Storage Area Networking

The Network Behind the Server

Dave Tang, Gadzoox Microsystems

The rapid growth in data intensive applications continues to fuel the demand for raw data storage capacity. Applications such as data warehousing, data mining, on-line transaction processing, and multimedia Internet and intranet browsing are resulting in the near doubling of the total storage capacity shipped on an annual basis[i]. Further fueling the demand for network storage, analysts have predicted that the number of network connections for server-storage subsystems will exceed the number of client connections by 1999.

The Problem: Limitations Loom Over Surge of Data

With the rise of client networking, data-centric computing applications and electronic communications applications, virtually all network stored data has become mission-critical in nature. This increasing reliance on the access to enterprise data is challenging the limitations of traditional server-storage solutions. As a result, the ongoing need to add more storage, service more users and back-up more data has become a monumental task.

Having endured for nearly two decades, the parallel Small Computer System Interface (SCSI) bus that has facilitated server-storage connectivity for Local Area Network (LAN) servers is imposing severe limits on network storage. Compounding these limits is the traditional use of LAN connections for server-storage backup which detracts from usable client bandwidth. To contend with these limitations, network managers are often forced to compromise on critical aspects of system availability, reliability and efficiency. To address the debilitating and potentially costly effects of these constraints, an infrastructure for server-storage connectivity which can support current and future demands is badly needed.

Impending Limitations of Existing Network Server Connectivity

- Bandwidth to service clients and maintain data availability
- Scalability for long term, rapid growth
- Flexibility to provide optimum balance of server and storage capacity
- Manageability for ease of installation and maintainability

The Solution: Storage Area Networking

The Storage Area Network (SAN) is an emerging data communications platform which interconnects servers and storage at Gigabaud speeds. By combining LAN networking models with the core building blocks of server performance and mass storage capacity, SAN eliminates the bandwidth bottlenecks and scalability limitations imposed by previous SCSI bus-based architectures.

The Storage Area Network Environment

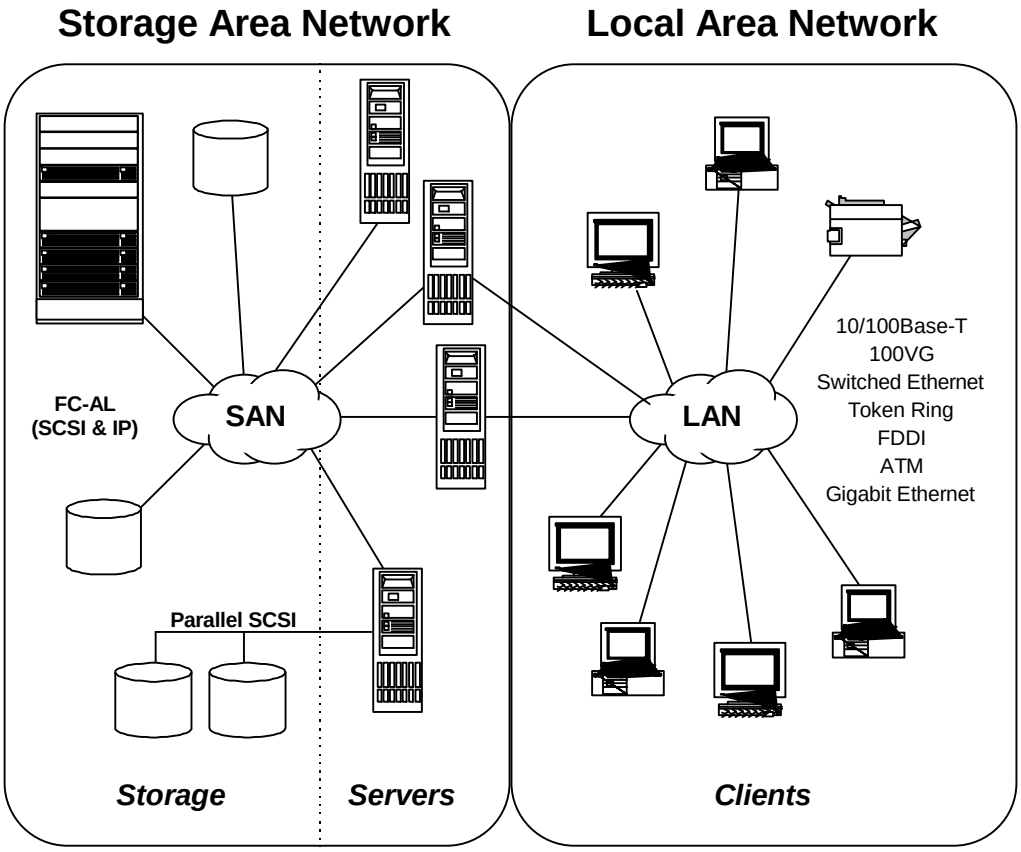

In addition to the fundamental connectivity benefits of SAN, the new capabilities, facilitated by its networking approach, enhance its value as a long term infrastructure. These capabilities, which include compute clustering, topological flexibility, fault tolerance, high availability, and remote management, further elevate SAN's ability to address the growing challenges of data-intensive, mission-critical applications. From a client network perspective, the SAN environment complements the ongoing advancements in LAN and WAN technologies by extending the benefits of improved performance and capabilities all the way from the client and backbone through to servers and storage.

Benefits of the Storage Area Network Environment

- High Bandwidth
- Modular Scalability
- High Availability & Fault Tolerance
- Manageability
- Ease of Integration
- Total Cost of Ownership

Fibre Channel: The Open SAN Solution

Over the past year, Fibre Channel-Arbitrated Loop (FC-AL) has emerged as the high-speed, serial technology of choice for server-storage connectivity. With over 70 companies, including industry leading disk drive, disk array, server and networking connectivity suppliers, supporting FC-AL, it has become the most widely endorsed open standard for the SAN environment. This broad acceptance is attributed not only to FC-AL's high bandwidth and high scalability but also to it's unique ability to support multiple protocols, such as SCSI and IP, over a single physical connection. This enables the SAN infrastructure to serve as both a server interconnect and as a direct interface to storage devices and storage arrays.

Complementary SAN, LAN and WAN Technologies

Technology	Current Bandwidth	Future Bandwidth	Applications
ATM	622 Mbps	1+ Gbps	LAN and WAN
Ethernet	100 Mbps	1 Gbps	LAN and WAN
FC-AL	1 Gbps	4 Gbps	SAN

High Bandwidth

FC-AL provides a 2.5 to 10 fold increase in effective data bandwidth over the traditional parallel SCSI storage interface. Additionally, FC-AL offers future expandability. While the current FC-AL standard for bandwidth is 1 Gigabaud, planned enhancements to 2 and 4 Gigabaud give FC-AL a solid platform to address longer term bandwidth requirements.

Comparison of FC-AL and UltraSCSI

Attribute	UltraSCSI Limit*	FC-AL SAN
Data Transmission	Half-Duplex	Full-Duplex
Effective Data Bandwidth**	40 MB/sec	200 MB/sec
Protocol Support	SCSI	SCSI, IP, others
Connection Scalability	15 drives per bus	126 nodes per loop
Connection Distance	25m	10km
Relative Storage Capacity***	136 Gbytes	9,172 Gbytes

* Based on wide differential UltraSCSI parallel interface standards

** Under FC-AL's 8b/10b encoding, 1 Gbaud = 100 MB/sec data rate

*** Based on 9.1GB disk drives and one, eight drive disk array per node for SAN

Server and Storage Scalability

The modular scalability of FC-AL is key to enabling an infrastructure for long term growth and manageability. Traditional parallel SCSI bus connections have been limited to a total of 7 or 15 storage devices. As bus bandwidth is pushed further and further this limit is compressed to even fewer devices per bus. In contrast, FC-AL supports up to 126 nodes per loop with a typical configuration consisting of a combination of servers and multi-disk arrays per node. By adding multiple loops, the overall scalability is limitless.

Scalability in terms of capacity management and capacity balancing is an area of significant differentiation between FC-AL and SCSI. Largely dictated by the limits on physical cable length, parallel SCSI storage connectivity requires close proximity to its host system, typically a server. This translates to a single, integrated server-storage enclosure that contains both server processing power and one or two SCSI buses of limited scalability.

Traditional Integrated Server-Storage Model

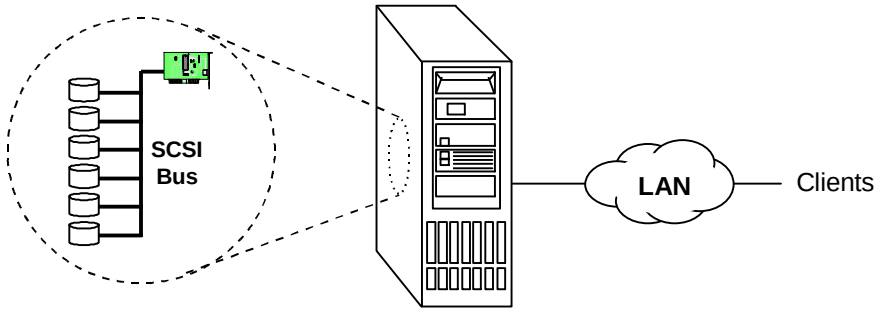

Under this single server-storage enclosure model, the scaling of server capacity and storage capacity becomes inflexible and inefficient. Single enclosures typically hold only 4-10 drives. In order to scale the storage capacity beyond this limit, additional server-storage enclosures, including the cost of the server processor board and peripherals, is required. The inter-dependence between the server and storage capacity in this single enclosure model leads to inefficient and costly scaling. With a diverse combination of data-intensive applications and server processing-intensive applications running concurrently in the enterprise, the need for more flexible and efficient scaling is needed.

Inter-Dependent Capacity Scaling with Integrated Server-Storage Model

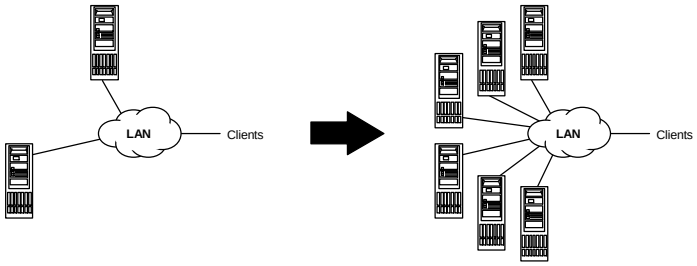

With less stringent cable length limitations, FC-AL enables the networking of separate server and storage enclosures within the SAN environment.

Server-Storage Networking Model of SAN

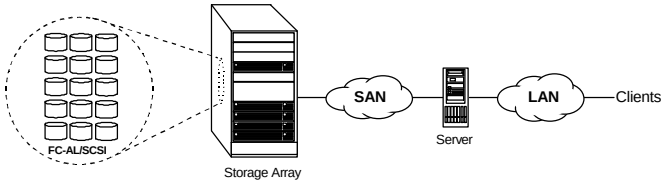

This capability provides a more flexible and cost-effective path for the independent scaling of server performance and storage capacity where either may be expanded independently to achieve an optimum balance.

Independent Capacity Scaling with Server-Storage Networking Model

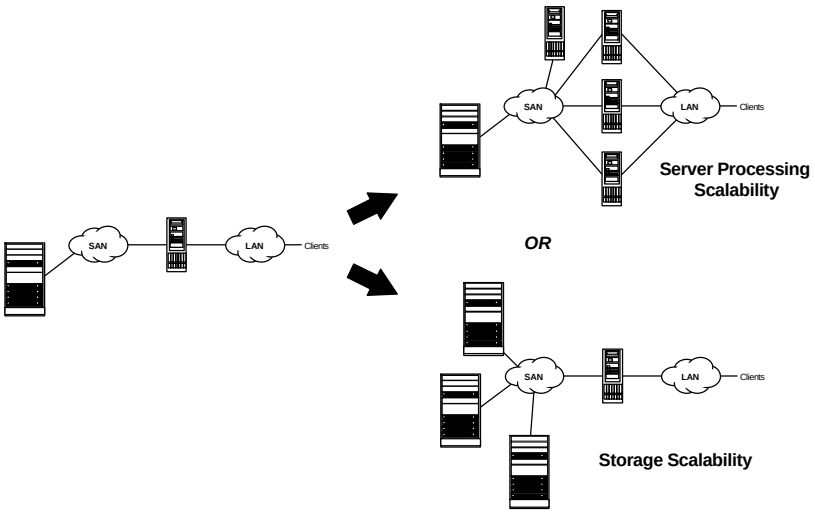

Modular Connectivity

In addition to superior flexibility in scaling server processing capacity and data storage capacity, the networking approach of FC-AL introduces aspects of interconnect scalability that have not been possible with previous architectures. Through the use of modular networking devices such as hubs, switches, bridges and routers, advanced SAN topologies can be created to scale overall bandwidth, enhance availability, and enable advanced SAN application capabilities in storage management and load balancing.

Scalability of Storage, Connectivity and Servers

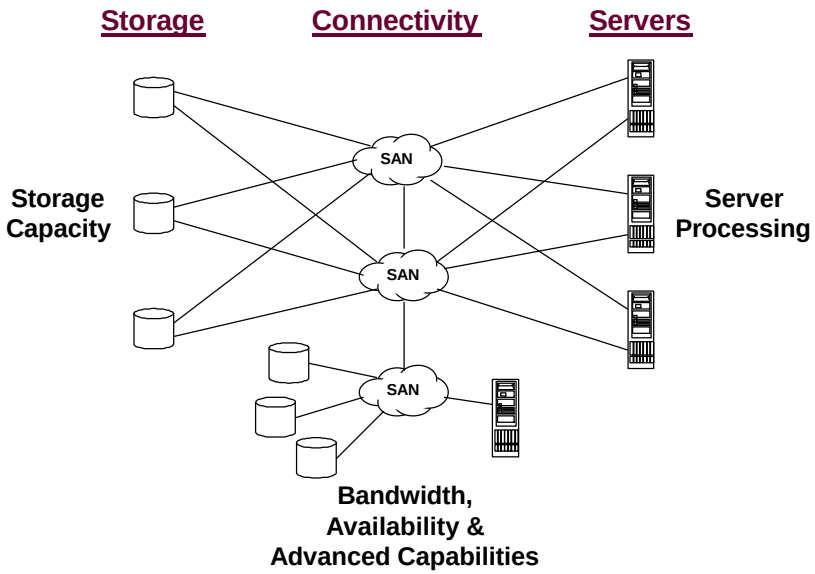

High Availability and Fault Tolerance

In addition to the high availability and fault tolerance provided by specific storage management and clustering configurations, many FC-AL devices provide features that ease the general deployment of fault tolerant SANs. One example of these on-board capabilities is the feature of dual porting, which is has become standard on FC-AL disk drives, to facilitate dual loop configurations. These dual loops provide a redundant path to each storage device in the array in the event that one of the loops is down or is busy.

Dual Loop Array Configuration

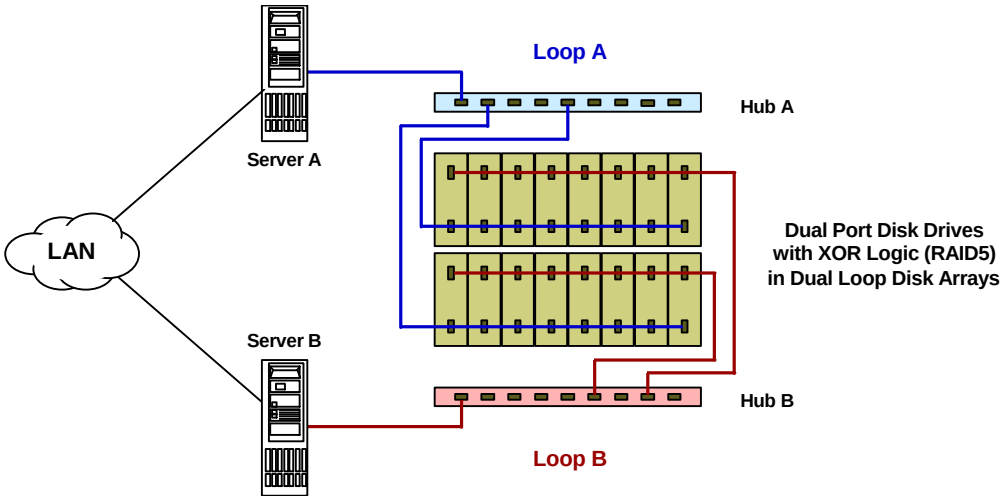

The implementation of Redundant Array of Independent Disks (RAID) configurations in storage arrays has become a standard approach for fault tolerance and is fully supported by the SAN environment. In fact, to even further embrace the RAID approach, FC-AL disk drives provide internal exclusive-or (XOR) logic which effectively provides Level 5 RAID capabilities from within the disk drive itself. This addition facilitates robust, proven fault tolerance while reducing the requirement for more complex and costly RAID controllers.

Manageability

Visibility down to the node and device level is essential to easing the efforts of installation, deployment and maintenance of any network. By embracing a network management approach, SAN connectivity devices, such as hubs and switches, have integrated highly evolved management capabilities modeled after proven LAN and WAN management techniques. A fully managed SAN platform can offer monitoring and bypass control of individual nodes, loops, enclosures, storage devices, and connectivity devices.

Open Standards Platforms for SAN Management

- SCSI command set
- SCSI Enclosure Services (SES)
- SCSI Self Monitoring Analysis and Reporting Technology (S.M.A.R.T.)
- SAF-TE (SCSI Accessed Fault-Tolerant Enclosures)
- Simple Network Management Protocol (SNMP)
- Web-Based Enterprise Management (WBEM)

By embracing the best-practice network management standards established by LAN and WAN platforms, information regarding SAN topology, status and alerts can be easily accessed by system administrators. This high visibility management not only helps reduce unplanned downtime, but can also simplify remote system recovery and restoration in the event of a failure. Traffic monitoring capabilities can also be embedded into the SAN management system to facilitate sophisticated, cost-effective load balancing and capacity planning.

Ease of Integration

Just as compelling as its ability to advance the capabilities of networked servers and storage, the ability to integrate SAN solutions into an existing network provides tremendous value in ease-of-integration. Since the SAN environment exists behind the server, existing server-LAN connections can easily be leveraged to facilitate a gateway between LAN and SAN, and allow utilization of legacy servers.

The broad cabling options supported by FC-AL also ease the introduction of SANs into existing campus networks. By leveraging pre-existing twisted pair, coax and optical cabling, SAN connection distances of up to 10km can be achieved without the need to pull new cable.

As a key building block of SAN deployment, SAN connectivity devices offer dynamically-configurable, hot-plugging capabilities. Combined with a graphical management interface, these features simplify troubleshooting and accelerate installation.

Total Cost of Ownership

Offering an infrastructure for cost-effective, long-term growth, fault tolerance and manageability, the SAN environment provides Total Cost of Ownership (TCO) advantages which have never before been possible with servers, storage or server-storage connectivity.

SAN Feature	***TCO Benefit***
• Connects to existing LANs	• Optimizes existing investment
• Fully managed environment	• Minimized support costs
• Integrated fault tolerance	• Minimized down time
• Distributes server-storage resources	• Complements Network Computer (NC) paradigm
• Independently scalable servers and storage	• Highly efficient scaling of resources

Advanced Application Capabilities

As an alternative to the traditional parallel SCSI storage interface, FC-AL's bandwidth and scalability provide highly compelling advantages for change. However, it is not until the more advanced facets of FC-AL are explored that the innovative strength of SAN comes into focus. By introducing the network-like features of extended connection distance, IP support, and use of hubs, bridges, switches and routers for complex topologies, the SAN infrastructure enables a broad range of new capabilities that were not previously possible for servers or storage. Examples of these capabilities include: advanced storage management and server-storage clustering.

Advanced Storage Management

Problem: Increasing amounts of network-stored data have become cumbersome, if not impossible, to maintain in a timely, secure, fault tolerant and restorable manner.

Solution: The high bandwidth and topological flexibility offered by the SAN environment accelerates the data backup process and facilitates new, innovative platforms for remote backup, mirroring and hierarchical storage.

Perhaps the biggest challenge facing storage management is the need to provide efficient, secure, high availability access to critical data. To effectively overcome these challenges, a number of fundamental issues must be addressed:

Storage Management Challenge	*SAN Solution*
• Length of time required to backup data	• Bandwidth and protocol efficiency accelerate backup
• Inability to backup, mirror or restore remotely	• Cable lengths up to 10km support remote operation
• Lack of alternatives to local backup and mirroring	• Ideal platform for distributed hierarchical storage management (HSM)
• Use of LAN connections for server backup consumes client network capacity	• Separation of server-storage connections from LAN connections reduces LAN traffic

The bandwidth and connectivity limitations imposed by server-to-storage parallel SCSI connections and server- LAN connections offer little to address these formidable tasks. Through its bandwidth, extended connectivity and transport efficiency, the SAN environment uniquely offers a broad range of solutions for storage management, including remote backup, mirroring, recovery, and distributed hierarchical storage management using a broad range of on-line and near-line storage devices.

Remote Backup

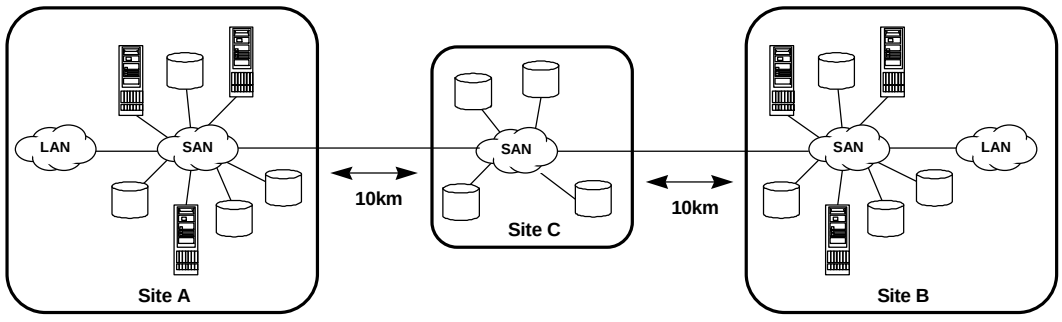

Server-Storage Clustering

Problem: Fault tolerance for server or disk array failures requires costly redundant systems.

Solution: Server-storage clusters facilitated by SAN provide high availability and fault tolerance using cost-effective, mainstream server and storage subsystems.

Once considered a solution for high-end distributed processing, server-storage clustering is quickly approaching mainstream markets. With the growing need for high bandwidth, high availability, fault tolerant servers for real time applications, such as on-line transaction processing, the distributed processing and automatic failover features of clustering are increasing in demand. The connectivity and performance benefits of the Storage Area Network make it the ideal open platform for mainstream clustering configurations such as those based on Microsoft's Windows NT Server 5.0 "Wolfpack" release.

Clustering

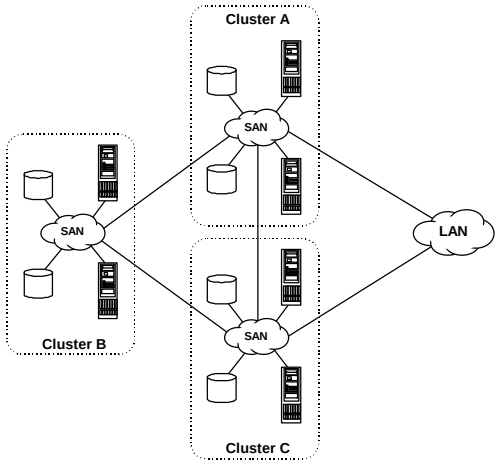

Summary

With the increasing complexity of networked computing systems and global enterprise solutions it is refreshing when a single technology yields both unmatched performance and exceptional Total Cost of Ownership benefits. In the case of Fibre Channel-Arbitrated Loop and the rapidly developing Storage Area Network, an evolutionary open technology promises to revolutionize the network-centric, data-intensive computing era through a new, innovative market space.

Open system SAN solutions and products are now available from leading OEMs, integrators, resellers and independent suppliers of FC-AL host adapters, intelligent hubs, disk drives and disk array enclosures. Additional information regarding suppliers of SAN products can be obtained from the Fibre Channel Loop Community industry trade organization at http://www.symbios.com/fclc/fclcmain.htm.

[i] Source: International Data Corporation, 1996

About the author: Dave Tang is the Director of Business Development at Gadzoox Microsystems, Incorporated. Based in San Jose, California, Gadzoox is a leading developer and marketer of advanced Fibre Channel-Arbitrated Loop Connectivity products for the open SAN market. The company was founded in 1994 and is privately owned.