

## Fibre Channel: A connection to the future

Clint Jurgens, Ancor Communications

The phenomenal increase in processor speeds has significantly outstripped the I/O capabilities of today's communications channels. With commodity microprocessors on the horizon capable of processing hundreds of MIPS, current transmission rates cannot accommodate "Amdahl's Law"—where 1 Mbps of I/O is required for every MIPS of processing power—and will become a bottleneck to system performance in data-intensive applications.

To remedy this shortcoming, in 1988 ANSI committee X3T11 initiated development of Fibre Channel, a switched protocol capable of transmitting at rates exceeding one gigabit per second, while still supporting existing protocols over both optical fiber and copper cables. Fibre Channel combines the best attributes of legacy channels and networks into a single standard that is a generic transport mechanism for data, voice, and video. It is the key to scientific and business applications implemented in open and distributed architectures, because it removes the barriers to performance presented by the old methods of data communications.

Fibre Channel FC-PH standard X3.230-1994 was made official with formal adoption by ANSI and is now at the International Standards Organization (ISO) for ratification. Today, there are over 100 companies developing products that provide storage, networking, and peripheral I/O for local/campus data communications.

### Fibre Channel

Despite the name, the Fibre Channel architecture represents a true channel/network integration. A Fibre Channel port manages a simple point-to-point connection between itself and an intelligent interconnection scheme called a fabric. Transmission is isolated from the control protocol, so different topologies can be implemented as needed.

Fibre Channel consists of five multilayered functional levels that define the physical media and transmission rates, encoding scheme, framing protocol and flow control, common services, and the upper-level protocol interfaces:

- FC-0: physical layer,
- FC-1: transmission protocol,
- FC-2: framing protocol and flow control,
- FC-3: common services, and
- FC-4: protocol mapping.

Figure 1 shows the hierarchy of these functional levels.

**Physical interface and media.** The physical media and data rates described in FC-0 are conservative guidelines for industry. Table 1 illustrates the relationship between the media type and the operating range for each defined Fibre Channel data transfer rate. For optical fiber interfaces, Fibre Channel uses a special low-cost duplex SC connector to ensure polarization and ease of use. For the shielded twisted-pair (STP) copper interface, a 9-pin D-type connector is used. For the coaxial connections, a TNC receiver connector and a BNC transmitter connector have been chosen.

**Transmission protocol.** FC-1 defines the byte synchronization and the encode/decode scheme. Fibre Channel uses a dc-balanced 8B/10B code scheme, selected for its superior transmission characteristics. In this scheme, eight internal bits are transmitted as a 10-bit group. A unique special character, called a comma character, ensures proper byte and word alignment. Other advantages of this code are its useful error-detection capability and a simple logic implementation for both encoder and decoder.




Figure 1. Fibre Channel layers.

This article is a condensation of the Fibre Channel Association's Fibre Channel: Connection to the Future. Interested readers can contact the Fibre Channel Association at 12407 MoPac Expressway North 100-357, PO Box 9700, Austin, TX 78758; e-mail FCA-info@amcc.com.

**Framing, signaling, and flow control.** Fibre Channel is hardware intensive, using hardware entities called N\_ports (node port) within nodes such as mainframes or peripherals. Physically, the link consists of a minimum of two nodes, each having at least one N\_port interconnected by a pair of links, one outbound and one inbound. The FC-2 transport mechanism defines a set of building blocks to transparently move the payload contents. Payload contents are transparent to FC-2 and visible to FC-3 and FC-4.

*Ordered set.* An ordered set consists of four 10-bit characters—a combination of data characters and special characters—that are used to provide certain very low-level link functions, such as frame demarcation and signaling between two ends of a link. This signaling enables link initialization after power-ON and provides certain basic recovery actions.

*Frame.* A frame is the smallest indivisible packet of data that is sent on the link. The frame protocol/structure contains:

- a start-of-frame delimiter,
- a header,
- an optional header,

- a payload (from 0 to 2112 bytes),
- a 32-bit CRC (cyclic redundancy check) error check for header and payload, and
- an end-of-frame delimiter.

*Sequence.* Fibre Channel places no limit on the size of the transfers between applications. A sequence is composed of one or more related frames for a single operation, flowing in the same direction on the link. It is the FC-2 layer's responsibility to break a sequence into the frame size that has been negotiated between the communicating ports and between the ports and the switching fabric. The sequence is also the recovery boundary in Fibre Channel. When an error is detected, Fibre Channel identifies the errant sequence and lets that sequence (and subsequent sequences) be retransmitted. This unique capability enables Fibre Channel to provide reliable data transfer without relying on an upper layer protocol like TCP/IP. It is also what makes Fibre Channel very desirable for storage and peripheral I/O, where normally, no upper layer protocol is used to provide reliable data transfer. Each sequence is uniquely identified by the initiator of that sequence.

*Exchange.* An exchange is composed of one or more non-concurrent sequences for a single operation. For example, an operation may consist of a command to read data, fol-

**Table 1. Relationship between the media type and the operating range for each defined Fibre Channel data transfer rate.**

Data transfer rate	Operating range	Speed	Delivery media
<b>9-micron single-mode fiber</b>			
100 MB/s	up to 10 Km	1062.5 Mbps	Longwave laser
50 MB/s	up to 10 Km	531.25 Mbps	Longwave laser
25 MB/s	up to 10 Km	265.625 Mbps	Longwave laser
<b>50-micron multimode fiber</b>			
50 MB/s	up to 1 Km	531.25 Mbps	Shortwave laser
25 MB/s	up to 2 Km	265.625 Mbps	Shortwave laser
<b>62.5-micron multimode fiber</b>			
25 MB/s	up to 500 m	265.625 Mbps	Longwave LED
12.5 MB/s	up to 1 Km	132.8125 Mbps	Longwave LED
<b>Video coax</b>			
100 MB/s	up to 25 m	1062.5 Mbps	ECL
50 MB/s	up to 50 m	531.25 Mbps	ECL
25 MB/s	up to 75 m	265.625 Mbps	ECL
12.5 MB/s	up to 100 m	132.8125 Mbps	ECL
<b>Miniature coax</b>			
100 MB/s	up to 10 m	1062.5 Mbps	ECL
50 MB/s	up to 20 m	531.25 Mbps	ECL
25 MB/s	up to 30 m	265.625 Mbps	ECL
12.5 MB/s	up to 40 m	132.8125 Mbps	ECL
<b>Shielded twisted-pair</b>			
25 MB/s	up to 50 m	265.625 Mbps	ECL
12.5 MB/s	up to 100 m	132.8125 Mbps	ECL

lowed by the data, then the completion status of the operation. Each phase is a separate sequence, but the sequences form a single exchange. This is an example of multiplexing supported by Fibre Channel. The exchange is uniquely identified by each participating N\_Port.

#### Fibre Channel upper layers

Fibre Channel also defines upper layer functions that provide common services and seamless integration of existing standards.

**Common services.** FC-3 is concerned with functions that span multiple N\_Ports, including the following:

- *Striping.* This is a method for achieving higher bandwidth. It uses multiple N\_Ports in parallel to transmit a single information unit across multiple links and permits such links simultaneously.
- *Hunt groups.* A hunt group is a set of associated N\_Ports attached to a single node. This set is assigned an alias identifier that lets any frames containing this alias be routed to any available (nonbusy) N\_Port within the set. This improves efficiency by decreasing the chance of reaching a busy N\_Port.
- *Multicast.* Multicast delivers a single transmission to multiple destination N\_Ports. This includes sending to all N\_Ports on the fabric or to a subset of N\_Ports on the fabric.

**Protocol profiles.** FC-4 Fibre Channel is equally adept at transporting video, channel, and network information. FC-4 provides a common and interoperable method for implementing these different protocols over Fibre Channel. Some of the protocols Fibre Channel supports are Small Computer Systems Interface (SCSI); Intelligent Peripheral

Interface (IPI); High-Performance Parallel Interface (HiPPI); Internet Protocol (IP); Link Encapsulation (LE); and Single Byte Command Code Set (SBCCS).

FC-4 profiles use the FC-PH physical and signaling protocols to transfer information. Through mapping rules, a specific FC-4 describes how processes of the same FC-4 interoperate.

#### Fabrics

Fibre Channel employs a switching fabric to connect devices. Fibre Channel fabrics relieve each N\_Port of the responsibility for station management. There are no complex routing algorithms to compute to send data to another port. The calling port simply does the equivalent of dialing a phone number. A fabric has an address space of  $2^{24}$  (more than 16 million addresses). A Fibre Channel fabric appears as a single interconnecting entity even though it may contain one or more fabric elements.

#### Classes of Service

Fibre Channel provides solutions to a wide variety of communications needs, as follows:

- Class 1 functions like current physical channels. It provides acknowledged connection service, guaranteed bandwidth and delivery, end-to-end flow control, and in-order delivery.
- Class 2 is a frame-switched, acknowledged connectionless service that provides guaranteed delivery and buffer-to-buffer flow control. If delivery cannot be made because of congestion, a busy signal is returned and the sender tries again.
- Class 3 is an unacknowledged connection service that lets data be sent rapidly to multiple devices attached to the fabric. Datagram transfers speed transmission by not requiring confirmation of receipt. They are useful for real-time broadcasts where timeliness is key and information not received has little value after the fact.
- Intermix is a concurrent Class 1, 2, and 3 service that enables parallel operation. It reserves full Fibre Channel bandwidth for dedicated Class 1 connection, but permits connectionless transmissions if bandwidth becomes available during idle Class 1 connections. A client can use connectionless service to set up a file transfer from one file server at the same time it is moving a file from a second server. Thus, the system can reduce the latency associated with file retrieval, resulting in more I/O operations per second.

#### Topologies

Fibre Channel supports point-to-point, switched, and loop topologies. The highest performance is derived from a switched topology, which lets users scale system bandwidth linearly by adding N\_Ports. The arbitrated loop topology provides a low-cost solution to attach multiple ports in a loop. As with other types of loops the bandwidth is shared. The loop is self-configuring and may operate with or without a fabric present. To ensure fairness, equal access to all N\_Ports is guaranteed.

FIBRE CHANNEL INTRODUCES HIGH-PERFORMANCE, easy-to-use, low-cost communications required by a new breed of processors and applications. Available today are new high-speed, scalable links to storage; high-performance networks enabling clusters, backbones, imaging, and visualization; and low-cost arbitrated loops providing efficient peripheral I/O. Fibre Channel is a "connection to the future."

*Clint Jurgens is vice president of Ancor Communications and is a member of the Board of Directors of the Fibre Channel Association. He can be reached via e-mail at [clintj@ancor.com](mailto:clintj@ancor.com).*

**Some of the protocols Fibre Channel supports are Small Computer Systems Interface (SCSI); Intelligent Peripheral Interface (IPI); High-Performance Parallel Interface (HiPPI); Internet Protocol (IP); Link Encapsulation (LE); and Single Byte Command Code Set (SBCCS).**