HDS Storage

HITACHI
DATA SYSTEMS

HDS Fibre Channel Implementation for Data-Intensive Applications

by Hubert Yoshida and Dennis Stein

For computing as critical as your business

Contents

HDS Fibre Channel Implementation for Data-Intensive Applications

by Hubert Yoshida and Dennis Stein

Improved connectivity arrives just in time for business

Information is recognized as any company's most important asset. That's why keeping information flowing across the enterprise, managing it effectively, and securing its safety are all critical business issues. Taking control of these issues requires advanced technology and the knowledge of how to use it. Now, as information overload begins to crowd old pathways, improved connectivity has made a welcome appearance on the business scene.

Business needs instigate a technological transition

SCSI is an acronym created by the American National Standards Institute (ANSI) to define Small Computer Systems Interface for low-cost, parallel bus-oriented interconnects. The computing industry is currently transitioning from SCSI-2 to SCSI-3 standards for computer input/output interfaces in order to meet the increasing demands of data-intensive applications. The first SCSI-2 devices became available in 1988 and helped to launch the client/server computing paradigm. Since then, the technology of processors, I/O devices, systems, and applications has been improved by orders of magnitude, and SCSI bus interconnect speeds have become a severe limitation. SCSI-3 standards are being introduced to overcome these limitations through the use of serial interconnects and special commands that optimize the use of RAID and automated media libraries. The transition from SCSI-2 to SCSI-3 is intended to provide new transport mechanisms for the delivery of existing commands while preserving the investment in current applications.

The SCSI-3 definition is available today in both parallel and serial modes. Parallel SCSI-3, also known as Ultra SCSI, can transmit up to 40MB/sec with 16 bits in parallel. An advantage of Ultra SCSI is that it can also support SCSI-2 devices up to the maximum speed of 20MB/sec. SCSI-3 also supports the following serial interfaces:

- **Fibre Channel** is a very high-speed, 100MB/sec protocol that can be used for networks as well as storage.
- **Serial Storage Architecture** is a serial interface available on IBM systems. It differs from Fibre Channel Standard (FCS) in that it has a node-to-node architecture and speeds which are currently 20MB/sec between nodes.
- **IEEE P1394** is a low-cost serial connection for desktop systems. It includes power so that several devices can be attached to a PC without the need for separate data and power cables. IEEE P1394 is also known as "Fire Wire."
- **Universal Serial Bus** (USB) is also a low-cost serial attachment for the desktop which will allow any device with a USB interface to attach to a single port.

Of all these serial standards, Fibre Channel Standard is the best enabling technology to make future storage and network applications and architectures a practical reality.

Fibre Channels combine the best of channels and networks

A channel provides point-to-point connectivity. Channels are hardware intensive, allowing data transport at the highest speed and the lowest latency, with error correction. Channels are usually limited in addressing and distance. Networks provide a shared service between many addresses and are designed to handle unpredictable burst traffic across long distances. Networks are very software intensive; they have a high latency and provide less reliability than channels. Fibre Channel combines the best attributes of a channel and a network to provide storage/server and high-performance networking environments with high bandwidth, low latency, scalability, flexiblility, and highly reliable data transfer.

Even though I/O interfaces have used parallel bus interfaces to improve data rates—going to 16 parallel data bits to get 20MB/sec with SCSI-2—problems with transmission skew and electrical interference require additional signals for clocking, shielding, and detection. It requires 68 wires to reliably transmit 16 bits of data in parallel across a maximum distance of 25m (82ft). In a serial fibre, bits are sent serially one bit at a time. This results in simplification of cables and connectors and improved distances, since skew is not a problem. The use of optical fibre cables makes it transmission-insensitive to external electronic noise.

Unlike SCSI-2, which transmits an entire block of data at a time, Fibre Channel divides the data transfer into frames, and interleaves the frames during transmission so that one I/O does not dominate the I/O path. The frames also include a header and trailer which contain the addresses of the source and destination ports, link control information, and cyclical redundancy checks, much like packets in network protocols. The data transfer portion, called the "payload," contains a maximum of 2112 bytes, which is encoded in an 8/10-bit scheme. This scheme encodes 8-bit data bytes into 10-bit characters to improve the transmission of characters in a serial data stream. Encoding the data prior to transmission facilitates the successful recovery of the data at the receiver. This is the same encoding used with ESCON.* Table 1 compares Fibre Channel characteristics with the best that SCSI-2 can provide.

Characteristics	**SCSI-2 Fast/Wide Differential**	**Fibre Channel**
Data rate	20MB/sec	100MB/sec per link
Data transmission	Half duplex	Full duplex
Distance	25m (82ft)	10km (6.2mi) per link
Connectivity	15 targets with 1 initiator	Unlimited switching
Configuration	String	Point-to-point, loop, fabric
Scalability	None	Loops and switches
Protocol support	SCSI	SCSI, IP, FDDI, ATM, others
Error handling	Parity checking	8/10-bit encoding detects and corrects
Transmission media	68 parallel copper wires	Single optical cable or one copper twisted pair

For more information on Fibre Channel definitions, see http://www.ancor.com

Table 1: Fibre Channel Compared with SCSI-2

Fibre Channel configurations

Fibre Channel consists of nodes, which act as transmitters or receivers, connected by a link that consists of two unidirectional fibres transmitting in opposite directions. (See Figure 1.) There are three basic configurations for Fibre Channel: point-to-point, loop, and switch or fabric.

- **Point-to-point** is the simple connection of two node ports called N_Ports.
- A **loop** can be considered a poorman's switch. It allows multiple nodes to be connected into a loop. However, only one conversation can occur at a time. A node must "arbitrate" for use of the loop before it can communicate with another node. A loop is referred to as a Fibre Channel Arbitrated Loop or FC-AL. If a loop is broken or one of the nodes fails, none of the other nodes on the loop can communicate. For redundancy, FC-ALs are configured as a pair of loops. Nodes on a loop are called L_ports. Additional fault tolerance can be provided with a **hub**. A hub allows the loop to remain intact by closing the connections to the failed node. A hub can also allow nodes to be plugged into or out of the loop without breaking communications to the rest of the nodes. An FC-AL can address up to 126 nodes.
- A **switch** extends the concept of a hub by allowing multiple paths to each loop. Switches may also be connected to other switches. This combination of one or more switches is called a "**fabric**," which allows multiple paths to a node. Nodes on the fabric are called F_Nodes. Many conversations can take place simultaneously across a fabric or switch. The fabric also provides path redundancy. A fabric might be compared to a telephone exchange, where a caller is connected to another phone by area code and phone number. The actual connection between the two phones could be different on each call. An FCS fabric can address up to 16 million nodes.

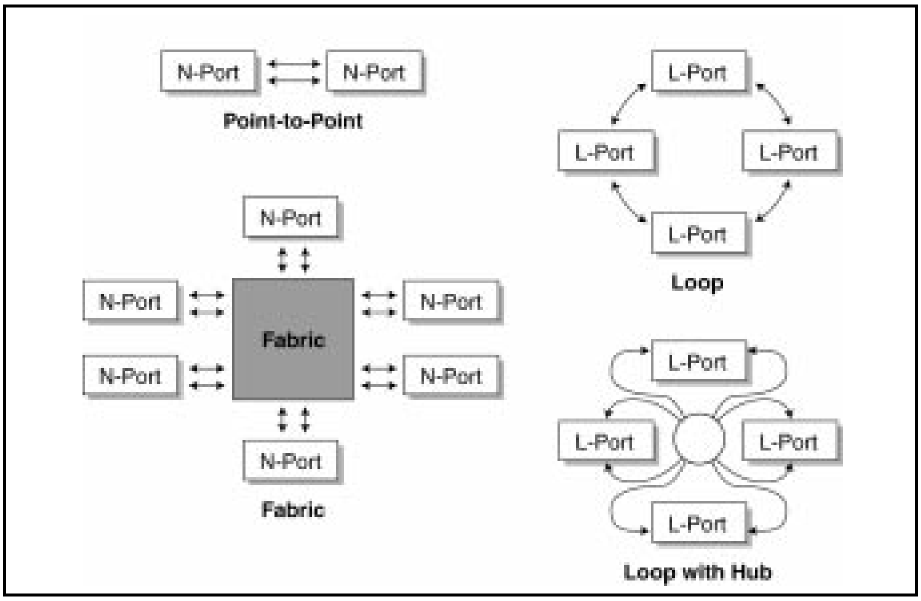

Figure 1: Fibre Channel Topologies

Employing fibre optic links

Fibre Channel can use metal or optical cables to transmit data. The optimum use of Fibre Channel is with optical fibre. Information is transmitted by sending pulses of light through the optical fibre. Turning the light on sends a 1 bit and turning the light off is detected as a 0 bit. The rate at which the light can be turned off and on and the capability of the optical fibre determine the data rate. The distance capability of the fibre is determined by how much light is injected into the fibre, how much degradation occurs to the pulses in transmission, and the detection capabilities of the receiver.

There are two common light sources: lasers and Light Emitting Diodes (LEDs). Lasers are generally used in high-speed applications and longer distances due to their faster switching capabilities and superior optical qualities. Lasers can be longwave in the infrared spectrum at 1300 nanometers, or shortwave at 780 nanometers. LEDs are used in lower speed applications (less than 25MB/sec) because of their lower costs. ESCON was implemented with LEDs.

Optical fibres are superior to copper transmission lines because they are insensitive to Electromagnetic Interference, and can provide greater distance. Copper fibre has the same distance limitation as SCSI-2 cables at 20m to 25m (65.6ft to 82ft). Optical fibre is more compact, and not susceptible to cross talk. However, optical fibre tends to be more expensive over short distances, can be affected by dirt and other contamination, and does not lend itself to backplane printed circuit wiring.

There are two types of optical fibre: single mode and multimode. In single mode, all the light is constrained to follow the same path, which has a core diameter of 9 micrometers (0.0000027ft). In multimode, the core diameter is much bigger—50 or 62.5 micrometers. This wider core results in multiple propagation paths. Some light takes one path and other light takes another path, resulting in a spreading of the pulse, which limits distance and data rate. A 62.5-micrometer multimode is cheaper and is commonly installed in existing ESCON and FDDI installations. A 50-micrometer multimode is preferred for new installations, as it creates less dispersion and travels longer distances. Single-mode fibre with longwave laser is required to get to 10km (6.2mi.) with a 100MB/sec data rate.

Table 2 summarizes the optical variants. Actual distances and data transmission rates will vary with transmitters, receivers, cable, plant, and loss budget (the difference between the transmitter and receiver). For a specific review of optical variants, consult your Fibre Channel vendor or the Fibre Channel Association web page at http://www.fibrechannel.com.

Media Type	**Transmitter**	**Speed**	**Distance**
Single-mode fibre	1300nm longwave laser	100MB/sec	2m (6.56ft) to 10km (6.2mi.)
Multimode fibre 50μm	780nm shortwave laser	100MB/sec	2m (6.56ft) to 500m (1,640ft)
Multimode fibre 50μm	1300nm LED	25MB/sec	2m (6.56ft) to 1.5km (0.93mi.)
Multimode fibre 62.5μm	780nm shortwave laser	100MB/sec	2m (6.56ft) to 300m (989ft)
Multimode fibre 62.5μm	1300nm LED	25MB/sec	2m (6.56ft) to 1.5km (0.93mi.)

Table 2: Summary of Optical Variants

Where is Fibre Channel today?

Today, over 200 vendors offer Fibre Channel products: host adapter boards, bus interconnects, software drivers, switches and hubs, peripherals and applications. The earliest adopters of Fibre Channel included companies in the film, video, broadcast, avionic, and publishing industries. In the storage industry, HDS, Clariion, Boxhill, Symbios, Sun, HP, and IBM have announced serial fibre products. Most products provide an upper interface attachment to host processors and retain SCSI-2 attachment for the back end because of the current cost and performance levels of Fibre Channel chipsets and firmware. Although FCS has a 100MB/sec instantaneous data rate, firmware and other overheads create sustained data rates for Fibre Channel implementations of about 30MB/sec to 60MB/sec for storage subsystems. Fibre Channel is usually implemented on the front end first, where distance, connectivity, and addressability are more important.

In terms of storage systems, Fibre Channel is considered to be in the early adopter phase of Geoffrey Moore's Technology Adoption Life Cycle model. Storage servers must provide Fibre Channel adapters to attach Fibre Channel storage subsystems. Sun and IBM have proprietary serial subsystems. Sun uses a quarter-speed Fibre Channel, while IBM uses Serial Storage Architecture. By year-end 1998, it is likely that all major storage and server vendors will have implemented Fibre Channel interconnects.

HBAs, hubs, switches, and edge routers

In addition to fibre ports on storage subsystems and servers, there will be a need for interconnect products similar to those found in network environments—such as host bus adapters (HBA), hubs, switches, and routers. HBA provides the attachment to a processor bus. Hubs allow multiple FCS nodes to be connected in a loop with the added protection of circuitry that can bypass a failed node on the loop. A switch provides multiple connections between FCS nodes so that multiple FCS conversations can take place between different nodes on the switch. Switches form the basis for a fabric (as discussed earlier in the "Fibre Channel Configurations" section). They are also the basis for establishing Storage Area Networks (SAN). "Edge router" is a term applied to a connection that allows different protocols to connect to each other, such as Copper SCSI to FCS fibre. This type of converter connection is sometimes called a MUX or bridge. Table 3 lists some of the vendors who offer products in this area.

Host Bus Adapters	**Converters/ Routers**	**Hubs**	**Switches**
Emulex	Crossroads	gadzoox	Brocade
Jaycor	Ancot	Vixel	McData
Tachyon		Emulex	gadzoox
			Ancor

Table 3: Connectivity Products and Their Vendors

Storage Area Networks

SAN was identified by *Byte Magazine*, December 1997, as one of the top new emerging technologies of 1997. It is generally defined as a data communications platform which interconnects servers and storage at Gigabaud speeds. By combining LAN networking models with fibre fabrics for mass storage capacities and high I/O performance, SAN eliminates the bandwidth bottlenecks and scalability limitations imposed by previous bus-based SCSI-2 architectures, and offloads the transfer of data from the LAN networks. SAN benefits include:

- Bandwidth efficiency to accelerate backup
- Cable lengths of 10km (6.2mi.) to facilitate remote operations
- Distributed hierarchical storage management capabilities
- More efficient use of LAN connections by separation of server/storage connections
- Scalable growth of servers and storage
- Fault tolerance and manageability.

This year may well mark the beginnings of a paradigm shift from bus-attached storage to network-attached storage for large subsystems. Fibre Channel fabrics offer low latency network connections which will help to create new standards for the capacity, performance, interconnects, and scalability required to meet increasing storage demand.

Hitachi Data Systems—For computing as critical as your business

Fibre Channel is an emerging technology and, as such, creates a whole new realm of opportunities. HDS intends to take advantage of these opportunities by assuming a leadership role in the Fibre Channel-based information-centric network market. By continuing to explore Fibre Channel technology, developing innovative ways to connect high-speed storage to LANs, and partnering with best-of-breed vendors, HDS will provide added value for businesses that are increasingly dependent upon data-intensive applications. Currently, with Fibre Channel technology as one of the core hubs, HDS is building a series of solutions that will fulfill our customers' business requirements as well as satisfy their IT challenges.

- **Centralized information management.** HDS' package of hardware and software components provides the tools needed to *manage the flow of data* throughout the enterprise and to *ensure the integrity of that data.*
- **Global information access.** An integrated series of components—services, cables, switches, host bus adapters, storage devices, and training—*delivers enterprise-wide information access* based on a Fibre Channel-networked information infrastructure.
- **Nonstop business processing.** To provide 100 percent availability of information within an enterprise, HDS carefully matches a suite of hardware and software tools that meet *business continuance* and disaster recovery requirements.
- **Business agility.** A mix of compute platforms and storage devices makes it difficult to compete in a fast-moving global market. Fibre Channel is now part of an enterprise-wide solution that makes *information rapidly available* to all users, supports *shorter time-to-market* for new products, and *enables faster customer service.*

Fibre Channel technology is available today on the Hitachi Freedom 5700E, and will be available on the Hitachi Freedom 7700E in the third quarter of '98. HDS is an active participant in the Fibre Channel Association and the Storage Networking Industry Association.

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

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 are conditioned on a maintenance contract with HDS being in effect, and 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.