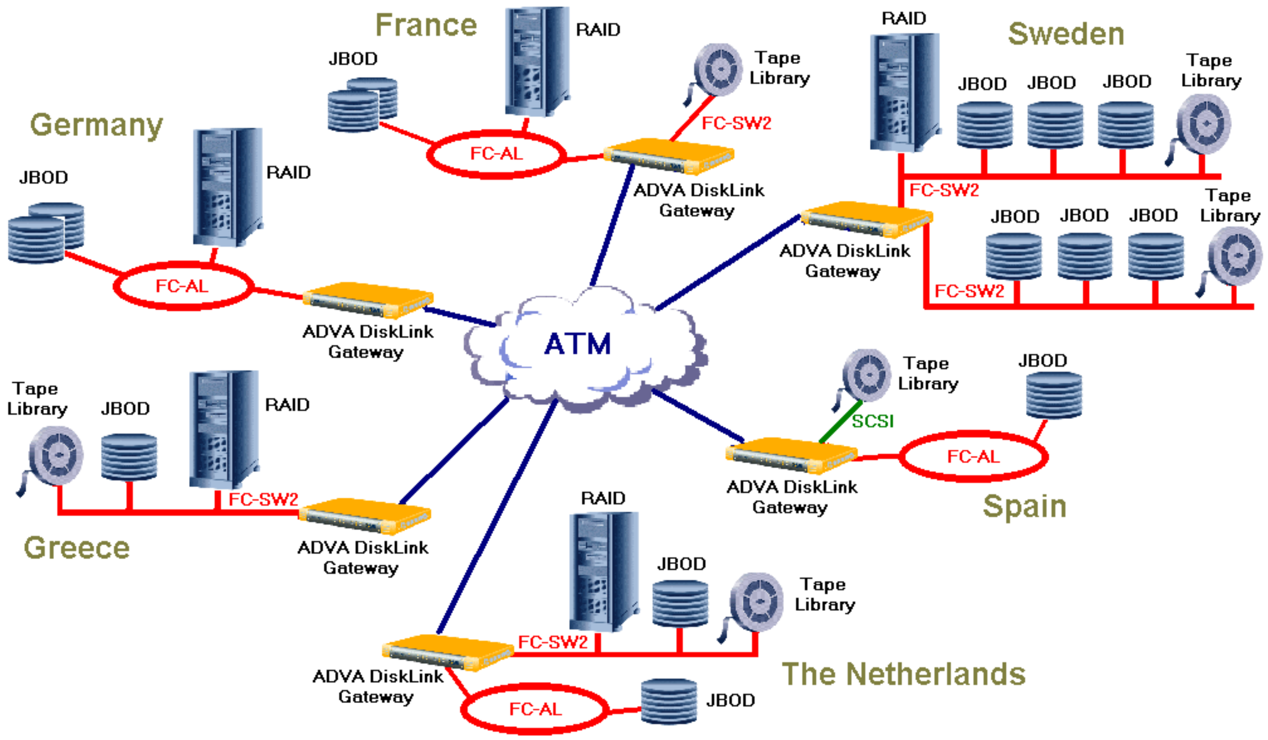

Figure 3. Internetworked Storage Area Networks

6 Conclusion

In addition to the fundamental connectivity benefits, deployed enterprise-wide SAN also features the new capabilities, facilitated by its networking approach, which enhances SAN's value as a long term infrastructure. These capabilities, which include compute clustering, topological flexibility, fault tolerance, high availability and remote management, further elevate deployed SAN's ability to address the growing challenges of data-intensive, mission-critical e-business applications.

As Storage Area Network infrastructures continue to expand, so does the need to connect storage devices over longer distances in heterogeneous environments [27]. In fact, many organizations are beginning to connect local SAN islands over existing high-speed public and private networks — an approach that enables new types of applications that leverage a geographically dispersed, yet interconnected SAN infrastructure.

SAN is certainly a technology that has proven useful, but remains relatively new and the market is still in flux. The technology is still relatively immature and there is also a lack of accepted industrial standards [28]. Despite the muddle, SANs generate considerable interest in the computer industry because, in this age of explosive growth in storage needs and the resulting complexity in managing storage, SANs provide a viable solution, managing a company's current storage needs, as well as its future growth.

References:

[1] Stanislav Milanovic, Nikos E. Mastorakis, "Architecting the Next Generation End-to-End e-Business Trust Infrastructure", Proceedings of WSEAS CSCC 2002, pp. CD-ROM, July 7-14, 2002, Crete, Greece, http://www.wseas.org/New_Books.htm

[2] Stanislav Milanovic, Nikos E. Mastorakis, "Delivering Enhanced Voice Services over the Internet", WSEAS TRANSACTIONS on SYSTEMS, Issue 1, Volume 1, pp. 74-80, January 2002, WSEAS Press,http://www.wseas.org/SYSTEMSJANUARY2002.htm .Proceedings of WSEAS MMACTEE 2001, pp. CD-ROM, December 29-31, 2001, Vouliagmeni, Athens, Greece, http://www.wseas.org/History.htm

[3] Stanislav Milanovic, Nikos E. Mastorakis, "A Transition Path to Gigabit Ethernet over WDM in Support of Emerging e-Business Applications", WSEAS TRANSACTIONS on SYSTEMS, Issue 1, Volume 1, pp. 80-87, January 2002, WSEAS

Press,http://www.wseas.org/SYSTEMSJANUARY2002.htm.Proceedings of WSEAS MMACTEE 2001, pp. CD-ROM, December 29-31, 2001, Vouliagmeni, Athens, Greece, http://www.wseas.org/History.htm
[4] Stanislav Milanovic, Zoran Petrovic, “Securing the Networked e-Business Throughout an Internet Distributed Organization”, Advances in Intelligent Systems, Fuzzy Systems, Evolutionary Computation, pp. 180-186, WSEAS Press, February 2002, http://www.wseas.org/New_Books.htm.Proceedings of the WSEAS EC’02, pp.CD-ROM, February 11-15,2002, Interlaken,Switzerland,http://www.wseas.org/conferences/2002/interlaken/program.pdf
[5] Stanislav Milanovic, Zoran Petrovic, “Deploying IP-based Virtual Private Network across the Global Corporation”,Communications World, pp. 13-17, July 2001,WSESPress,http://www.wseas.org/8052386.doc. Proceedings of WSES/IEEE CSCC 2001, pp. CD-ROM, July 8-15, 2001, Crete, Greece, http://www.wseas.org/New_Books.htm
[6] Stanislav Milanovic, Zoran Petrovic, "A Practical Solution for Delivering Voice over IP", Proceedings of IEEE/IEE/WSES ICN’01, July 9-13, 2001,Colmar,France,http://iutsun1.colmar.uha.fr/pgmICN01.html, Lecture Notes in Computer Science (LNCS #2094), Part II, pp. 717-725, July2001, Springer-Verlag GmbH & Co. KG, http://link.springer.de/link/service/series/0558/tocs/t2094.htm
[7] Stanislav Milanovic, Zoran Petrovic, "Building the Enterprise-wide Storage Area Network". Proceedings of IEEE EUROCON 2001, Vol.1, pp. 136-139, July 5-7, 2001, Bratislava, Slovakia,
http://www.ktl.elf.stuba.sk/EUROCON/program.htm
[8] Stanislav Milanovic, “At the Front End in Migrating to Gigabit Ethernet”, Proceedings of IEEE SoftCOM 2000, Vol.1, pp. 369-378, October 10-14, 2000, Split, Rijeka (Croatia), Trieste, Venice (Italy), http://www.fesb.hr/SoftCOM/2000/IE/Network_Architectures.htm
[9] Stanislav Milanovic, Alessandro Maglianella, "ATM over ADSL Probe in Telecom Italia Environment", Computer Networks, Vol. 34, No. 6, pp. 965-980, December 2000, Elsevier Science, http://www.elsevier.nl/gej-ng/10/15/22/49/32/show/toc.htt.Proceedings of TERENA Networking Conference 2000, pp. CD-ROM, May 2000, Lisbon, Portugal, http://www.terena.nl/tnc2000/proceedings/10A/10a3.pdf
[10] Tom Clark, “SANs for the Internet”, White Paper, Vixel Corporation, 2000.
[11] Ravi Kumar Khattar, Mark S. Murphy, Giulio John Tarella, Kjell E. Nystrom, “Introduction to Storage Area Network, SAN”, White Paper, IBM Corp. 1999.
[12] “Network Technologies Investigation - NASA/GSFC High Speed Fiber Optics Test Bed”, Advanced Component Technology Group, Swales Aerospace, The Thomas Consultancy Group, 4 October 1998.
[13] Fibre Channel Physical and Signaling Interface (FC-PH), ANSI X3.230-1994, Rev. 4.3, 1994.
[14] “Designing Storage Networks with Fibre Channel Switches, Switching Hubs, and Hubs”, White Paper, Vixel Corporation, 2000.
[15] Steve Nowack, “Storage Networking with Fibre Channel Connectivity: Fibre Channel Basics”, White Paper, StorageTek, June 1998.
[16] “Storage Networking, The Next Network”, White Paper, Strategic Research Corporation, 2000.
[17] Tom Clark, “Why SANs”, White Paper, China Computerworld, March 2000, Vixel Corporation.
[18] “LAN-Free and Server-Free Backup Overview”, Brocade Communications Systems, Inc, 2001.
[19] SilkWorm 3800, Data Sheet, Brocade Communications Systems, Inc., 2001.
[20] David Merrill, “End-to-End 2 Gigabit per Second Fibre Channel SAN”, White Paper, Hitachi Data Systems Corporation, October 22, 2001
[21] ATTO FibreBridge 4500C/R/D, Technical Sheet, 2000 ATTO Technology, Inc.
[22] Roland Leins, Lloyd Dieter, Wim Feyants , Leos Stehlik, “Using Tivoli Storage Manager in a SAN Environment”, IBM Corporation, International Technical Support Organization, IBM Redbooks, December 2000.
[23] "SAN Interconnect Gateway DiskLink", Data Sheet, ADVA Optical Networking, 2000.
[24] “Brocade SAN Solutions: A More Effective Approach to Infirmation Storage and Management”, White Paper, Brocade Communications Systems, Inc., 2000.
[25] “The Performance of SAN traffic over Wide Area Networks”, White Paper, ADVA Limited, June, 2000.
[26] Andrew Rowney, "Storage at a distance — the Virtual SAN", ADVA Limited, June, 2000.
[27] “Connecting SANs over Metropolitan and Wide Area Networks”, White Paper, Brocade Communications Systems, Inc., 2001.
[28] Alan S. Horowitz, “The Basics Of Storage Area Networks - Have We Reached the Time of SANs?”, CMP Media Inc, 2000.