HDS Storage

HITACHI
DATA SYSTEMS

Hitachi Storage Central and Storage Area Networks

by Hu Yoshida

For computing as critical as your business

Contents

Background 1

- Fibre Channel architecture 2
- The benefits of SANs 4

Fibre Channel Topologies 5

- Point-to-point 6
- Fibre Channel-Arbitrated Loop 6
- Fibre Channel hubs 7

Fabric Switch: The SAN Backbone 8

Routers and Bridge Connections 9

SAN Management 10

- Backup 10
- In-band versus out-of-band management 11

Hitachi Freedom Storage™ and Hitachi Storage Central 12

- Hitachi Remote Copy (HRC) 12
- Hitachi ShadowImage 13
- Hitachi On-line Data Migration (HODM) 13
- Hitachi LUN Manager 13
- Hitachi GRAPH-Track™ 13
- SOMMET 13
- SNMP 13

Hitachi Storage Central with SANs 14

- HDS fibre management services 14
- Open SAN management 14

Hitachi Data Systems—For Computing as Critical as Your Business 15

Hitachi Storage Central and Storage Area Networks

by Hu Yoshida

Storage area networks (SANs) are based upon storage protocols and interconnects that enable “any-to-any” connectivity between servers and storage. ESCON® is an example of a SAN that has been available for some time. It uses ESCON directors, which can connect multiple S/390® processors to storage. However, until recently, the UNIX® and NT worlds have not been able to implement SANs due to the limitations of their parallel SCSI interconnects.

The advent of Fibre Channel technology has opened up exciting new opportunities for SANs to bring large-systems storage configuration and management capabilities to campus or metro environments. Hitachi Data Systems contributes to these new opportunities with Hitachi Storage Central, which enables the management of SANs across a global enterprise from a central location. Through alliances with the many providers of the Fibre Channel infrastructure, which includes hubs, switches, and bridges, Hitachi Storage Central will offer SAN-wide management. Hitachi Freedom Storage products are managed through Hitachi’s unique out-of-band communications architecture, which protects user data while enabling management of configuration, performance, events, and security. Hitachi Storage Central also provides proactive management, which is much more extensive than the SNMP alert management offered by most of today’s network and storage management systems.

This paper explains the benefits customers can realize through the combined advantages of Fibre Channel SANs and Hitachi Storage Central. It also provides a road map for using Hitachi Storage Central with other storage and systems management software.

Background

When client-server systems were introduced in the 1980s, they rapidly began to capture applications from centrally configured mainframes. The costs and capabilities of these systems enabled user departments to install their own compute systems, and corporate data began to be distributed far and wide.

Today, a company's corporate data can be distributed across many separate compute centers. (See Figure 1.) These consist of a number of clients attached via a LAN to one or more UNIX or NT processors with SCSI-attached storage devices. ANSI standards committees defined SCSI protocol standards to enable the "open" attachment of peripheral devices to any server that adhered to these specifications. While the widely accepted SCSI standard offers the benefits of a low-cost, high performance storage interconnect for small servers, it has inherent limitations in bandwidth, addressing, and cabling distances.

This SCSI paradigm creates islands of computing that must be configured and managed separately. Although other networks may connect compute centers for centralized network and application management, storage remains with the servers and must be separately configured and managed. To back up this information, the data is sent over the network to a central backup server, or backup resources are allocated to each server. The former lowers the cost of management, but it puts an extra load on the network and impacts the production use of that resource. Such networks are becoming overloaded due to increasing application demands for more and more data. As application files grow larger, more network bandwidth is required for management of data, backup, and replication, leaving less bandwidth for productive use. Server-attached storage leads to duplication of hardware, software, and data management resources, adds to the total cost of computing, and restricts the company's ability to meet the growing data demands of business.

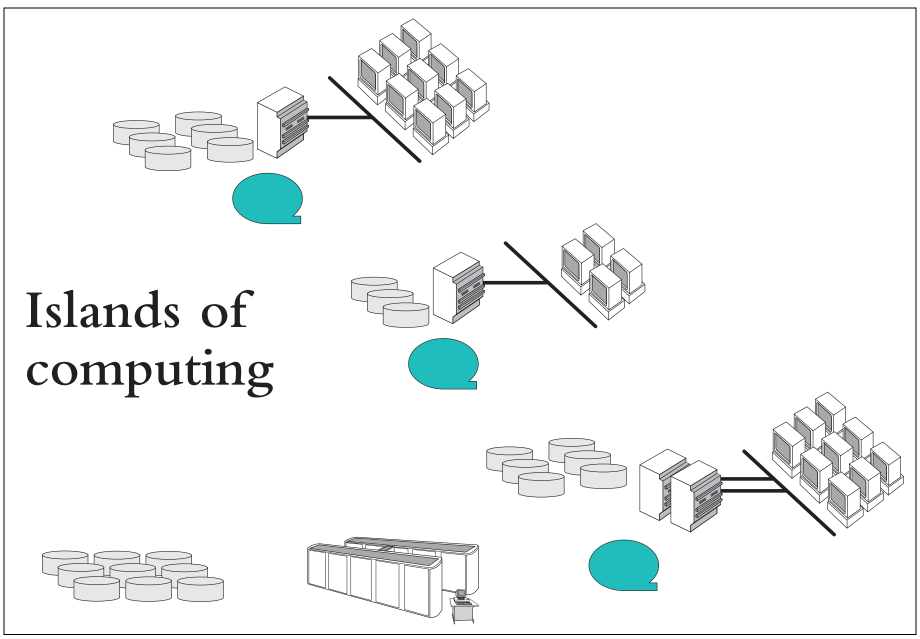

Figure 1: Distributed Computing (Islands of Computing) at XYZ Company

Fibre Channel architecture

The IT community has recognized the limitations of parallel SCSI and has worked to define new standards around a serial implementation that could eliminate or reduce those limitations. Fibre Channel is the result of these efforts. (See Table 1.)

	Parallel SCSI-2	Fibre Channel
Transfer speeds	20MB/sec Fast and Wide 40MB/sec Ultra SCSI 80MB/sec LVDS	100MB/sec
Maximum cabling distance	25m (82ft)	10km per link (6.2mi)
Address space	15 targets 8LUNs per target	16 million nodes
Transmission mode	8 or 16 bits in parallel	Serial by bit
Protocols	SCSI	SCSI, IP, ATM, ESCON
Unit of transfer	Block	Frame
Error encoding	Parity	8/10 encoding, CRC

Table 1: Parallel SCSI-2 and Fibre Channel Comparison

Fibre Channel has been broadly accepted, not only for its higher bandwidth, cabling distances, reliability, and scalability, but also for its ability to combine the best characteristics of channels and networks. (See Figure 2.)

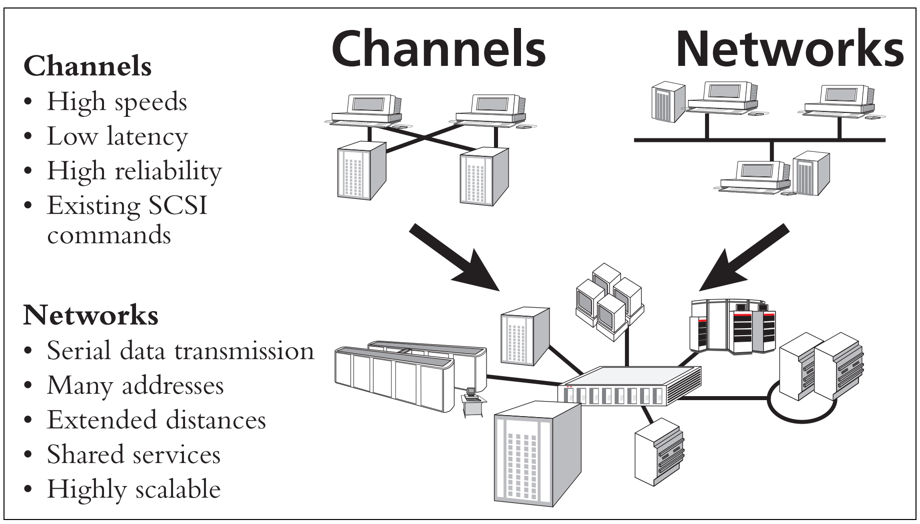

Figure 2: Fibre Channel—The Best of Both Worlds

Fibre Channel transmits data in fixed-length data packets that are prefixed with a start-of-frame identifier (SOF ID) and a header that allows frames to be navigated through a connection of loops, hubs, and switches. The payload can be a maximum of 2,112 bytes. The size of the payload is communicated during Loop Initialization (LIP) in the case of a loop configuration or Fabric Login in the case of a switched fabric. Data blocks are mapped across multiple frames, which are called a sequence. A cyclical redundancy check (CRC) and an end-of-frame identifier (EOF ID) follow the payload. (See Table 2.) In addition to the CRC, Fibre Channel uses an 8in10 encoding scheme to ensure that data is transmitted correctly.

SOF ID	HDR	2112 byte payload		CRC	EOF ID
		Optional HDR	2048 byte payload		

Table 2: Fibre Channel Fixed-length Data Packet

An optional feature is the WorldWide Name, a unique 64-bit name that a manufacturer can purchase through the IEEE. A 24-bit company name is assigned by the IEEE, 36 bits are a vendor-specified ID, and the other 4 bits serve as a format ID. This combination provides a unique identifier for storage elements, wherever they might exist in the world. Switches and directors can use WW Names to establish connections between storage elements that are independent of specific fabric connections. During Fabric Login, the WW Names are registered in a Simple Name Server table, where connections between WW Names are recorded. This capability enables the establishment of a SAN that can provide "any-to-any" connections between servers and storage, using high-speed, low-latency, fibre connections.

The benefits of SANs

A SAN can be employed to provide connectivity from any server to any tape or RAID subsystem in the company. (See Figure 3.) The benefits of this implementation include the ability to pool storage devices in one location, utilize a common automated tape library (ATL) to back up all the servers, consolidate management resources, and offload data traffic from the front-end LAN.

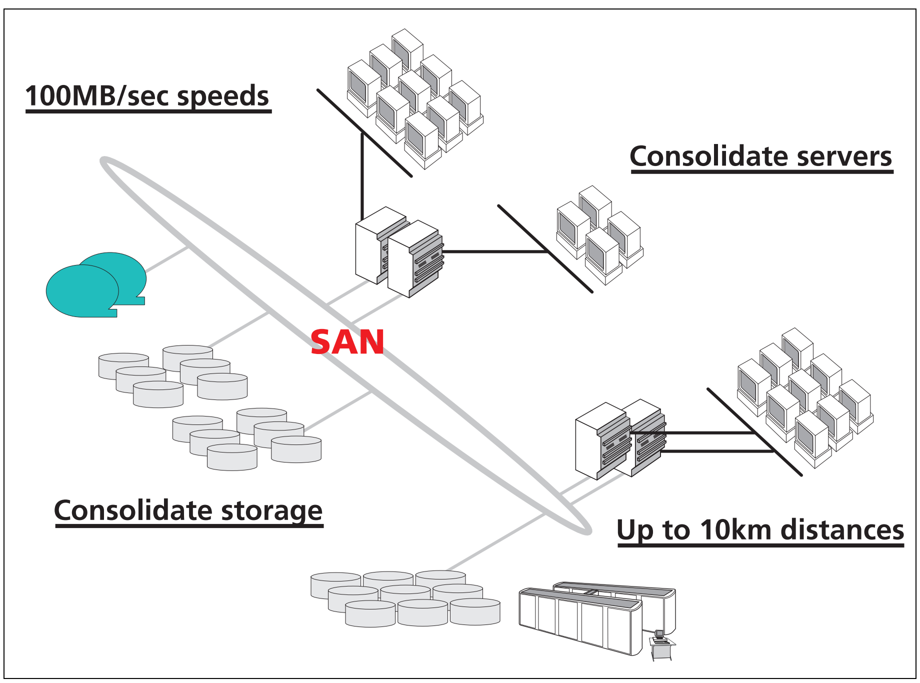

Figure 3: SAN-attached Storage

The benefits of a Fibre Channel SAN include:

- High bandwidth: full duplex 100MB/sec today, increasing to 400MB/sec
- Expanded address space: up to 16 million nodes
- Increased distances between nodes: from 500m with shortwave multimode fibre to 10km with longwave single-mode fibre

- Modular scalability and connectivity
- High availability and fault tolerance with hubs and switches
- Manageability through network management standards (SNMP) established by LAN and WAN platforms
- Ease of integration: support of SCSI protocols that allow attachment to existing applications
- Reduction in total cost of ownership
  - Minimizes support costs by pooling storage resources
  - Minimizes downtime with integrated fault tolerance and error checking
  - Avoids duplicating hardware and software upgrade costs
  - Enables more time for production use due to faster backup
  - Reduces LAN traffic by offloading data movement from the network
  - Increases cable lengths to enable remote operations at minimal costs
  - Employs clustering configurations for high availability and load balance

Requirements for a SAN include servers with Fibre Channel host bus adapters (HBAs), storage subsystems with Fibre Channel port adapters, and Fibre Channel interconnect products like hubs, switches, directors, and bridges.

Fibre Channel Topologies

The three standard Fibre Channel topologies are point-to-point, loop, and fabric switch. In Figure 4, a hub is illustrated as a separate topology, although the industry considers it to be a variation on the loop.

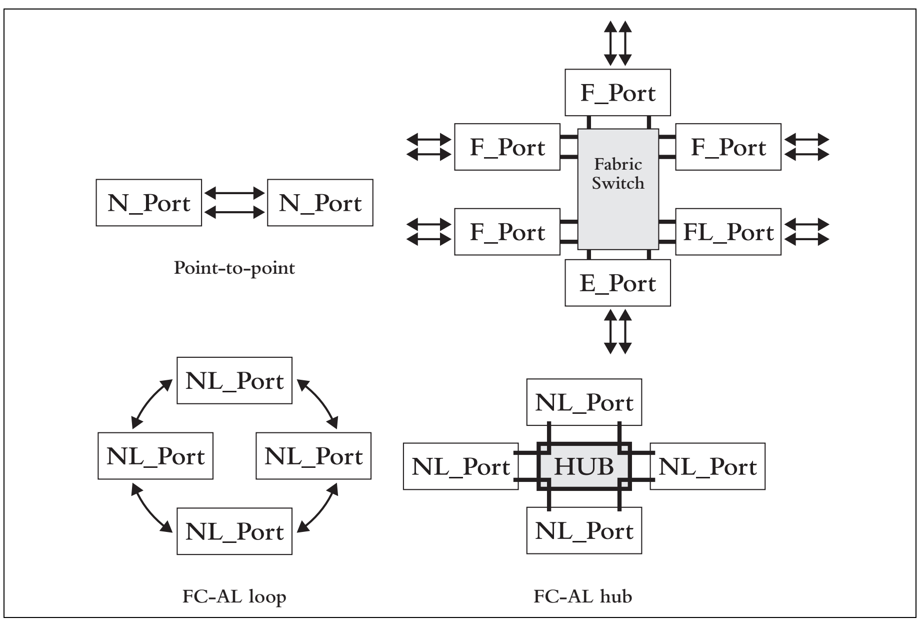

Figure 4: Fibre Channel Topologies. Point-to-point offers one conversation and two nodes; FC-AL loop offers one conversation and up to 126 nodes; fabric switch offers a non-blocking switch, multiple conversations, any-to-any connectivity, and up to 16 million nodes; and a hub provides FC-AL loop availability.

Point-to-point

Point-to-point topology is a single, full-duplex 100MB/sec connection between an HBA and a fibre port on a storage device. This simple configuration consists of one initiator (HBA) and one target (device port) in SCSI-addressing terminology. As such, this configuration can address up to eight logical units (LUNs), which is the limitation for SCSI-2 targets (but has been relieved in SCSI-3). To maximize the data accessible through a point-to-point configuration, the current addressing limitation requires that the target storage device be a RAID subsystem that can be mapped into very large LUN sizes.

Since today's server HBAs and fibre ports on storage subsystems use shortwave lasers, fibre distances are limited to 500m with 50μm multimode cables and 300 meters with 62.5μm cables. A few vendors provide lower-cost copper fibre cables using DB9 connectors. Copper fibre supports up to 30 meters, which is similar to parallel SCSI-differential-ended cables. And copper cables transmitting gigabit frequencies will tend to act like big antennas. For these reasons, HDS has opted not to provide copper fibre ports for Hitachi Freedom Storage products. Conversion from optical fibre to copper fibre can be done through hubs or switches that support both types of fibre connections.

A point-to-point topology is relatively easy to install and configure. Users who are entering the fibre world might consider this configuration as a first step toward implementing a SAN.

Fibre Channel-Arbitrated Loop

The Fibre Channel-Arbitrated Loop or FC-AL protocol is required to allow the attachment of up to 126 fibre devices in one loop configuration. The HBA and the storage fibre port must support the FC-AL protocol to participate in an FC-AL configuration. This protocol provides for arbitration to allow a node (a fibre port on an HBA or storage device) to gain control of the loop for communication with another node in the loop. While these two nodes are in conversation, no other traffic is allowed on the loop. The nodes in between must pass the data packet on to the destination node. If a break occurs in the loop or a node fails, no data trans-mission can occur.

Loops are difficult to manage when attached to multiple hosts, since all hosts have access to all devices in the loop. In open systems, the devices have no knowledge of who owns the data they hold, so data can be overwritten. Software is required in hosts to map devices and mask or unmask them to the appropriate host. NT should not be intermixed in a loop with other systems since it assumes every device it sees belongs to NT and needs to be reformatted.

In FC-AL loops as well as hubs, it is the responsibility of the attaching nodes (HBA or storage port adapter) to map the addresses in a loop. When a node attaches to the loop, it issues a LIP, which causes all nodes in the loop to be placed in an open-init state while their arbitrated loop addresses (ALPA) are reassigned. Each node then "walks" the loop to identify and build a map of other nodes in the loop.

When any changes are made to a loop configuration, the loop must be re-initialized through a LIP. The loop is stopped and, if nodes were added or deleted, the relative addresses of the nodes on the loop are reassigned. Due to the probable occurrence of a LIP, the attachment of tape to an FC-AL loop is not recommended. If a LIP occurs during a tape process, the tape becomes useless, since it cannot recover or write a tape mark. FC-AL loop failures are difficult to diagnose. When a node fails, unlike

Ethernet loops, the FC-AL loop is dead, and diagnostic information is not available unless other external tools like the FINISAR Gigabit Link Analyzer are used.

FC-AL loops are an economical way to attach many fibre devices together. The early adopters of FC-AL were the video and pre-press industries, where low-cost/high-bandwidth storage was required. Most storage vendors offer FC-AL front-end interfaces to their storage arrays for increased cabling distances. Some storage vendors also offer FC-AL on the back end of their storage subsystems, although benefits are currently limited and do not outweigh the additional costs and availability considerations.

Increased cabling distances, which are available with fibre, are not required for back-end connection. Although bandwidth is a consideration, current fibre chip sets can only sustain a throughput of 40 to 80MB/sec. Parallel SCSI speeds are at 40MB/sec with Ultra SCSI and 80MB/sec with Low Voltage Differential (LVD) SCSI. Back-end fibre is usually implemented with copper connections to a backplane, which generates high electromagnetic interference due to the gigabit frequencies required for 100MB/sec fibre speeds. To ensure that a disk failure does not take down the loop, a redundant loop with bypass switches on each disk is required.

Although the transfer rate of an FC-AL loop is 100MB/sec, the bandwidth of the loop is limited to the speed of the devices on the loop, since only one conversation can occur at any time. Devices on a loop must be adequately cached or buffered in order to maintain optimum throughput. FC-AL architecture supports 126 nodes. However, with the addition of arbitration overheads, the practical attachment limitation is on the order of five to eight high-performance disk drives. Direct attachment to an FC-AL loop requires two fibre ports, an in and an out. Attachment to more than one HBA and one storage port requires a hub.

Fibre Channel hubs

To overcome the limitations of FC-AL, companies such as Vixel, Gadzoox, Emulex, G2 Networks, and others, have developed hubs like those created in the networking world to solve similar problems. Hubs support the FC-AL protocol and provide high-availability protection against failures in the loop, automatically bypassing failed HBAs or storage nodes. This protection is not nondisruptive, however. Anytime a failure occurs or the loop configuration is changed, the loop must go through a LIP. This initialization process amounts to the equivalent of a reset on a SCSI bus.

Hubs are required to attach an HBA to more than one storage port in an FC-AL loop. Hubs can also act as a bridge between copper and optical fibre, and between shortwave and longwave ports for communications to other longwave ports located up to 10km away. By using a pair of hubs connected through longwave ports, a server with a shortwave HBA can connect to a storage device residing beyond its shortwave transmission capabilities. The numbers of ports in a hub range from 5 to 12.

Most HUB vendors use modular “snap-in” Gigabit Interface Converters (GBICs) to provide the optical-to-electrical conversions for transmitting and receiving Fibre Channel signals. Different GBICs are used for longwave or shortwave transmissions and can be intermixed in a hub. Although modular GBICs provide configuration flexibility, they also introduce reliability exposures. One hub vendor, G2 Networks, provides hardwired ports to reduce reliability exposures at the expense of configuration flexibility.

There are essentially three types of hubs: unmanaged hubs, managed hubs, and switched hubs. The first hubs in the market were low-cost unmanaged hubs that required external analyzers plugged into adjacent ports on either side to handle hub port failures. These hubs were difficult to maintain and manage and, for availability reasons, most experienced users have since upgraded to managed hubs.

Vendors like Vixel, Gadzoox, and G2 Networks provide a switched hub. This third type of hub can be divided into multiple subloops that can maintain separate conversations at 100MB/sec. However, these are still limited to 126 addresses and do not support switch features like Fabric Login, WW Names, Simple Name Server, and so on.

The three basic concerns with hubs are jitter (the degradation of signal transitions when they are propagated across nodes in a loop), diagnosis of failed nodes in a loop, and the ability to remotely monitor and manage the hub. Vendors manage jitter in different ways, such as re-clocking (Gadzoox), digitizing the signal (Emulex), and eliminating unused hub ports (G2 Networks). Some hub manufacturers, like Vixel, have added intelligence to scan the loop to detect errors and diagnose failures. Remote monitoring and management is provided with LAN- or telnet-attached programs. Some vendors provide SNMP. However, standard SNMP Management Information Blocks (MIBs) are still to be defined for hubs.

Fabric Switch: The SAN Backbone

The backbone of a high-performance, scalable SAN is the fabric switch. Fabric is the term used to describe the combinations of connections between nodes on the switch. Like the cross weave of clothing fabric, fibre switches provide connectivity paths between nodes in a SAN. Each connection supports the full 100MB/sec bandwidth. Theoretically, a fabric made up of multiple switches could support up to 16 million nodes.

Storage fibre ports and HBAs must support the Fabric Login facility to fully participate in a switched fabric. Fibre ports, which support Fabric Login, log into a switch to exchange information about packet sizes, buffer credits, and WW Name, and are assigned an address by the switch. This information is registered into a Simple Name Server table, which eliminates the need to "walk the bus" (as was required in parallel SCSI) to discover what devices are accessible. However, this would be impractical in a fabric that could have up to 16 million nodes. Since the switch maintains the addresses for the nodes—not the HBA, as in the case of FC-AL—changes in node configurations do not disrupt other nodes on the switch.

Fibre nodes in a fabric can be nondisruptively added to or deleted from the rest of the fabric. A failure in a host or storage device attached to a switch should have no impact on the rest of the fabric. If more than one path is available between two nodes, the switch can reroute the frames through an alternate path in the event of a failure.

Loops can be attached to switches through a special port called an FL port. If the attaching port on the loop does not support Fabric Login, it is attached as a private loop and cannot access any nodes on the other side of the switch. The Brocade Switch has a translative mode that creates a pseudo WW Name for FL ports that do not have one. This allows the loop to act as a public loop that can be accessed by any node in the fabric. If an FC-AL node has the Fabric Login capability, it can log in to a switch through an FL port as a public loop. The login allows this FC-AL node to be seen by the other nodes in the fabric. Access to a device on the public loop still requires an arbitration process.

Fabric switches provide the reliability of hubs in addition to alternate path capability. Security is enhanced by the ability to zone certain nodes so that they are only visible to selected servers. This is particularly important when intermixing NT nodes on the switch. Currently, the only two vendors who provide switches with full fabric capability are ANCOR and Brocade. Other vendors provide switching hubs that were described above. McData provides a fibre director, which also supports the fabric capability. Fibre directors, like ESCON directors, contain multiple switches with full failover capability and hot-pluggable components. Alternate controllers are provided for nondisruptive microcode upgrades. Other vendors of ESCON directors, like InRange, are expected to provide fibre directors in the near future.

As shown in Table 3, fibre switches provide the maximum bandwidth and addressability. As a rule of thumb, unmanaged hubs cost about $400 per port, managed hubs cost up to $1,000 per port, switches cost about $2,000 per port, and directors about $3,000 per port. Unmanaged, low-cost hubs are used for configurations with low port populations that do not require higher levels of control. Managed hubs are the next step up for FC-AL loops, and directors provide the larger number of ports within one package—up to 32. As port densities increase, the multipath capabilities of switches and directors become important.

Feature	Point-to-point	FC-AL loop	Hubs	Switches
Maximum number of FCS ports	1	126	126	16 million
Number of transmissions	Single	Single	Single	Multiple
Scalable bandwidth	No	No	No	Yes
Hot-pluggable devices	No	No	Yes	Yes
Bypass failed drive	No	No Requires bypass circuitry	Yes	Yes

Table 3: Comparison of Fibre Channel Topologies

Routers and Bridge Connections

Routers or bridges provide a connection between Fibre Channel and non-Fibre Channel interfaces. The distinction between routers and bridges in a SAN is not very clear and depends on the vendor's definition. Both provide conversion between SCSI on a parallel interconnect and SCSI on a fibre interconnect. Vendors like Crossroads and Chaparral Technologies provide SCSI-to-fibre as well as fibre-to-SCSI conversions, which they call routers. Vendors like ATTO provide a one-way fibre-to-SCSI interface and refer to their products as bridges. Another definition of a router refers to the ability to perform a conversion between protocols, such as the conversion between ESCON and FICON that enables ESCON devices to participate in the Fibre Channel world. There are many flavors of routers supporting the different combinations of SCSI and fibre nodes. A SAN solution must consider routers for non-Fibre Channel interconnects. Bridges are primarily used to connect SCSI tapes and SCSI tape libraries to host fibre HBAs or to fibre fabrics. They can also be used to attach legacy SCSI devices and subsystems.

SAN Management

The introduction of a SAN between servers and storage also presents an array of new management challenges. Nodes and devices must be installed, deployed, and maintained in the storage network. Security features must be provided to ensure that a server on the SAN does not have access to LUNs that belong to another server. A fully managed SAN platform must be able to monitor and provide control down to the individual node and device. Since SAN infrastructure can be spread over a large metro area, remote management must be provided from a central site.

Increasing port densities also require more management capabilities. Most managed hubs, switches, and directors support SNMP, which was developed in the networking world to solicit status and set operating parameters for basic availability management in a network. The common language of this IP protocol is supported by multivendor platforms. Device status information is organized in a management information block, which is maintained by an SNMP agent in the device. The protocol also enables devices to generate unsolicited status information (alerts) if a preconfigured condition or threshold is reached. To participate in a managed SAN environment, hubs, switches, and directors, as well as storage devices, must be able to support SNMP. In addition to SNMP, most fabric vendors provide their own graphical interfaces to focus on vendor-specific parameters and real-time diagnostics. A comprehensive SAN management strategy must be able to integrate these vendor-specific management tools.

Proactive management should be able to both monitor and avoid outages and provide automatic load balance and asset management. Other essential features are security, which includes login authorization, zoning for access control, and encryption. Data backup and recovery is also key in a SAN management solution.

Although most hub, switch, and director vendors support SNMP, there are no standard MIBs for many node interfaces, such as FL nodes. There are a lot of other parameters—as well as vendor unique information—that are needed for management of a SAN. The Storage Network Industry Association (SNIA) is leading the effort to define a set of common reporting parameters and interfaces.

Backup

One of the major benefits of a SAN is the ability to do LAN-free backups over 100MB/sec fibre links. Another important benefit is the ability to consolidate tape backup resources. In a server-attached storage world, backup tapes are duplicated at each server or backups are completed over a LAN using products like Legato NetWorker or VERITAS NetBackup. Performing a backup over a 10/100bT communications link is highly disruptive to applications—the primary LAN users. Moving backup traffic off the LAN frees it for application use, and executing the backup at 100MB/sec shortens the backup window. Consolidating backup resources reduces hardware, software, and management costs. The extended distances available with SAN enable high-speed backups to disaster recovery sites.

High speeds and longer distances also provide faster replication of data for use by other applications or for point-in-time copies.

Fibre Channel ports are not generally available for tapes today. SCSI tapes are attached to fibre through bridges like the Crossroads CrossPoint fibre-to-SCSI router products. Although some early implementations attached the bridges to SANs through hubs, HDS recommends that tapes be attached through a switch, for availability. Hubs are subject to LIPs if any nodes on the hub experience a failure. As previously mentioned, tapes cannot recover from a LIP since that is equivalent to a bus reset. In the near future, intelligent agents may be put into the nodes to do third-party backups without the continued intervention of the host server.

Network Data Management Protocol (NDMP) is an emerging industry standard that will enable plug-and-play backup between management software running on different hosts and a network file server or a network-attached storage device acting as a file server. The protocol promises to enable users to back up network file system servers or network-attached devices without a general-purpose operating system running. In addition, the protocol provides a path for fast backup that lets data go straight from a file server's disk to tape instead of moving through a general-purpose server. HDS supports the NDMP initiative. See: http://www.ndmp.org for more information.

In-band versus out-of-band management

In traditional network management, SNMP commands, SCSI Enclosure Services (SES), and other management traffic is intermixed with data and sent along the primary data path. This is called in-band management, and is preferred by network managers because it allows them to configure devices and monitor status from any management station in the network. The disadvantage of in-band management is that all traffic, including management information, is unavailable if the network is down. Also, management information contends with the data for network bandwidth.

A SCSI fibre SAN is different from a LAN in that it does not support the ability to broadcast like Ethernet. Today, in-band management in a SAN only allows the node on the other end of the data path to see the management information. In order to provide a consolidated management view, each receiving node would have to resend the information across a LAN to a central manager.

Another difference is that LANs are optimized to transfer small packets of information, while fibre SANs are optimized for large data transfers. Mixing these two different types of data in the same data path becomes a much greater issue for arbitrated loops. These require a great deal of overhead to repeatedly solicit information from the loop's node, and then arbitrate, open, and close, in order to transfer small packets of management information. This overhead impacts the bandwidth for mission-critical data.

For this reason, most SAN hubs and switches provide an out-of-band management capability, either through Ethernet or serial RS232 interface. Out-of-band management is important for multiplatform storage subsystems that can attach to many different hosts. A single host view would not show the impact of shared back-end resources, such as array groups, cache, and internal busses. Also, in-band management does not give the user the ability to centrally manage the entire subsystem.

Hitachi Freedom Storage™ and Hitachi Storage Central

All Freedom Storage products, including the 5700/5700E and the 7700/7700E, are network-enabled for out-of-band management. They all have an IP address and provide SNMP reporting. (See Figure 5.)

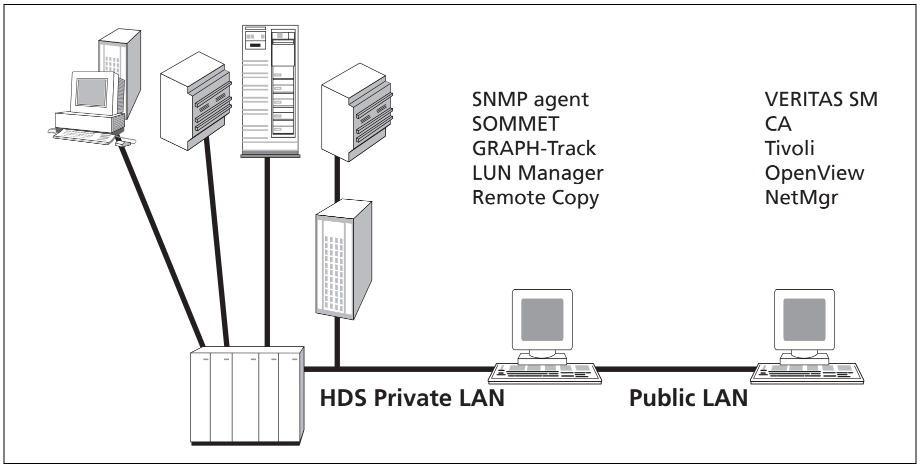

Figure 5: Hitachi Freedom Storage—Network Enabled

The processors in Hitachi Freedom Storage report status to an optional remote console over an internal Ethernet LAN. The processor data is captured over internal command busses, which are independent from the internal data busses. No data is ever exposed to the internal LAN. In addition to monitoring status, management applications running on the LAN-attached remote console can configure and invoke various features of Freedom Storage. The remote console can also report status to higher-level storage or system management software like Tivoli, CA Unicenter, Sun Net Manager, HP OpenView, or VERITAS Storage Manager. This architecture and the suite of management applications on the remote console are known as Hitachi Storage Central.

Hitachi Storage Central provides centralized management of storage across an out-of-band private LAN. Whether the storage is distributed across many locally attached SCSI storage subsystems or centralized in a SAN, Hitachi Storage Central provides the capability to manage it from one central point. The following Hitachi Storage Central programs are available today for management of Freedom Storage subsystems.

Hitachi Remote Copy (HRC)

This real-time copy capability enables nonstop solutions, such as disaster recovery and data migration. Graphic and command line interfaces (PPRC commands) control synchronous and semi-synchronous mirroring of logical volumes. HRC supports mainframe and open systems host platforms. Synchronous S/390 HRC can be managed using IBM PPRC and P/DAS support. Hitachi Extended Remote Copy (OS/390®-based asynchronous remote copy) to/from other vendor subsystems is also supported.

Hitachi ShadowImage

Business agility can be enhanced with ShadowImage's point-in-time copy capability. Graphic and command-line interfaces control data replication and fast resynchronization of logical volumes. Three copies of a logical volume may be maintained within the same subsystem. ShadowImage also works in concert with HS-DataPlex and Remote Copy to provide additional copies in the same or other subsystems, and it supports mainframe host platforms.

Hitachi On-line Data Migration (HODM)

HODM provides fast, on-line device migration to Hitachi Freedom Storage from other vendors' subsystems. It updates both source and target subsystems to ensure data availability if the migration process is interrupted, and it supports mainframe data.

Hitachi LUN Manager

LUN Manager is available for configuration and assignment of LUNs to SCSI or Fibre Channel ports on the 7700E. This does not require a service call. LUN Manager is also available for the 7700, which does not have Fibre Channel ports.

Hitachi GRAPH-Track™

GRAPH-Track on the 7700/7700E provides a view of the entire performance of the storage subsystem, which is important in multiplatform configurations. Traditional host-driven performance monitoring like RMF on MVS mainframes can only see the activity on host-attached channels. It does not see the contention for internal storage resources, such as like cache, busses, and array groups, that is caused by other host accesses. Out-of-band performance monitoring is the only way to see the complete performance picture. GRAPH-Track provides a "drill-down" view of performance within the subsystem and allows modification of parameters, such as cache, write line, and bus access modes. GRAPH-Track offers scheduling as well as a logging facility.

SOMMET

SOMMET, or Storage On-line Management and Media Error Transmission, is a program that is used to configure and monitor the 5700 and 5700E line of Freedom Storage products from a central location. It provides the invaluable ability to extract subsystem usage and performance information for storage analysis tasks.

SNMP

The 5700/5700E and 7700/7700E products support SNMP. A robust status and event monitoring capability is available for support of call-home maintenance, which is much more robust than that available with SNMP. The subsystem reports system information messages (SIMs), which are monitored and reported to the HDS call center for pre-emptive maintenance. In the 7700/7700E, the remote console converts the SIMs to SNMP MIBs for reporting over a public LAN to a storage-management or systems-management software like VERITAS Storage Manager, Tivoli, TNG, etc. This out-of-band reporting of SNMP ensures that the storage or systems manager can determine the status of the HDS storage subsystem even if the data channel path is down. In the 5700/5700E, SNMP is reported through an optional LAN card in the subsystem.

Hitachi Storage Central with SANs

Hitachi Storage Central provides a way to manage a SAN from a central control point for maximum availability, scalability, and performance. As with any new technology, there will be a transition period where the SAN and legacy systems will co-exist. Hitachi Storage Central will also provide management for non-SAN-attached storage subsystems throughout this transition.

Since Hitachi Freedom Storage connects to Hitachi Storage Central through a private LAN, there is no need to wait for standards to be defined to implement a managed SAN. This can be accomplished today. HDS is currently working with selected hub and switch vendors to connect their out-of-band management interfaces with Hitachi Storage Central. (See Figure 6.)

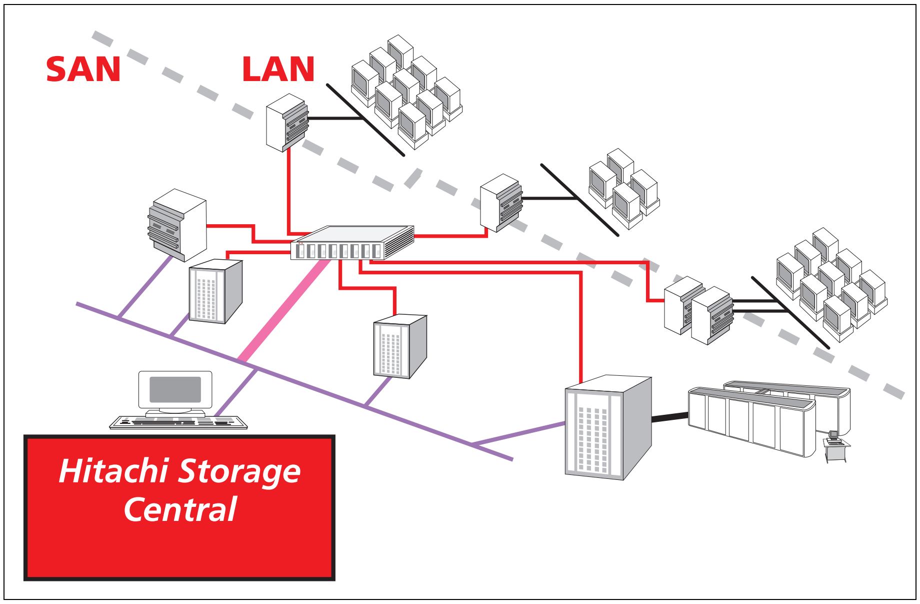

Figure 6: XYZ Company Employing Hitachi Storage Central with a SAN

HDS fibre management services

HDS' experience in fibre development and testing has shown that integration of all SAN elements requires testing of many combinations and permutations of vendor products. HDS will consider a select list of fibre vendors—for hubs, switches, and directors—as possible partners to help build an end-to-end services contract for design and implementation of a SAN. HDS' current services have been defined based on our experiences with fibre management service for ESCON fibre implementation.

Open SAN management

Because customers need to manage multiple vendor devices in a SAN, HDS is working through organizations like SNIA (http://www.snia.org) to define a common architecture and interface that will provide a centralized management capability. We envision a three-tier approach to storage management. The top tier would be systems management, provided by such companies as CA Unicenter, Tivoli, OpenView, Net Manager, and so on. Hitachi Storage Central would report directly to them through SNMP for alert management. The second tier would be storage management provided by vendors, such as VERITAS, who can manage file systems and volume managers, as well as storage. Hitachi Storage Central would provide information to this tier or be launched from this tier for specific functions, like GRAPH-Track. The third tier would be product-specific management, such as LUN Manager for the 7700E. (See Figure 7.)

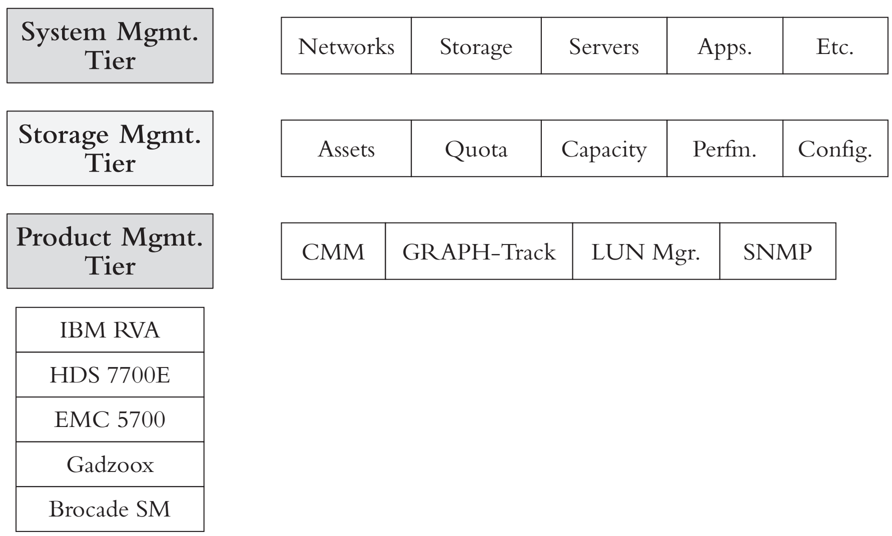

Figure 7: HDS Storage Management Model

Hitachi Data Systems—For Computing as Critical as Your Business

Hitachi Storage Central and Freedom Storage support full fibre fabric capabilities, and position HDS to be a leader in the SAN marketplace. Hitachi Storage Central provides centralized SCSI and SAN management. SANs provide the scalability, configurability, availability, performance, and total cost of ownership (TOC) that comes from the ability to pool the storage, regardless of server location. Working with leading providers of SAN hubs, switches, directors, and routers, HDS will provide an end-to-end services answer for rapid and reliable implementation of SAN solutions. HDS is committed to working with industry associations like SNIA to develop and implement open standards for SAN management.

Fibre Channel SANs and Hitachi Storage Central are important enablers of HDS' "No-Limits" storage philosophy. By delivering comprehensive, customizable, manageable, and application-enabling hardware and software, HDS helps to build and protect the mission-critical environments of the world's largest corporations—ensuring access to any information, from any computer, anywhere, anytime.

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