FIBRE CHANNEL
BACKGROUND, TODAY, FUTURE

HORST L TRUESTEDT
TrueFocus, Inc.
12 Elmwood Dr NE
Pine Island, MN 55963-9740
USA
(507)356-4701
hotrues@.attglobal.net

06 June 2001
University of Minnesota

AGENDA

- Timeline
- Purpose
- Characteristics
- Current
- Future
- Summary

TIMELINE

- FC began as part of IPI Enhanced Physical
- Scope widened to include HIPPI, SCSI and other interface protocols
- Chartered with wider scope by ANSI in 1988
- First year spent looking at implementations
- Selected 8B/10B in Oct89
- Added a low-cost Copper fibre variant

TIMELINE
(continued)

- Work started on FC-PH in earnest in late 1989
- FC-PH ANSI Standard in 1994
- FC-AL ANSI Standard in 1996
- FCP ANSI Standard in 1996
- FC-LE ANSI Standard in 1996
- FC-PH-2 ANSI Standard in 1997
- FC-PH-3 ANSI Standard in 1998
- FC-AL-2 ANSI Standard in 1999

PURPOSE

- Use Serial Media
  - smaller connectors
  - longer distances
  - easier switching
  - lower cost to install and maintain
  - higher reliability
- Provide a single *pipe* to carry Multiple Protocols
- Be asynchronous--allow maximum throughput, even when busy

PURPOSE
(continued)

- Be scalable
  - allow simultaneous transmission of Multiple Protocols
  - allow transfers over copper (for short distances) and optical (for long distances)
  - allow a single interface to support work stations to supercomputers to disk drives
  - allow the media to ride the technology curve
- Be very, very fast--25MB/s to 400MB/s per fibre per direction; 1000MB/s being developed

PURPOSE
(continued)

- Be switched for highest performance and largest connectivity
- Separate Information Transfer from Physical Transport system
- Minimize Error Checking and Recovery
  - parts need to be designed for 10 E^{-15} BER
  - parts need to be tested for 10 E^{-12} BER
  - errors must be an exception (1 bit per day)
  - errors must be detected
    - 8B/10B code -- 99% of all errors
    - CRC -- no undetected errors in 1.2 million years

PURPOSE
(continued)

- Be an **ENABLER**
  - allow full duplex operations
  - allow highest transfer speeds available
  - allow long distances
  - allow hardware protocol assists
  - allow legacy operations to work better than before, but allow new constructs to improve performance

CHARACTERISTICS

- Network and Channel interface
- 16+ Million direct-attached Nodes
- 50 km per Hop
- 25MB/s; 100BM/s; 200MB/s; 400MB/s; ...
- Low-Latency Protocol
- Flow Control for Distance Insensitivity
- Connection and Connection-less service
- Guaranteed Bandwidth **and** Delivery

CHARACTERISTICS
(continued)

- Virtual Connections and Isochronous Transfers
- Multiplexed Transfers
- Merger of Networking and Channel Cultures
- Multiple Topologies
  - Point-to-Point
  - Switched
  - Loop (Arbitrated, Slotted, and Register Insertion)
- Multiple Upper-Level Protocols

SCSI	**IP**	**SB**	IPI-3	**VI**
ATM	HIPPI	AE	AV	...

PROJECTS AND DOCUMENTS

- FC-AL Arbitrated Loop Topology
- FC-ATM ATM (AAL-5)
- FC-FG Fabric Generic Requirements
- FC-FP Mapping to HIPPI Framing Protocol
- FC-LE Link Encapsulation (IEEE 802.3 & IP)
- FC-PH Fibre Channel Physical and Signaling Interface
- FC-SB Mapping to Single-Byte Command Code Sets
- FC-GS Generic Services
- FC-SW Switch Fabric Topology
- FC-IPI Mapping to Intelligent Peripheral Interface
- SCSI-FCP Mapping to SCSI-3 Phase Emulation Protocol
- FC-VI Mapping to Virtual Interface Architecture

LEVELS

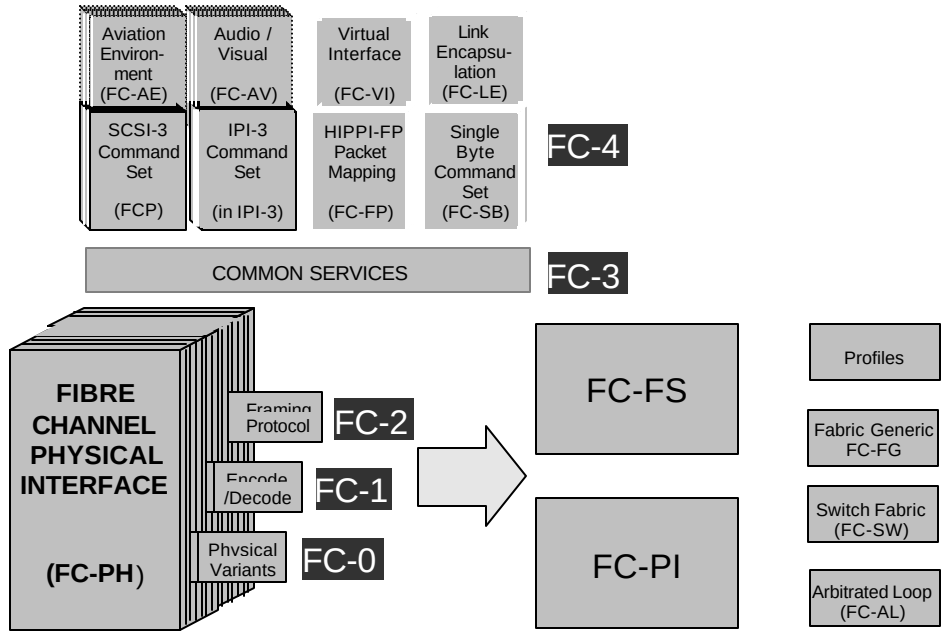

BASIC TERMS

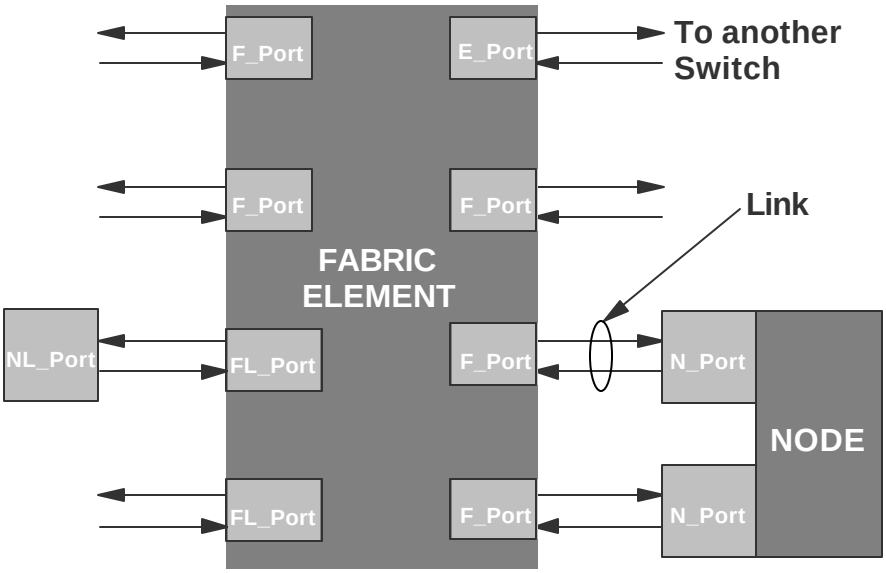

BASIC CONCEPTS

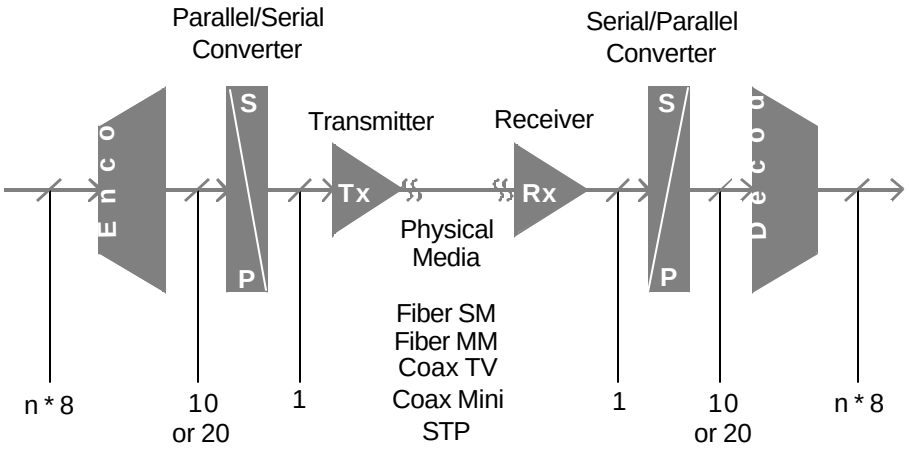

100 MB/s = 800 Megabaud
8B/10B Encoded = 1000 Megabaud
+ Framing (6.2%) = 1.062 Gigabaud

BASIC TRANSFER CONCEPTS

- Transmission Word (40 bits)
  - smallest unit which can be transmitted by FC-0
  - Data Word or Ordered Set
- Frame
  - the smallest unit which can be transmitted by FC-2
- Sequence
  - relates to smallest unit which the ULP can transmit (called an Information Unit at FC-4)
- Exchange
  - relates to an operation (e.g., SCSI command)

BASIC TRANSFER CONCEPTS

(continued)

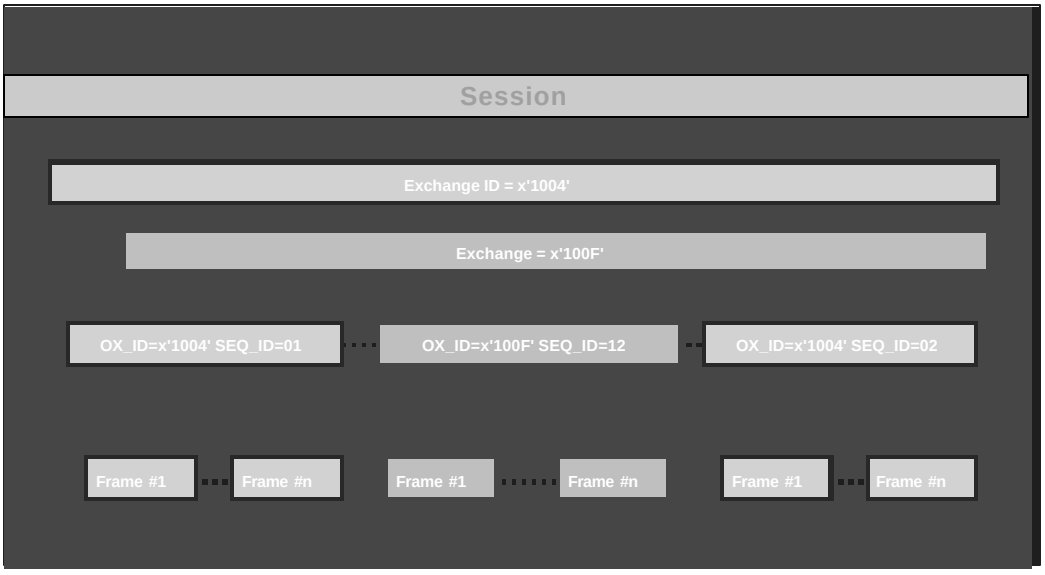

BASIC TRANSFER CONCEPTS
(continued)

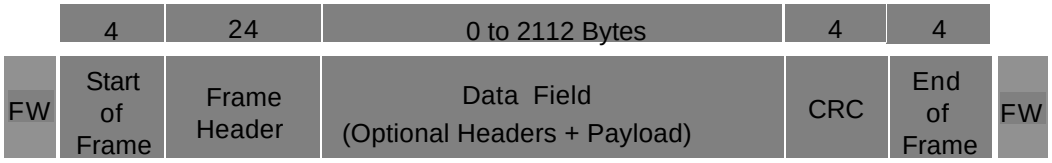

Frame Structure consists of:

- Four byte **SOF** delimiter
- Twenty-four byte fixed-format **Frame Header**
- Variable-size **Data Field:**
  - 0 to a maximum of 2112 bytes (2048+64)
  - May contain Optional Headers as well as payload
- Four byte **CRC** (same as used in FDDI)
- Four byte **EOF** delimiter

CLASSES OF SERVICE

- Class 1
  - Dedicated Connection between two N_Ports
  - Guaranteed delivery
  - Frames received in transmitted order
- Class 2
  - Frame Switched
  - Buffer-to-Buffer flow control
  - Guaranteed delivery
  - Order not guaranteed in the general case
- Class 3
  - Datagrams or "Ship and Pray"
  - Neither Delivery or receipt order guaranteed

CLASSES OF SERVICE
(continued)

- Class 4
  - Fractional Bandwidth
  - Guaranteed delivery
  - Frames received in transmitted order
- Class 6
  - Like Class 1 for Reliable Multicast
  - Guaranteed delivery
- Class 1 Enhancements
  - Class 1 with buffers for speed matching
  - Class 1 simplex
  - Class 1 connect request (camp-on)

TOPOLOGIES

- Point-to-Point
  - Exactly two N_Ports connected together
  - No "Fabric" Services available
- Switched (Fabric)
  - Cross-point switch
  - Packet switched
  - "Fabric" Services available
- Arbitrated Loop
  - Low cost attachment of 3 -126 ports
  - Special Environments
  - "Fabric" Services available via Fabric or F/NL_Port

POINT-TO-POINT

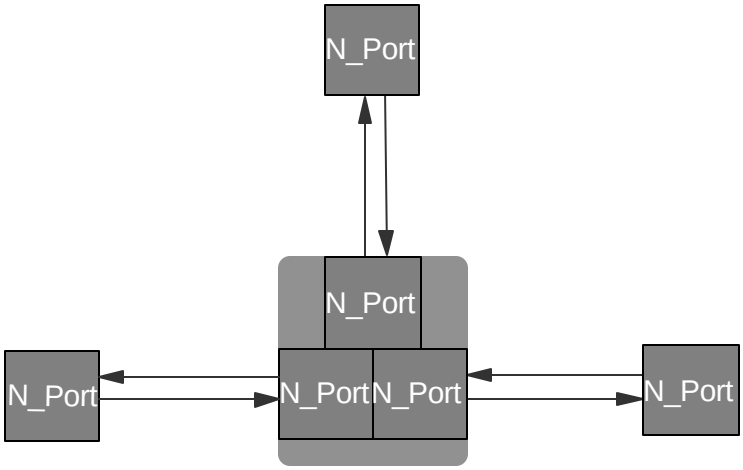

FABRIC

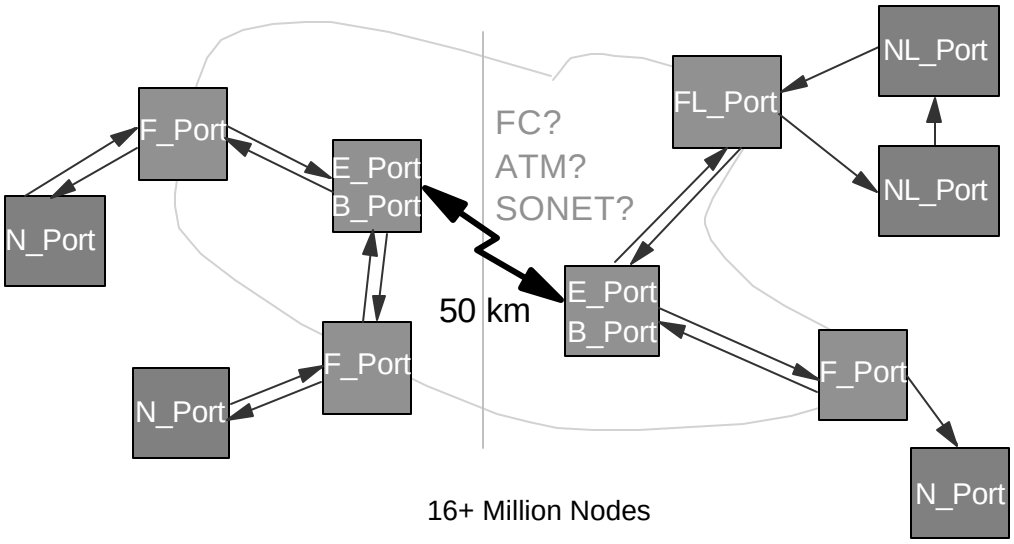

16+ Million Nodes

HUBS AND LOOPS

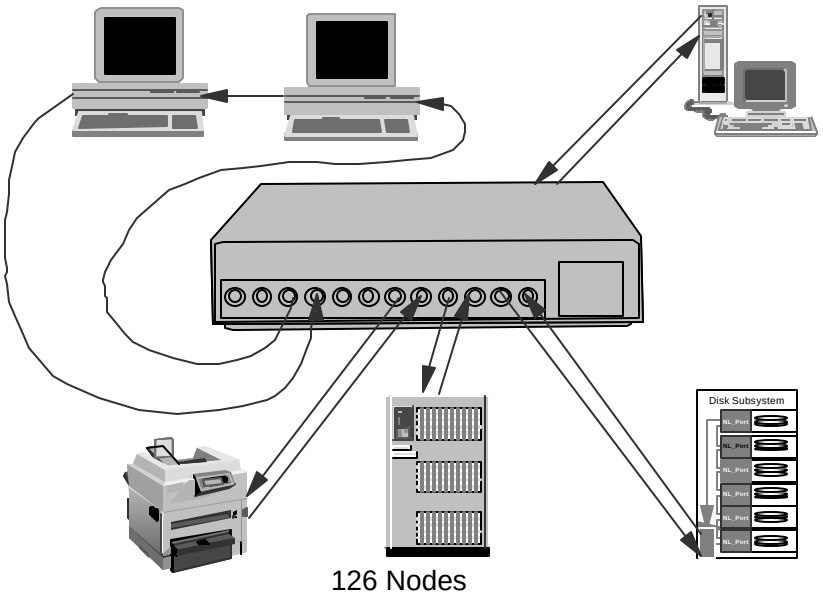

126 Nodes

FABRIC and LOOPS

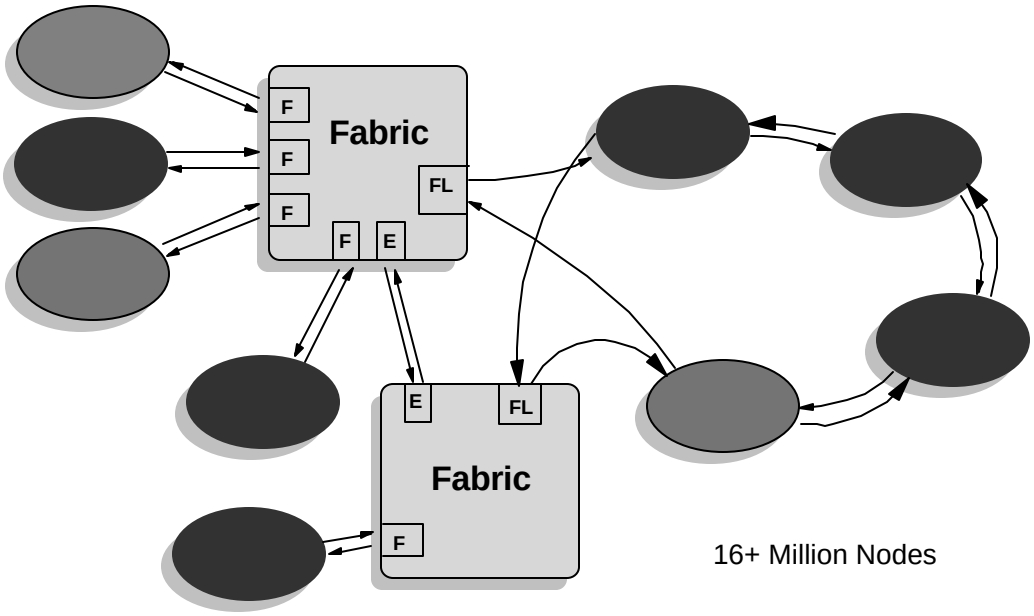

16+ Million Nodes

CONFIGURATIONS

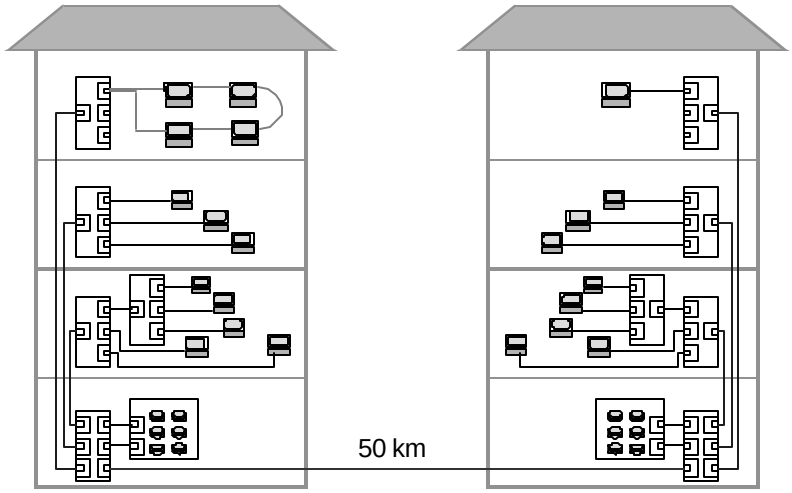

"OFFICE" LOOP

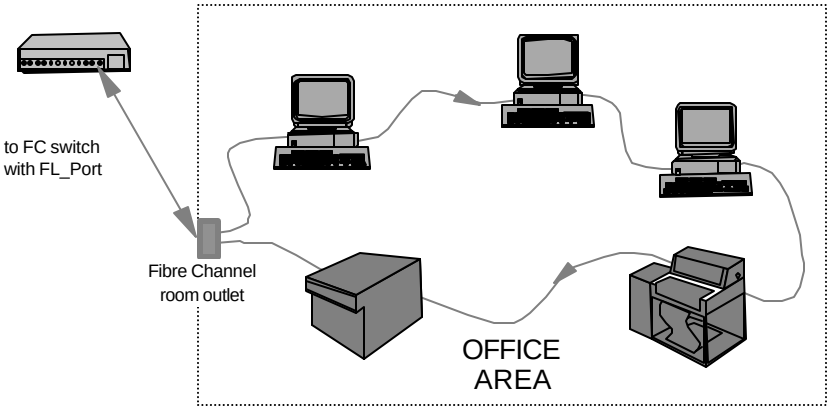

- Additional NL_Port Fibre Channel devices can be added to office area without the need to run additional wiring to FC switch
- Public Loop consists of one FL_Port and multiple NL_Ports

"TRUE THIRD-PARTY"

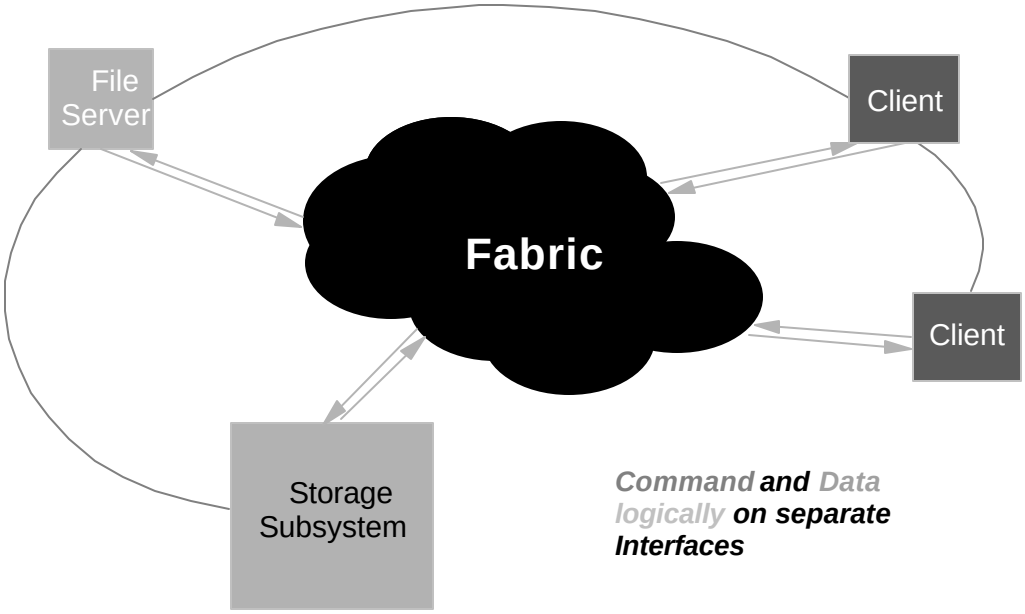

Command and Data logically on separate Interfaces

"DISK" LOOP

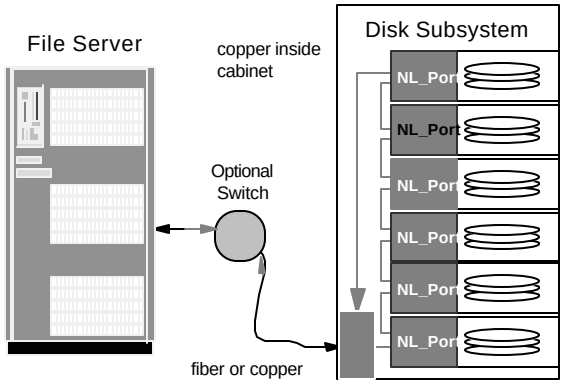

- Allows connection of large number of high performance disks
- Provides high performance / low-footprint interconnect
- Supports multiple Loops and full / half duplex operation on each
- Private Loop comprised of only NL_Ports

LOOP TOPOLOGY

- **Requirements:**
  - "**Dirt Cheap**", allows configuration to operate without separate switch
  - Self discovery of connection to Loop, else revert to point-to-point mode
  - No Loop master (distributed control)
  - Loop can "grow" by attaching to a fabric FL_Port
  - **Key point: Simple protocol** addition to FC-PH

LOOP TOPOLOGY
(continued)

- **Other characteristics:**
  - 126 NL_Ports + 1 FL_Port maximum per Loop
  - Supports all FC-0 variants
  - Supports all FC-PH protocols and classes
  - Supports Access Fairness and Preferred Address Selection
  - Uses arbitration to access the Loop

LOOP TOPOLOGY
(continued)

- L_Port (NL and FL_Port) is a superset of an N and F_Port:
  - One chip(set) with L_Port functionality
  - Simple state machine controls normal operation (11 states)

FC-AL CHARACTERISTICS

- **Connectivity:** *n* Number of L_Ports; 127 Active Addresses
- Distance: 2680 km
- **FC-PH classes:** Class 1, 2, and 3 (4 and 6)
- **Low Cost:** Requires no fabric; reduces number of transmitters / receivers

FC-AL CHARACTERISTICS
(continued)

- **End Node:** Can operate as both N_Port and NL_Port; may connect to an F_Port
- **Fabric Node:** Can operate as both F_Port and FL_Port; may connect to an N_Port
- **Fairness:** L_Port cannot arbitrate again until all other arbitrating L_Ports have won arbitration

FC-AL CHARACTERISTICS
(continued)

- **Addressing:** Assigned automatically, hard-coded, or Fabric
- **Replicate**: Supports broadcast, multicast, selective multicast

LOOP TOPOLOGY ISSUES

- **Resilience of Loop to failures is procurement issue:**
  - In lowest-cost procurements (e.g., single Loop), a power failure of a single Node may cause the entire Loop to fail.

LOOP RESILIENCY

- **Various degrees of resilience are possible:**
  - Redundant loops
  - Multi-ported Nodes
  - Coax relays, optical by-pass switches
  - Active Hubs
  - Active Port Bypass Circuits (PBC's) or Loop Redundancy Circuits (LRC's)

LOOP ADVANTAGES

- Simple addition to protocol defined in FC-PH
- Lowest cost when combined with coax physical variants; requires half the number of transmitter / receivers.
- No separate Fabric is required to connect *n* Nodes

LOOP ADVANTAGES
(continued)

- Backpanel daisy-chain scheme eliminates wiring clutter in configurations with many peripherals (e.g., disk arrays)
- Supports Broadcast and Selective Replicate today
- Easily expands by adding a Fabric (switch)
  1024 port network (16 port switch with 64 loop ports)

FUTURE FC-AL

- **FC-AL-3**
  - **MCM work completed, but viewed to compete with fabric and therefore rejected by the committee**
  - **Minor changes and corrections from FC-AL-2 when appropriate**

INDUSTRY ACTIVITY

- LAN (Local Area Network)
  - Buildings are wired for FC for teleconferencing and multi-media
  - Aviation Environment expects that FC Loops will be on all planes built within 2 years; prototypes began 4Q95
- CHANNEL
  - System interconnect
  - Storage attached (especially SAN)

INDUSTRY ACTIVITY
(continued)

- DEVICE
  - Disk attached profile completed
  - Backplane attached disk drives available
- SAN (Storage Area Network)
  - Storage added and addressed as part of a network
  - Storage may be shared between several computer systems
  - Storage may be located at long distances from computer systems

INDUSTRY ANNOUNCEMENTS

- LSI Logic's Proven Merlin(TM) Fibre Channel Protocol Controller Now Available In Third Generation Architecture
- QLogic Selected by Intel to Provide Interface Technology for Internet Based Enterprise Application
- Emulex Corp announced its LightPulse(TM) LP8000 host adapter and LH5000 digital hub have been selected for use with Intel's new 8-way platform...

INDUSTRY ANNOUNCEMENTS
(continued)

- CNN Continues Partnership with Tektronix for Multimillion Dollar, 44-channel MPEG-2 Systems...
- McDATA To Provide S/390 Fibre CONnection (FICON) Support for IBM 9032 Model 5 Director...
- Adaptec Introduces Fibre Channel RAID Subsystem With Industry-Leading Features...
- GS Networks Offers High-Speed Backup and Restore from the Mainframe to the Desktop...

INDUSTRY ANNOUNCEMENTS
(continued)

- Gadzoox Networks Appoints Storage Industry Leaders to Key Posts: Gadzoox Networks, Inc., the industry leader in Gigabit-speed Fibre Channel storage area networking products...
- Sony and Avid Technology... digital recording and editing systems for broadcast and film, have selected Seagate's newest generation Cheetah drives for their latest digital media solutions. The Seagate Cheetah and Barracuda lines of A/V professional recording products have revolutionized post-production work in the studio.

INDUSTRY ANNOUNCEMENTS
(continued)

- Seagate's latest product offerings include six Fibre Channel configurations, including 18, 36, and 50 Gbyte Barracuda disc drives delivering 7,200-rpm performance and 9, 18, and 36 Gbyte Cheetah disc drives delivering 10,000-rpm performance. All models are currently s**hipping in volume. Newest disks up to 180 Gbytes.**
- **Today at Computex Taipei 2001, IBM will be displaying its new 2Gigabit Fibre Channel technology as well as subsidiary Mylex's full line of RAID controllers.- Jun 05 2:00 PM ET**

INDUSTRY TODAY

- Products have been dribbling into the market place since 1993
- Pre-1996 most emphasis was on 25MB/s
- Post-1996 most emphasis is on 100MB/s
- Disk drives shipping in volumes (now at 200MB/s)
- 11 movies in 1997 were made using FC
- Wall Street is wired for FC

INDUSTRY TODAY
(continued)

- Over 180 companies shipping product
  - switches--FL_Ports cost same as F_Ports
  - disks--FC dual port disks cost same as SCSI
  - adapters supporting IP, SCSI, VI, and SB
  - hubs
- Fall 1998 Comdex showed 60 vendors operating on 2 km long FC-AL
- Jay Kramer of Unisys

  Never before in our industry have so many leading companies come together to form a consensus on the next evolutionary step in interfaces.

INDUSTRY TODAY
(continued)

- Mike Fitzpatrick of Fujitsu on NAB'98

  There were 17 FC exhibitors, including seven different RAID products. They were feeding video streams and post production processing. Seven-hundred leads were generated. Interest was shown by TV networks, TV studios, the movie industry. Many mentioned that they are wiring buildings today for FC and will be on-line within the next two years. I no longer have to sell FC, just tell them where to find it. Fujitsu plans to have FC disks next year. (now available)

INDUSTRY ASSOCIATIONS

- FCSI -- Fibre Channel System Initiative
  - 3 companies: HP, IBM, and SUN
- FCIA -- Fibre Channel Industry Association
  - Replaced FCA and FCLC in July 1999
  - 180+ companies (including Japan & Europe)
- Test Agencies (Standards Police)
  - University of New Hampshire
  - Interphase plug fests in Dallas, TX
  - SNIA lab in Colorado Springs, CO

FUTURE

- Fibre Channel is an enabler. Look for:
  - full duplex transfers
  - focus on using new constructs for optimization and better performance
  - adapters with multiple protocols
  - higher speed (200MB/s, 400MB/s, 1000MB/s)
  - more peripherals (e.g., tape, scanners, printers)
  - more SAN applications
  - more LAN applications

FUTURE
(continued)

- Higher speed; 200MB/s, 400MB/s, 1000MB/s
- More acceptance as a network
- More Upper Level Protocol (ULP) Mappings
- More acceptance of FC hardware:
  - Gigabit Ethernet
  - Gigabit 1394
  - Common 10Gbs hardware for 10GFC and 10GbE

SUMMARY

- A new era in performance:
  - FC's full-duplex operation can cut number of connections in half
  - Multiplexing permits more device and media sharing
- Distance opens up new application opportunities:
  - Direct device connection instead of via File Server
  - Locate processors near people, but protect and isolate data

SUMMARY
(continued)

- Fibre Channel is a comprehensive physical interface for multiple existing and new protocols -- it may be the only interface a system needs!
- Fiber-optic technology is available to meet FC requirements in a cost-effective manner
- Copper proven to operate at 200 MB/s
- Fibre Channel began shipping in 1998 in volumes

BIOGRAPHY

Horst L Truestedt joined the IBM Corporation in Rochester, MN in 1968 as an Associate Engineer after conducting physics research for six years at the 3M Co in St. Paul, MN. He has worked in several divisions of IBM (including assignments to Germany and Japan) in the design and development of Optical Character Recognition equipment; as manager of Diagnostics Engineering on System/36 and System/38; and, on storage product attachments for S/36, S/38, 4300, 9370, and AS/400.

Before retiring from IBM in 1997, Truestedt had been involved with DASD interfaces and attachment architecture. He represented IBM at ANSI T11 - Intelligent Peripheral Interface (IPI), High Performance Parallel Interface (HIPPI), Fibre Channel (FC), and Enterprise System Connectivity (SBCON); represented IBM SSD on the Fibre Channel Association (FCA); represented IBM on the Redundant Array of Independent Disk (RAID) Advisory Board; represented ANSI T11 at the International Standards Organization (ISO); and, consulted for large IBM systems worldwide.

Prior to working at 3M and IBM, Truestedt attended Gustavus Adolphus College, St. Peter, MN in pre-Engineering and received a BS in Physics and Mathematics from the University of Minnesota, Minneapolis, MN.

Truestedt is the President/CEO of TrueFocus, Inc. which is a consulting firm focusing on industry Fibre Channel and Storage Product implementations.