

Storage Area Network (SAN)

By

Hardik Jani

Subtopics

- Need of Storage Area Networks
- What is SAN
- SAN Protocols
- Storage and Data Sharing
- SAN Backup
- SAN and Data Availability
- SAN Management

Need of SAN

- High-speed data protection
- Growing data
- Need to quickly restore and recover critical information


Organizations need a storage management solution that provides:

- Fast backup speeds matching allotted backup window without adversely affecting network band-width.
- Flexibility to connect and share remote devices and servers.
- Improved scalability to expand the storage infrastructure without rebuilding.
- Interoperability between disparate systems.
- Centralized management to lower the overall total cost of ownership.

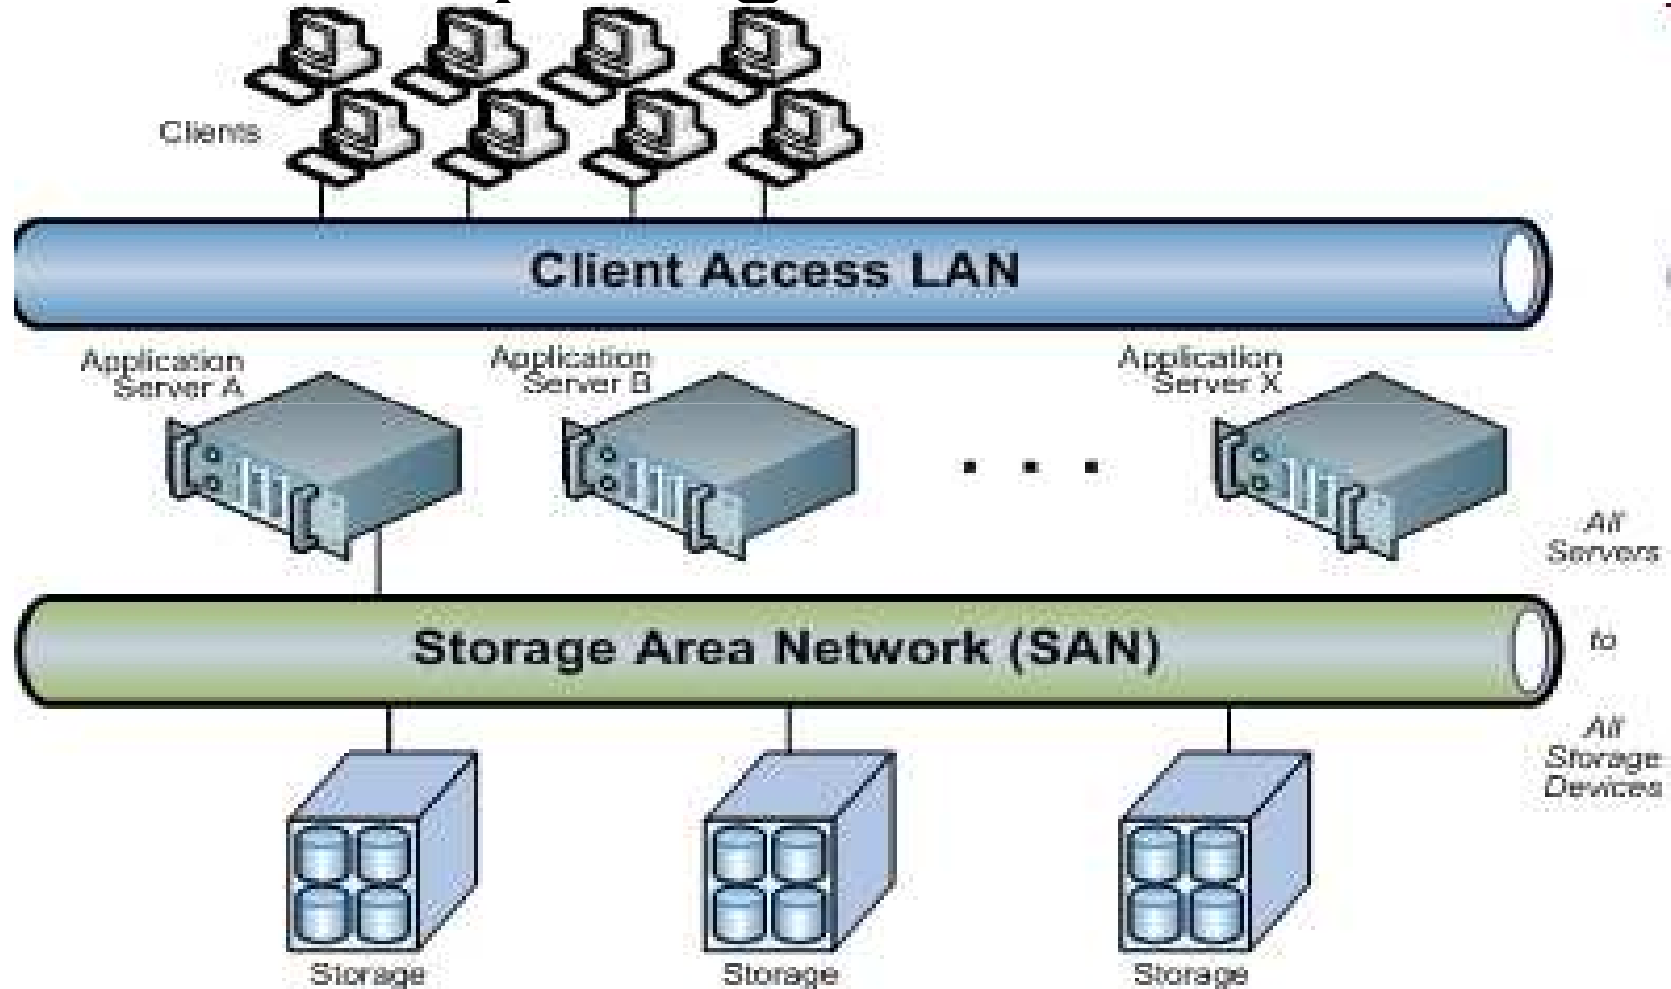
SAN Definition

- Any high-performance network whose primary purpose is to enable storage devices to communicate with computer systems and with each other.
- Involves moving network storage from its traditional location—inside, or directly connected to, file servers—to a separate network of its own.
- Disk, tape, and optical storage can then be attached directly to this network.
- Dedicated storage network that carries I/O traffic only between servers and storage devices.
- Enables direct storage-to-storage interconnectivity.

Conventional Client Server Computing system



SAN based distributed computing architecture



SAN Protocols

- FibreChannel and FibreChannel Protocol
- FibreChannel over TCP/IP
- FibreChannel over Ethernet
- Internet FibreChannel Protocol
- Internet Small Computer System Interface
- Internet Storage Name Space and iSNS Protocol
- Network Data Management Protocol

FibreChannel and FC Protocol

- Fibre Channel Standards (FCS) defines a high-speed data transfer mechanism that can be used to connect workstations, mainframes, supercomputers, storage devices and displays.
- FCP is the interface protocol of SCSI on the Fibre Channel.
- Key Characteristics: Performance, Support, Small Connectors, High B/W utilization

FCIP

- Describes mechanisms that allow the interconnection of islands of FC SANs over IP-based networks to form a unified storage area network in a single Fibre Channel fabric.
- Relies on IP-based network services to provide the connectivity between the storage area network islands over LANs, MANs or WANs.

FCoE

- Standard of mapping of Fibre Channel frames over selected full duplex IEEE 802.3 networks.
- Allows Fibre Channel to leverage 10 Gigabit Ethernet networks while preserving the FCP.
- Main application: data centers.

Internet Fibre Channel Protocol

- Gateway-to-gateway protocol
- Provides fibre channel fabric services to fibre channel devices over a TCP/IP network.
- Uses TCP to provide congestion control, error detection and recovery.
- Primary objective is to allow interconnection and networking of existing fibre channel devices at wire speeds over an IP network.

iSCSI

- TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts and clients
- The SAN makes possible to use the SCSI protocol in network infrastructures for high-speed data transfer at the block level between multiple elements of data storage networks.

Internet Storage Name Service and iSNS Protocol

- iSNS facilitates scalable configuration and management of iSCSI and FC storage devices in an IP network, by providing a set of services comparable to that available in Fibre Channel networks.
- iSNS allows the administrator to go beyond a simple device-by-device management model, where each storage device is manually and individually configured with its own list of known initiators and targets.
- Using the iSNS, each storage device subordinates its discovery and management responsibilities to the iSNS server.

Network Data Management Protocol

- Open protocol for enterprise-wide network based data management.
- Defines a network-based mechanism and protocol for controlling backup, recovery, and other data transfers between primary and secondary storage.

| | Ultra Wide SCSI | FibreChannel SAN |
|----------------------------------|---------------------------------|---|
| Data Transfer Rate | 40 Mbps | 100 Mbps |
| Scalability | 15 Devices | 126(FC-AL), Virtually Unlimited(Switched) |
| Max. Length | 10 feet,inflexible cable | 6.25 miles, easy to interconnect |
| Hot Swap Support | No | Yes |
| Manageability | Server-dedicated device | Load balancing multiple servers across multiple devices |
| Connectivity | Costly reconfiguration required | Hot swap new devices into hub/switch |
| Availability 4/16/2009 | None | Easily redirect job to another server on the loop |


Storage Sharing

- Enables multiple computers across a corporate n/w to access a common set of storage devices.
- Resources are moved out of the server and attached directly to the network, thus allowing them to be addressed directly by multiple servers.
- Exclusive access is provided while a device is assigned.

Data Sharing

- Data is moved onto a much faster infrastructure allowing multiple computers to transfer large files concurrently at rates comparable to locally attached disks over the SAN without adversely affecting the corporate LAN.
- Data sharing requires that the participating computers be able to find and use the contents of a file.
- Computers with different O.S. must use protocol translation modules and other software to establish a common communication dialect.

SAN Backup



LAN-free backups

- Occurs when several servers share a single tape library.
- Each server connected to the SAN can back up to tape drives it believes are locally attached.
- Data transferred via the SAN using the SCSI-3 protocol, and thus doesn't use the LAN.
- All that is needed is software that will act as a "traffic cop."

Client-free backups

- If a client has its disk storage on the SAN, and that storage can create a mirror that can be split off and made visible to the backup server, that client's data can be backed up via the backup server; the data never travels via the backup client. Thus, this is called *client-free* backup.

Server-free backups

- If the SAN to which the disk storage is connected supports a SCSI feature called **extended copy**, the data can be sent directly from disk to tape, without going through a server.
- Newest area of backup and recovery functionality being added to SANs.

Data Availability

- Provides continuous client availability to storage devices if a server in the loop fails
- Capable of replicating data and application files in real time to secondary servers on the SAN
- Tool to remotely mirror data to an alternative site at local, metropolitan, or worldwide locations, providing further levels of data protection and redundancy.

SAN Management

- Centralized management of all physical and logical storage resources via a single console.
- Storage Resource Management (SRM) defines applications that monitor and manage physical and logical resources.
- Effective SAN management necessitates SRM tools be integrated with the SAN management solutions.

Contd...

- Ideal enterprise management tool includes the SAN as a storage network topology and as a sub-network to the communication LAN or WAN.
- Management solutions should offer SAN discovery topology mapping and end-to-end management for fibre channel devices in the loop.

SAN Summary

- SANs represent a huge stride toward a cost-effective solution, providing increased performance, fault tolerance, and scalability for long-term growth.
- In addition, SANs provide total cost of ownership benefits such as:
 - Minimized down time
 - Improved LAN performance
 - Ability to connect to existing LANs
 - Reduced administrative costs
 - Leveraging of existing hardware
 - Improved fault tolerance
 - Maximized storage resources through load balancing
 - Total SAN management.