

An Introduction to Fibre Channel

Fibre Channel is a flexible, scalable, high-speed data transfer interface that can operate over a variety of both copper wire and optical fiber at data rates up to 250 times faster than existing communications interfaces. Networking and I/O protocols, such as SCSI commands, are mapped to Fibre Channel constructs, encapsulated, and transported within Fibre Channel frames.

by Meryem Primmer

Fibre Channel is a standard, efficient, generic transport mechanism whose primary task is to transport data at the fastest speeds currently achievable with the least possible delay. It is a flexible, scalable method for achieving high-speed interconnection, communication, and data transfer among heterogeneous systems and peripherals, including workstations, mainframes, supercomputers, desktop computers, and storage devices. It handles both networking and peripheral I/O communication over a single channel using the same drivers, ports, and adapters for both types of communication.

Fibre Channel began in the late 1980s as part of the IPI (Intelligent Peripheral Interface) Enhanced Physical Project to increase the capabilities of the IPI protocol. That effort widened to investigate other interface protocols as candidates for augmentation. The first year of the project was spent looking for existing implementations to adopt, but none were found to be sufficient. The focus then changed to develop a new implementation. That implementation became Fibre Channel. Fibre Channel was approved as a project in 1988 by ANSI X3T9.

During the first year of investigation the ANSI working group decided to adopt a serial rather than a parallel bus interface. IBM's 8B/10B encode/decode scheme was adopted, and a decision was made to support both copper cable and optical fiber. Copper can be used for low cost while optical fiber can be used for distance. *Fibre* is a generic term used by the Fibre Channel standard to refer to all the supported physical media types.

The first draft of the Fibre Channel standard was developed in 1989. The standard addresses the need for very fast transfers of large volumes of data, while at the same time relieving systems of the need to support the multitude of channels and networks currently in use. The Fibre Channel standard covers networking, storage, and data transfers. In October 1994 the Fibre Channel physical and signaling interface standard, FC-PH, was approved as ANSI standard X3.230-1994.

Fibre Channel is structured as a set of hierarchical functions that support a number of existing protocols, such as SCSI (Small Computer System Interface) and IP (Internet Protocol), but it does not have a native I/O command set. It is not a high-level protocol like SCSI, but does contain a low-level protocol for managing link operations. Fibre Channel is not aware of, nor is it concerned with the content of the user data being transported. Networking and I/O protocols, such as SCSI commands, are mapped to Fibre Channel constructs and encapsulated and transported within Fibre Channel frames. The main purpose of Fibre Channel is to have any number of existing protocols operate over a variety of physical media and existing cable plants.

Fibre Channel is a high-speed data transfer interface that can operate from 2.5 to 250 times faster than existing communications interfaces. Its performance is both scalable and extendable and it supports multiple cost/performance levels, from small configurations such as disk arrays and low-cost, low-performance I/O devices and small systems to high-performance supercomputers and large distributed systems.

Fibre Channel runs at four speeds (actual data throughput): 100 megabytes per second (Mbytes/s), which translates to 1062.5 megabaud, 50 Mbytes/s or 531.25 megabaud, 25 Mbytes/s or 265.625 megabaud, and 12.5 Mbytes/s or 132.812 megabaud. A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw throughput. Fibre Channel provides a total network bandwidth of about one gigabit per second.

Fibre Channel operates over a variety of both copper wire and optical fiber at scalable distances, as shown in Table I. Distances are easily extendible using repeaters or switches.

Fibre Channel provides full duplex operation with separate transmit and receive fibers.

Another advantage of Fibre Channel is that it uses small connectors. The serial connectors used for Fibre Channel are a fraction of the size of SCSI parallel connectors and have fewer pins, thereby reducing the likelihood of physical damage. Also, depending on the topology, many more devices can be interconnected on Fibre Channel than on existing channels.

Topologies

Fibre Channel can be implemented in three topologies to interconnect varying numbers of devices, called *nodes* in Fibre Channel terminology. The topologies are point-to-point, arbitrated loop, and crosspoint switched, or *fabric* (a Fibre Channel term for a network of one or more switches connecting multiple nodes). Nodes contain one or more ports, such as an I/O adapter, through which they communicate over Fibre Channel. A generic node port is called an *N_Port*. The connections between ports are called *links*.

Table I
Fibre Channel Media, Data Rates,
Distances, and Transmitters

Media Type	Data Rate (Mbytes/s)	Maximum Distance	Transmitter Type
150-ohm Twinax or STP	100	30 m	ECL
	50	60 m	ECL
	25	100 m	ECL
75-ohm Video Coax	100	25 m	ECL
	50	50 m	ECL
	25	75 m	ECL
	12.5	100 m	ECL
75-ohm Miniature Coax	100	10 m	ECL
	50	20 m	ECL
	25	30 m	ECL
	12.5	40 m	ECL
105-ohm Type-1 Shielded Twisted-Pair Electrical	25	50 m	ECL
	12.5	100 m	ECL
62.5- μ m Multimode Optical Fiber	100	300 m	SW Laser
	50	600 m	SW Laser
	25	1 km	LW LED
	12.5	2 km	LW LED
50- μ m Multimode Optical Fiber	100	500 m	SW Laser
	50	1 km	SW Laser
	25	2 km	SW Laser
	12.5	10 km	LW LED
9- μ m Single-Mode Optical Fiber	100	10 km	LW Laser
	50	10 km	LW Laser
	25	10 km	LW Laser

ECL = Emitter-Coupled Logic, LW = Longwave, SW = Shortwave, LED = Light-Emitting Diode, STP = Shielded Twisted-Pair

Point-to-point (Fig. 1) is a direct channel connection between two *N_Ports*, typically between a processor and a peripheral device controller. In this topology exactly two devices are connected together. No fabric elements exist and no fabric services, such as name mapping, are necessary. Point-to-point is the default topology.

Fibre Channel arbitrated loop, or FC-AL, is a method for interconnecting from two to 126 devices through attachment points called *L_Ports* in a loop configuration. *L_Ports* can consist of I/O devices and systems of various performance levels. FC-AL is a low-cost solution because it does not require hubs and switches. FC-AL is a good choice for small to medium-sized configurations and provides an upward growth path by interconnecting the loop with a fabric through an *FL_Port*. Arbitrated loop is the most common Fibre Channel topology.

Fig. 2 shows the Fibre Channel arbitrated loop topology. A *private loop* (Fig. 2a) consists only of nodes, called *NL_Ports*, and does not connect with a fabric. A *public loop* (Fig. 2b) connects with a fabric via an *FL_Port*. A *disk loop* uses the loop topology to interconnect a number of high-performance disks, for example, a RAID (Redundant Array of Inexpensive Disks) device. Fig. 3 shows an office configured in a public arbitrated loop topology, and Fig. 4 shows a private disk loop.

All devices on the arbitrated loop share the bandwidth of the loop and the management of the loop. No dedicated loop master exists, and any node is capable of being the loop master. Which node performs the loop master functions is negotiated when the loop is initialized.

Each node has equal opportunity to communicate with another node by arbitrating for temporary ownership of the loop. An arbitration scheme using a fairness algorithm is used to establish a circuit between two *NL_Ports* on the loop before they

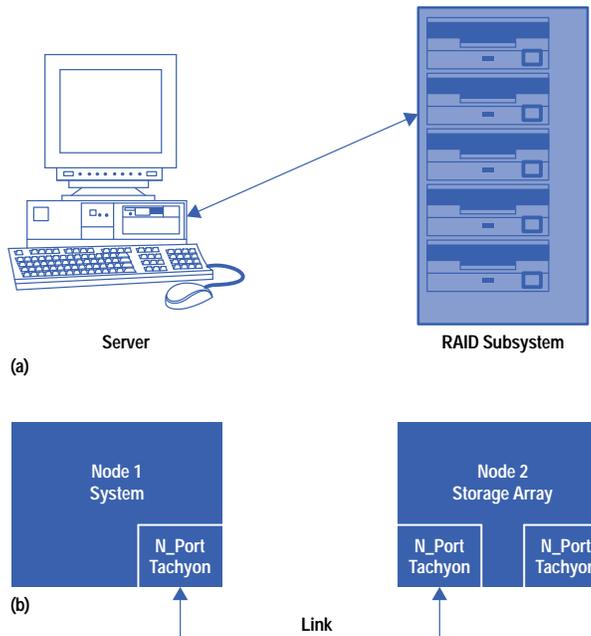


Fig. 1. (a) Two devices connected point-to-point. (b) Fibre Channel point-to-point topology. Tachyon is HP's gigabit Fibre Channel controller chip.

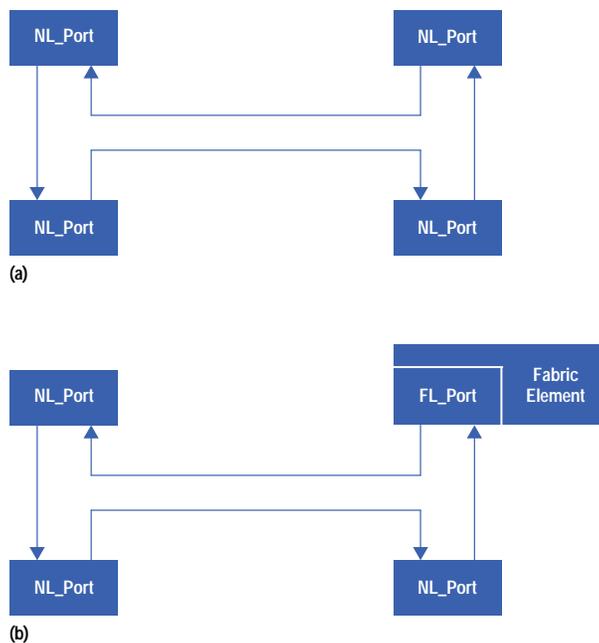


Fig. 2. Fibre Channel arbitrated loop topology. (a) Private Loop. (b) Public loop.

can communicate. Only one communication, or loop circuit, can be active at a time. After relinquishing the loop, an NL_Port cannot win arbitration again until all other arbitrating ports have had their turn.

The third Fibre Channel topology is crosspoint switched, or fabric. Fig. 5 shows a generic fabric topology, and Fig. 6 shows the Fibre Channel fabric topology with a single switching or fabric element.

A fabric topology is implemented as one or more switching elements. A fabric appears as a single entity to attached nodes, called *F_Ports*, even though the fabric can consist of multiple switches. Typically, a switch has from four to 16 *F_Ports* attached to it. In theory, there is no size limit to the number of nodes that can interconnect in a fabric, but addressing space limits the number to a maximum of 2^{24} . The fabric topology is good for interconnecting large numbers of devices and complex configurations.

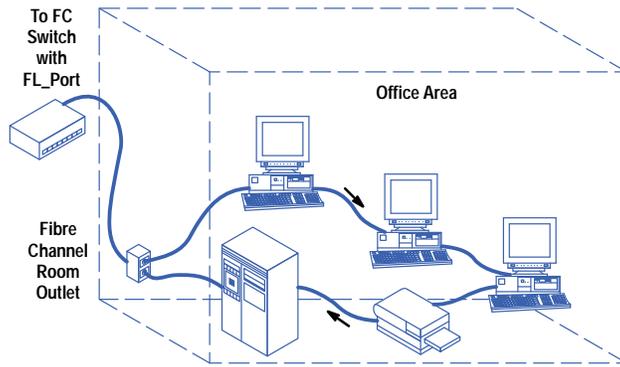


Fig. 3. An office configured in a public arbitrated loop topology.

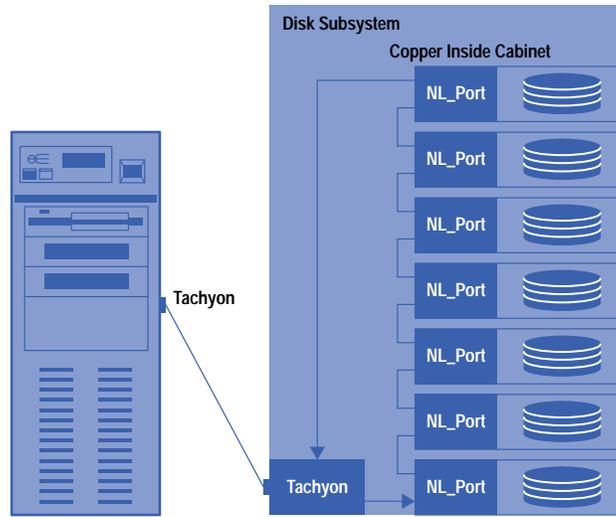


Fig. 4. A private disk loop.

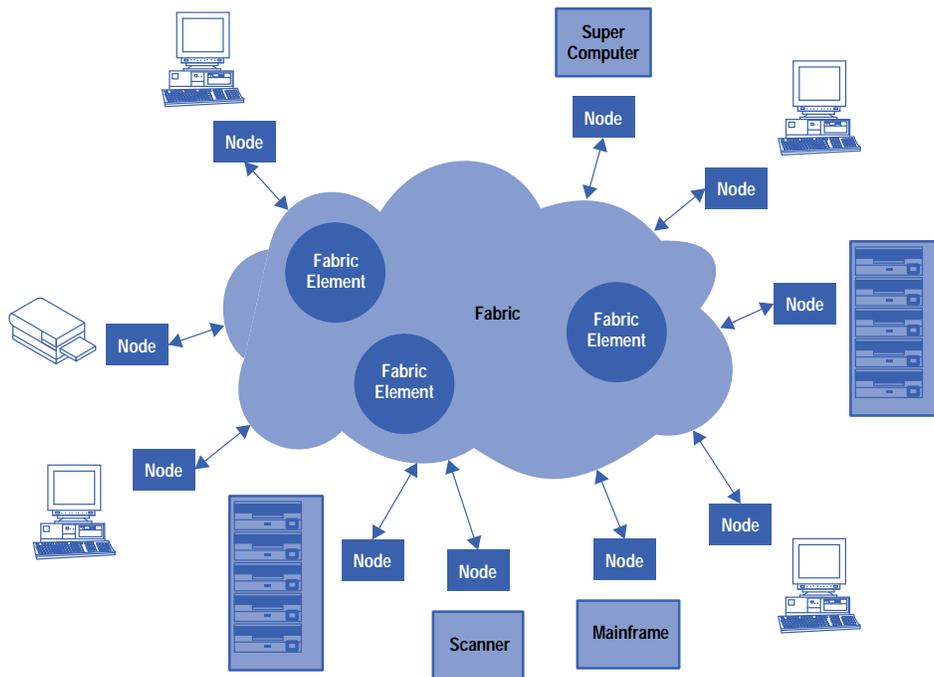


Fig. 5. A generic fabric topology.

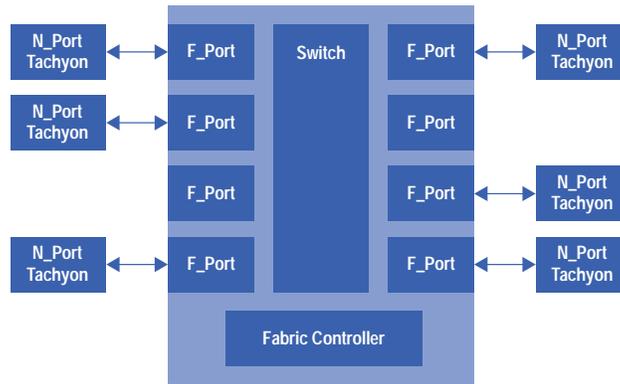


Fig. 6. Fibre Channel fabric topology with a single switching element.

The structure and operations of the fabric are transparent to the F_Ports attached to it. The fabric topology is self-managed, with the fabric performing station management functions and the routing of frames. Each port only needs to manage a point-to-point connection between itself and the fabric.

A Layered Approach

Fibre Channel is structured as a set of five hierarchical functional levels (see Fig. 7). The user protocol being transported over the Fibre Channel—SCSI or IPI (Intelligent Peripheral Interface), for example—is known as the *upper level protocol* (ULP) and is outside the scope of the Fibre Channel layers. The Tachyon Fibre Channel protocol chip described in **Article 12** implements the FC-1 and FC-2 layers, which are shaded in Fig. 7. Tachyon also implements SCSI assists and IP checksumming, shown as shaded boxes at the FC-4 level.

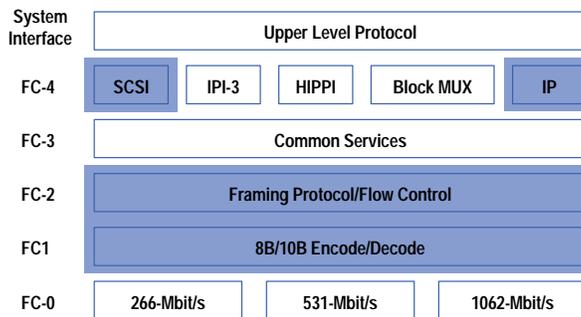


Fig. 7. Fibre Channel's five layers.

FC-4: The Protocol Mappings Layer. This topmost Fibre Channel level defines the mapping of the ULP interfaces to the lower Fibre Channel levels. Fibre Channel supports multiple existing protocols, including SCSI, IP, and IPI. Each ULP supported by Fibre Channel requires a separate FC-4 mapping and is specified in a separate FC-4 document. For example, the Fibre Channel protocol for SCSI, which is known as FCP, defines a Fibre Channel mapping layer that uses the services of the lowest three Fibre Channel layers to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target. ULPs are not tied to a particular physical medium or interface. For example, SCSI is supported without requiring a SCSI bus.

FC-3: The Common Services Layer. Nodes can be computer systems or peripheral devices. The FC-3 level defines a set of services that are common across multiple ports of a node. The FC-3 layer is still being formulated in the ANSI committee and no functions have been formally defined.

FC-2: The Framing Protocol Layer. This level defines the signaling protocol, including the frame and byte structure, which is the data transport mechanism used by Fibre Channel. Included in this level is the framing protocol used to break sequences into individual frames for transmission, flow control, 32-bit CRC generation, and various classes of service.

The FC-2 layer also handles hardware disassembly and reassembly of sequences of data. Defined in this layer are a few built-in command primitives, called *ordered sets*, for handling such functions as configuration management, error recovery, frame demarcation, and signaling between two ends of a link.

A *frame* (Fig. 8) is the smallest indivisible unit of user data that is sent on the Fibre Channel link. Frames can be variable in length, up to a maximum of 2148 bytes long. Frame size depends on implementation, not hardware or software. Each frame contains a four-byte Start of Frame delimiter, a 24-byte header, up to 2112 bytes of FC-4 payload consisting of zero to 64 bytes of optional headers and zero to 2048 bytes of ULP data, a four-byte CRC, and a four-byte End of Frame delimiter.

4 Bytes	24 Bytes	0 to 2112 Bytes		4 Bytes	4 Bytes
Start of Frame	Frame Header	FC-4 Data Payload		CRC	End of Frame
		0 to 64 Bytes	0 to 2048 Bytes		
		Optional Headers	Data Payload (e.g., IP Packet, SCSI Command)		

Fig. 8. A Fibre Channel frame.

A *sequence* is a set of one or more related frames. For example, a large file transfer would be accomplished in a sequence consisting of multiple frames.

An *exchange* contains one or more sequences. It is comparable to a SCSI I/O process, and is the mechanism for coordinating the exchange of information between two communicating N_Ports in a single operation.

In general, the sequence is the Fibre Channel error recovery boundary. That is, selective retransmission of frames for error recovery is not supported in the Fibre Channel physical and signaling interface, FC-PH. If an error is detected in a transmitted frame and the error policy requires error recovery, the sequence to which the frame belongs may be retransmitted.

Fibre Channel provides three classes of service, which are managed by the FC-2 layer. Class 1 dedicated connection service provides a dedicated or circuit-switched connection between two N_Ports. The connection must be established before communication can begin and must be torn down when communication is completed. Class 1 guarantees delivery of frames in the order in which they were transmitted. Confirmation of delivery also is provided. Class 2 multiplex service provides a connectionless, frame-switched link. Delivery is guaranteed, but not necessarily in order if multiple routes exist through the fabric. Class 2 also provides acknowledgement of receipt. Class 3 datagram service is a connectionless service similar to class 2, but without confirmation of receipt. Neither delivery nor receipt order is guaranteed in class 3.

FC-1: The Encode/Decode Layer. This layer defines the transmission protocol, including the 8B/10B encode/decode scheme, byte synchronization, and character-level error control. 8B/10B is a dc-balanced encode/decode scheme that provides good transition density for easier clock recovery and character-level error detection. In this scheme, 8-bit internal bytes are encoded and transmitted on the Fibre Channel link as 10-bit transmission characters. The transmission characters are converted back into 8-bit bytes at the receiver. Using 10 bits for each character provides 1024 possible encoded values rather than only the 256 values that are possible for 8-bit characters. Not all of the 1024 possible values are used. To maintain a dc balance on the link, only those that contain four zeros and six ones, six zeros and four ones, or five zeros and five ones are used. Some of the extra 10-bit characters are used for low-level link control. One special character called a *comma* is used for byte synchronization.

FC-0: The Physical Layer. FC-0, the lowest of the five levels, defines the physical characteristics of the media, including cables, connectors, drivers (ECL, LEDs, shortwave lasers, longwave lasers, etc.), transmitters, transmission rates, receivers, and optical and electrical parameters for a variety of data rates and physical media. Reference 1 describes HP products that implement the FC-0 layer.

Collectively, the three lowest layers constitute the Fibre Channel physical and signaling interface, FC-PH. FC-PH is a channel/network hybrid. It supports channel interfaces for peripherals—for example, SCSI, IPI, and HIPPI (High-Performance Parallel Interface)—as well as network protocols such as TCP/IP. FC-PH is similar enough to a network to gain connectivity, distance, and serial interfaces, while being enough like an I/O channel to retain simplicity, reliability, and hardware functionality.

Reference

1. J.S. Chang, et al, "A 1.0625-Gbit/s Fibre Channel Chipset with Laser Driver," *Hewlett-Packard Journal*, Vol. 47, no. 1, February 1996, pp. 60-67.

Bibliography

1. *Fibre Channel—Physical and Signaling Interface (FC-PH)*, X3.230-1994, Rev. 4.3, American National Standards Institute.
2. *Fibre Channel—Arbitrated Loop (FC-AL)*, X3.272-199x, Rev. 4.5, American National Standards Institute, June 1995.
3. *Fibre Channel Protocol for SCSI (FCP)*, X3.269-199x, Rev. 012, American National Standards Institute, May 30, 1995.
4. *Fibre Channel: Connection to the Future*, The Fibre Channel Association, 1994.
5. The Fibre Channel Association server URL: <http://www.amdahl.com/ext/CARP/FCA/FCA.html>

-
- ▶ [Go to Article 12](#)
 - ▶ [Go to Table of Contents](#)
 - ▶ [Go to HP Journal Home Page](#)