

Storage Area Networking

Frequently Asked Questions

Author: Andrew G. R. Hogg
Dept: Storage Line of Business
Area: SCHQ, USA
Date: October 23rd, 1998


1. What is a SAN?

SAN stands for storage area network. The current Hitachi Data Systems definition for a SAN is as follows. Note that this definition does not depend on operating systems, intelligent storage, or a particular topology or data transfer medium.

A storage area network is

1. A network of storage resources
2. that are attached to one or more host servers
3. via a storage protocol
4. that provides any-to-any access between servers and storage.

The diagram below shows a typical SAN configuration and how it differs from a basic LAN.



Traditional LAN with server-captive storage

SAN with any-to-any connectivity between servers and storage

Present-day SANs are homogenous. Their hosts employ the same operating systems, and the connecting host dictates their file structure. Heterogeneous, multi-operating-system SANs of the future will incorporate robust methods for cross-platform data sharing and file systems that are not operating-system-centric.

2. What is driving the increasing levels of SAN adoption?

The limitations in speed, distance, and connectivity of SCSI technology prompted the search for an alternate method to access storage devices. The need for LAN-free backups and data sharing started the initial move toward SAN technology. These necessities and the desire to keep all data sets on-line and accessible 24 hours a day to an increasingly global and/or Internet-based user population are driving current adoption.

- **Backup Capabilities:** Increasing data storage requirements and the need for 100-percent availability of applications have overwhelmed SCSI-based backups across the LAN.
- **System Flexibility/Cost:** SANs are storage-centric networks that provide easy scalability, allowing servers or storage to be added

individually and independently of one another. There is no need to add both servers and storage to the network at the same time.

- Availability/Performance: The use of a storage (rather than a network) data transmission protocol permits the transfer of large amounts of data with minimal latency and overhead.

3. What are the primary benefits of a SAN?

The primary benefits of a SAN are:

- Availability: A single copy of data is accessible to any and all hosts via multiple paths.
- Reliability: Dependable data transportation ensures a low error rate, and an ability to recover from failures is provided.
- Scalability: Servers and storage devices may be added individually and independently of one another, and do not depend on proprietary systems.
- Performance: Fibre Channel has a 100MB/sec bandwidth and low overhead, and it separates storage and network I/O. This removal of the storage I/O to a separate network further improves the transmission of network I/O on the LAN.
- Manageability: Emerging software and standards for both FC-AL and Fibre Channel fabric allow single point-of-view (centralized) management and proactive error detection and correction.
- Total Cost of Ownership: Due to path redundancy and superior manageability, there are reductions in failures and downtime. Out-of-band management requires fewer personnel and terminals to manage the system. Storage and servers can be added independently of one another, permitting the initial configuration of smaller, lower cost, systems that can then be easily expanded. Pooling storage resources and allowing multiple-server access reduce storage requirements and costs.

4. What benefits do Hitachi Freedom Storage™ products bring to the SAN environment?

The 7700E and 5700E offer several technologies that provide key benefits to the SAN end user.

- Out-of-band (OOB) Management: This occurs when management protocols are provided on a separate, lower cost network such as Ethernet. In contrast, in-band management occurs when management data is transmitted on the storage data network, between data transmissions or simultaneously, as a lower frequency waveform. OOB management offers several key advantages over in-band management.

No data network impact: No management traffic is passed across the data network, which ensures that business-critical data transmissions are not negatively impacted.

High availability: Management data is available when the network is down. In-band systems cease functioning when the network fails—at the time when management data is most needed.

Accessibility: Due to the broadcast nature of the TCP/IP protocol used in OOB, management data for every node is available at any point in the network. In-band management, using nonbroadcast storage protocols, only allows access to that data that is passing through the connection point.

- **Fabric Capabilities:** Every HDS storage system that supports Fibre Channel will offer Fabric Login capability. This allows the storage system to make use of fabric services such as Fabric Login, WorldWide Name, and State Change Notification. These features enable many of the benefits of fabric and allow true sharing of storage systems across multiple host servers.
- **Scalable Architecture:** HDS offers storage systems that scale from the department up to the enterprise, including open systems UNIX® and NT as well as mainframe S/390® capability. Additional storage resources can be added without shutting down the system.
- **Centralized Management:** HDS offers Hitachi Storage Central (HSC), a vision of nonproprietary, common management of Hitachi Freedom Storage in both centralized and distributed environments. HSC is being developed with industry groups and leading application partners (such as Computer Associates) to satisfy customer expectations of product interoperability and the provision of an end-to-end solution.

5. When should I use a switch and when should I use a hub?

- **Hubs:** Ideal for small scale, entry-level systems, hubs typically cost less and offer a lower throughput than switches.
- **Switches:** Data-intensive, high-bandwidth applications such as file backup, video editing, and document scanning can make full use of switches. Due to their redundant data paths and superior management capabilities, switches are also better suited for environments where high availability and reliability are necessary.

6. Are there reasons to use switches instead of hubs in a SAN?

The use of switches instead of hubs leads to several advantages for a SAN.

- **Failover Capabilities:** In a switched fabric, if a single switch fails, other switches in the fabric maintain operation. A hub-based loop typically fails if a single hub on the loop fails.
- **Increased Manageability:** Switches support the Fibre Channel Switch (FC-SW) standard, which makes addressing independent of the subsystem's location on the fabric and provides superior fault isolation and high availability. FC-SW also allows hosts to better identify subsystems connected to the switch.

- **Superior Performance:** Switches offer multiple-transmission data flow, and each fabric connection can simultaneously maintain a maximum 100MB/sec throughput. A hub-based system offers a single-transmission data flow and a fixed, maximum aggregate throughput of 100MB/sec.
- **Superior Scalability:** Interconnected switches provide thousands of connections without degrading bandwidth. A hub-based loop is limited to 126 device connections, although between five and eight are typically used for ideal system performance.
- **Increased Availability:** Switches support the hot addition of subsystems (hosts and storage) without the need to re-initialize or shut down. Hubs require a Loop Initialization (LIP) to reacquire subsystem addresses whenever changes (additions or subtractions of storage devices or hosts) occur on the loop. A LIP typically takes 0.5 seconds and can disable a tape system that is in the process of doing a backup.

7. How does a SAN differ from a LAN/WAN?

A SAN is similar to a LAN, in that it is a methodology of connecting multiple subsystems together with standardized hardware and software protocols. A SAN differs from a LAN in two key ways.

- **Server Captive Storage:** LAN-based systems connect servers to clients, with each server owning and controlling access to its own storage resources. Storage must be added to a server rather than directly to the LAN. A SAN connects servers and storage, allows storage resources to be added to the network, and enables any server to directly access storage resources.
- **Storage versus Network Protocol:** A LAN uses network protocols that send smaller “chunks” of data with increased communication overhead. This reduces bandwidth. A SAN uses storage protocols that send larger “chunks” of data with reduced communication overhead. This significantly increases bandwidth.

8. Can legacy systems be used with a SAN?

Yes, but many of the benefits of Fibre Channel would not be available. Systems using SCSI-based components, ESCON[®], and hubs can be incorporated into a SAN through bridges and routers in order to preserve previous investments. The Crossroads 4100 is a good example of a SCSI-to-FC bridge.

9. How do I manage a SAN?

There are two basic methods for SAN management.

- **Simple Network Management Protocol (SNMP):** The first method uses the ubiquitous SNMP, which is based on TCP/IP and offers basic alert management. This allows a node to alert the management system of failed fans, drives, power supplies, and so forth, and is available on almost all network-attachable devices. However, it does not offer proactive management, and it lacks security.

- Proprietary Management Protocol: The second way to manage a SAN is to use the typically proprietary management protocol provided by the manufacturer of the node. In the case of HDS, out-of-band management is accomplished via an Ethernet connection using a TCP/IP based protocol. This enables additional capabilities, such as zoning (security) and LUN management, as well as backup and restore, disaster recovery, and fault management.

In either case, to run the management software, a typical implementation involves the use of a separate terminal (such as an NT server) connected either to the out-of-band Ethernet or in line with the network itself.

10. What other companies are involved in the SAN concept?

Many key industry players support the evolution of the SAN concept and are involved in the development of operating systems, interconnects, storage systems, and application software. These companies include, but are not limited to: Sun, HP, Dell, Compaq, StorageTek, Brocade, ADIC, EMC, Gadzoox, VERITAS, Ancor, Vixel, McData, Computer Associates, SAP, and Oracle. In addition, analysts such as Gartner and Dataquest actively report on the SAN market.

11. Where can I get additional information?

Several excellent sources of information on Fibre Channel and SANs can be found on the Internet. The Fibre Channel Association at www.fibrechannel.com, the Fibre Channel Loop Community at www.fclloop.org and the Storage Networking Industry Association at www.snia.org, are all good starting points.

Hitachi Data Systems
www.hds.com

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia-Pacific Headquarters
11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325-3300
info@hds.com.au

Canada Headquarters
380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdscanada.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-61-8000
info@hds.co.uk

Latin America Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
lad@hds.com

U.S. Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1066
ussales@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd.

The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Freedom Storage is a trademark of Hitachi Data Systems Corporation.

S/390 and ESCON are registered trademarks of International Business Machines Corporation.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

All other brand or product names are or may be trademarks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by HDS. This document describes some capabilities that may be configuration-dependent, and features that may not be currently available. Contact your local HDS sales office for information on feature and product availability.

©1998, Hitachi Data Systems Corporation.
All Rights Reserved/Printer's Code/JOB-123