

FIBRE CHANNEL

By Robert C Burton, burton@cis.ohio-state.edu

Abstract

Fibre Channel is the general name of an integrated set of standards being developed by the **American National Standards Institute (ANSI)**. It is designed to significantly improve the speed at which data is transmitted between workstations, mainframes, supercomputers, storage devices and displays, while providing one standard for networking, storage and data transfer. In general Fibre Channel provides speed, reliability, distance, and connectivity; all at a low cost. It guarantees bandwidth and delivery while being able to handle any protocol or topology. This is a survey paper on Fibre Channel.

Table of Contents

- [Fibre Channel Overview](#)
 - [Fibre Channel Topologies](#)
 - [Fibre Channel Products](#)
 - [Fibre Channel Standards Organizations](#)
 - [Helpful Fibre Channel Links](#)
 - [Recent Events in Fibre Channel](#)
 - [References](#)
-

Fibre Channel Overview

The Universal Interface

Fibre Channel is designed to combine the best of both **channel** and **network** data communication. A channel is a direct or switched point-to-point connection. One benefit of a channel is that it is mostly hardware-intensive for speed and efficiency of data transport. This is compared to the higher overhead of network data transfer, which is slower because it is software-intensive. The benefit of networks is that they can handle a greater range of tasks because they operate in an environment of unanticipated connections. Fiber Channel avoids the downside of network data transfer by providing a way to transfer data from a buffer at the source device to a buffer at the destination device. Fibre Channel does not need to know what the data is or how it is formatted. What the individual protocols do with the data before or after it is in the buffer is independent of the function of Fibre Channel. All a Fibre Channel port has to do is to manage a simple point-to-point connection between itself and the Fabric. A Fabric is a **circuit switch** ([diagram](#)) , **active hub** ([diagram](#)) , or a **loop** ([diagram](#)) .

Fibre Channel is a switched medium that works similar to a telephone network: any user will have a temporary, direct connection that provides the option of the full bandwidth of the Fibre Channel as long as the connection is established. Fibre Channel's acknowledgment and flow control supports connection-less traffic by using time division multiplexing.

Fibre Channel is designed to transport many protocols, such as **FDDI**, **serial HIPPI**, **SCSI**, **IPI**, and many more that will be listed in the section describing the [FC-4](#) layer.

The connections provided for by Fibre Channel are:

- **High performance workstations**
- **Intelligent mass-storage systems**
- **Heterogeneous workstation clustering**
- **Parallel processing**
- **Load leveling**
- **Host-to-Host**
- **Host-to-mass-storage**
- **Server-to-mass-storage**
- **server-to-server**
- **Bulk data transfer**
- **Multi-media**
- **Printers linked to a LAN without a cpu**
- **Disk drives linked to a LAN without a cpu**

The transfer rates of Fibre Channel are currently (133 Mbps,266 Mbps,530 Mbps, and 1 Gbps). However, data rates of 2 to 4 Gbps should be available soon.

Fibre Channel will allow simultaneous transmission of different protocols over a single optical-fiber pair and it can allow a number of existing services, such as network, point-to-point, and peripheral interfaces, to be accessed over a single medium using the same hardware connection. Fibre Channel also provides control and complete error checking.

The Fibre Channel structure is defined as a multi-layered stack of functional levels, not unlike those used to represent network protocols, although not mapping directly to OSI layers. The layers of the Fibre Channel standard define the physical media and transmission rates, encoding scheme, framing protocol and flow control, common service, and the upper-level applications interfaces. The five layers are:

[FC-0](#),[FC-1](#),[FC-2](#),[FC-3](#),[FC-4](#)

[Return to table of contents](#)

FC-0, the lowest layer, specifies the physical link in the system, including the media, transmitters, receivers, and connectors that can be used with Fibre Channel. This also includes electrical and optical characteristics, transmission rates, and other physical components of the standard. The physical level is designed to be able to use a large number of technologies to meet the widest range of system

requirements. An end-to-end route can use different link technologies for increased performance and decreased cost, while at the same time, systems integrators can tailor an installation to meet the specific needs of their customers.

FC-0 also specifies the **(OFC) Open Fibre Control system**, which is a safety system used to control the optical power level of SW laser data links in which an open fibre condition occurs. This safety feature is required because the optical power levels in this kind of system exceed the limits defined by the laser safety standards. Whenever this open fibre condition occurs in the link from the sending port, the receiver port detects it and pulses its laser at a low duty cycle within the laser safety requirements. The receiver at the other port detects the pulsing signal and itself sends pulses within the specified laser safety range. If the open fibre condition is restored, receivers from both ports receive the pulsing signals which will result in a double handshaking procedure to restore normal transmission after a couple of seconds.

[Return to "fc-0" link](#)

[Return to table of contents](#)

FC-1 defines all transmission protocols including serial encoding and decoding rules, special characters and error control. Every 8 bits of data are encoded into a 10 bit **Transmission Character**. A **Transmission Word** is composed of four contiguous Transmission Characters. The transmission code is DC balanced and the Transmission character is used to ensure that clock recovery is possible by having enough transitions present in the serial bit stream.

Character conversion is accomplished by taking an unencoded information byte that is made up of eight information bits logically labeled A,B,C,D,E,F,G,H and a control character Z. Each binary value represented by E,D,C,B,A is converted into the form xx which is a decimal representation of the binary value, since $2^5-1=31$ requires at most two decimal digits. H,G similarly is converted into the form y which is a one digit decimal value. The control character Z is represented by D for data-type or K for special-type. The resulting combination of this information forms a name in the form Zxx.y that represents a valid Transmission Character. After transmission the D-type Transmission Characters are decoded into one of 256 eight-bit combinations. Any K-type Transmission Characters are used for protocol management functions. All other codes besides D or K types are invalid.

The physical layer uses a **(RD) Running Disparity**. This is a binary number that is calculated based on the number of 0's and 1's in the two **sub-blocks**. The first sub-block is the first six bits of the Transmission Character and the second sub-block is the last four bits of the Transmission Character. A new RD is calculated at the transmitter and receiver and if the the RD values are not the same a disparity violation condition is indicated.

[Return to "fc-1" link](#)

[Return to table of contents](#)

FC-2 is called the signaling protocol level. It is responsible to break the data to be transmitted into frame size, and reassemble the frames after transport. It specifies the framing rules of the data to be sent from one port to another, ways for controlling the three [service classes](#), and controlling the sequence in which

the data is transferred. All frames in a transfer have a sequence number from 0 to N so that the receiver is able to tell if a frame is missing and exactly which frame(s) is missing. There are five different building blocks defined to provide efficient means of data transport across a link:

- **Ordered Set**

Ordered Sets are Transmission Words that allow bit or word synchronization and word boundary alignment. The three types of Ordered Sets are:

- **Frame Delimiters**
 - (SOF) Start-of-Frame
 - (EOF) End-of-Frame
- **Primitive Signals**
 - Idle
 - (R_RDY) Receiver Ready
- **Primitive Sequence**
 - (OLS) Off-line
 - (NOS) Not Operational
 - (LR) Link Reset
 - (LRR) Link Reset Response

- **Frame**

The frames contain the data(payload), address of the source and destination ports, and the link control information. Each frame payload has a maximum 2112 byte long. Frames are broken down into:

- **Data frames**
 - Link data frames
 - Device data frames
- **Link control frames**
 - (ACK) Acknowledge
 - (Busy and Reject) Link response

[frame diagram by CERN](#)

- **Sequence**

A Sequence is a set of one or more related frames, each uniquely numbered for use by an upper level error recovery function.

- **Exchange**

An Exchange is one or more non-concurrent sequences for a single operation.

- **Protocol**

There are five different protocols for Fibre Channel not to be confused with upper level protocols that will be listed later.

- **Primitive Sequence Protocol**
-specified for link failure
- **Fabric Login Protocol**
-interchanging service parameters between an N_Port and a fabric
- **N_Port Login Protocol**
-interchanging service parameters between two N_Ports
- **Data Transfer Protocol**
-describes ways of transferring Upper Layer Protocols
- **N_Port Logout Protocol**
-N_Port wants to disconnect from another N_Port
- **Flow Control**

The type of Flow Control depends on which class is used (described next section)

- **class 1 frames**
-end to end
- **class 3 frames**
-buffer to buffer
- **class 2 frames**
-both end to end and buffer to buffer

[Return to "fc-2" link](#)

[Return to table of contents](#)

FC-3 is in the process of being defined for use by common services required for advanced features such as:

- **striping**
- to multiply bandwidth
- **hunt groups**
- the ability for more than one port to respond to the same alias address
- **Multicast**
- which delivers a single transmission to multiple destination ports

[Return to "fc-3" link](#)

[Return to table of contents](#)

FC-4 This layer defines the application interfaces that can be run over Fibre Channel. Both network and channel protocols can concurrently run over the same physical interface providing seamless integration of standards. A list of current or proposed protocols is shown below:

- (SCSI) Small Computer System Interface
- (IPI) Intelligent Peripheral Interface

- (HIPPI) High Performance Parallel Interface
- (IP) Internet Protocol
- (AAL5) ATM Adaptation Layer
- (FC-LE) Link Encapsulation
- (SBCCS) Single Byte Command Code Set Mapping
- IEEE 802.2


FCA connections diagram

Message Types--To accommodate a range of communications needs, Fibre Channel defines three different classes of service. Each class can be utilized between as few as two points, and each class can be integrated within the framework of other network protocols.

Class 1 is used when the need is for dedicated circuit-switched connections. This service guarantees the maximum bandwidth between two N_Ports and is good for high-throughput transactions.

Class 2 is a frame-switched, connection-less service where the bandwidth can be shared by different sources by multiplexing frames. Frames may be delivered out of order, but this service is good when the connection set-up time is greater than the latency of a short message. This class also sends an acknowledgment to confirm frame delivery.

Class 3 is the same as class 2 except that frame delivery is not confirmed. This is a datagram type of service that is good for real-time broadcasts.



[Return to "diagram" link](#)

[Return to "services classes" link](#)

[Return to "fc-4" link](#)

[Return to table of contents](#)

Fibre Channel Topologies

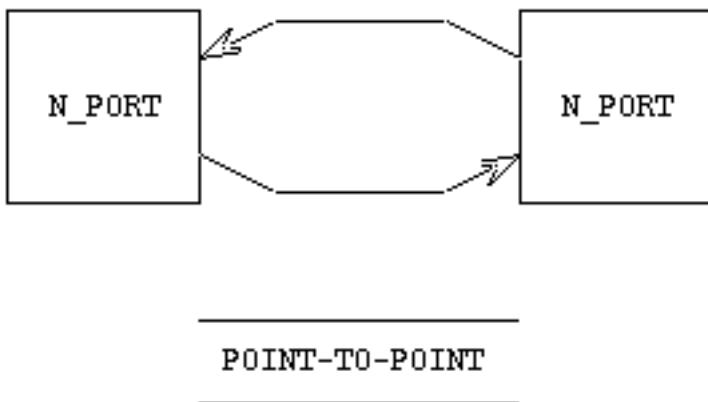
One of Fibre Channels most admired features is its' flexible topology. Fibre Channel includes three connection methods; Point-to-Point, Arbitrated Loop, and Switched Fabric. The ports in a point-to-point connection are called N_Ports; loop connections are called NL_Ports. An FC switch, or a network of switches, is called a fabric. The ports that correspond to it are called F_Ports.

Point-to-Point

The simplest of all connection methods is point-to-point. This type of connection gives the greatest possible bandwidth of all topologies.

The point-to-point topology can connect:

1. computer to computer
2. computer to disk



[Return to "diagram" link](#)


[Return to table of contents](#)

Switched Fabric

The switched fabric topology gives the greatest connection capability and largest total combined throughput. Each device is connected to a switch and receives a non-blocking data path to any other connection on the switch. This would be equivalent to a dedicated connection to every device. As the number of devices increases to occupy multiple switches, the switches are in turn connected together.

The switched fabric connection can:

1. interconnect large numbers of devices
2. sustain high bandwidth requirements
3. connect devices that run at different speeds
4. provide cable matching



[Return to "diagram" link](#)

[Return to table of contents](#)

Arbitrated Loop

An Arbitrated loop connects up to 126 devices in a ring. Arbitrated loop is similar to token ring, where each device arbitrates for loop access, and once granted, has a dedicated connection between sender and receiver. The available bandwidth of the loop is shared between all devices. The primary reason to use arbitrated loop is for cost. Since no switch is required to connect multiple devices, the per connection is significantly less.


Arbitrated loop is used for:

1. storage interconnection
2. interconnection of small numbers of devices
3. intermittent high bandwidth situations

Arbitrated loops are inherently inefficient in large configurations. This is because every node in the loop must look at the data regardless of the destination. Because each node is logically connected in a circle,

significant cable induced delays are possible with large configurations.

Emulex suggests that arbitrated loop be used when the total number of connected nodes is less than 30, the requirements for high bandwidth is intermittent, and the maximum distances between nodes is 100 meters or less.



[Return to "diagram" link](#)

[Return to table of contents](#)

Note that in EACH topology the connected devices must:

- support the same data link speed
- have the same cabling scheme

[Fibre Channel Reference Card: fold-in card in Postscript format from FCA \(670 KB\)](#)

[Return to table of contents](#)

Fibre Channel Products

- [Components \(Chips, cables, etc.\)](#)
- [Fabrics](#)
- [Interfaces](#)
- [Routers \(Ethernet, FDDI, Token Ring, ATM\)](#)

- [SCSI converters](#)
- [Software and Services](#)
- [Storage devices](#)
- [Testers](#)

[Products by manufacturer](#)

[Return to table of contents](#)

Fibre Channel Standards Organizations

- [American National Standards Institute \(ANSI\) X3T11 committee](#)

[Return to table of contents](#)

Helpful Fibre Channel Links

- [American National Standards Institute](#)
- [CERN Fibre Channel homepage](#)
- [Fibre Channel Consortium \(University of New Hampshire's InterOperability Lab\)](#)

[Return to table of contents](#)

Recent Events in Fibre Channel

In March of 1995 The Fibre Channel Systems Initiative (**FCSI**) and the Fibre Channel Association (**FCA**) announced the latest increase possible in data communication speeds with the ANSI Committee adoption of Fibre Channel standards for **2 and 4 gigabit per second (Gbps) data rates**. This is now the fastest data communication speed possible for Fibre Channel. The previous limit was 1 Gbps.

This advance was made possible because of **Vertical Cavity Surface Emitting Laser (VCSEL)** technology, which is new, practical and cost-effective for ultra-high data transfers. Motion picture and video production will be two of the many applications that will benefit from this recent advance.

Fibre Channel News

- [CERN news page](#)

[Return to table of contents](#)

References

- X3T9.3 Task Group of ANSI: Fibre Channel Physical and Signaling Interface (FC-PH), Rev. 4.2 October 8, 1993
- Fibre Channel Overview-CERN High Speed Interconnect pages-15 August 1994-Erik van der Bij (<http://www1.cern.ch/HSI/fcs/spec/overview.htm>)
- Fibre Channel Association: Fibre Channel: Connection to the Future, 1994, ISBN 1-878707- 19-1
- Gary Kessler: Changing channels, LAN Magazine, December 1993, p69-78
- ANSI X3.230-1994, Fibre Channel Physical and Signaling Interface (FC-PH) Rev 4.3
- EGDE, on & about AT&T, March 20, 1995 v10 n347 p25(1)
FCSI standards: fibre channel extends information highway speed limit to 4 gigabits per second;
ANSI Fibre Channel Committee adopts 2 & 4 gigabit data communication rates.*
- Special Report: Fibre Channel Association(FCA)
The LocalNetter, May 1994 v14 n5 p57(1)*

note:The starred (*) references can be found in Computer Select (June 1995)

[Return to table of contents](#)

[Other Reports on Recent Advances in Networking 1995](#)

[Back to Raj Jain's Home Page](#) *Last updated on August 28th, 1995*

Raj Jain is now at Washington University in Saint Louis, jain@cse.wustl.edu <http://www.cse.wustl.edu/~jain/>