Rivier College

CS553A Introduction to Networking Technologies

Project: A Gigabit Ethernet

Author: Faron Ducharme
Author's E-mail   fducharme@rivier.edu
Due Date: 8th of December 2005
Professor: Vladimir V. Riabov, PhD

Table of Contents

Introduction

Company X is as medium sized enterprise of 300 people (terminals) running and requiring access to the following:

- Enterprise Resource Planning (ERP) Database and Application
- E-mail
- Internet browsing
- File Transfer Protocol (FTP)
- Off-Site Backup
- Storage Area Network (SAN)

Along with proving the above, security is a paramount concern.

The Ethernet operates at layers 1 and 2 of the OSI networking model but as an academic exercise we will also examine the Network, Transport and Application layers.

STORAGE AREA NETWORK

The ever expanding volume of data produced by businesses today, is being driven by new technologies, customer demand and government regulation:

- Publishing companies moving towards digital pre-press publication.
- Retailers selling to online customers.
- Healthcare companies abiding by the Health Insurance Portability and Accountability Act (HIPAA).
- Organizations that mandate long-term archiving of everything from e-mail messages to corporate records.

The Storage Area Network (SAN) was developed in response to this changing climate and the requirements of High Availability (HA), scalable growth and system performance. Some estimates project enterprise storage to double every 12-18 months. According to the IMEX Research "SANs will eventually be at the core of every enterprise's data center." [i]

Most of today's SANs relay on Ethernet to transport data via Internet Small Computer System Interface (iSCSI). ISCSI is a block-level storage protocol that lets users create a separate storage network using Ethernet. ISCSI uses Ethernet as a transport for data from servers to storage devices or storage-area networks. Because iSCSI uses Ethernet, it doesn't suffer from some of the complexity and distance limitations that encumber other storage protocols.

Ethernet: A local-area network (LAN) architecture developed by Xerox Corporation in cooperation with DEC and Intel in 1976. Ethernet uses a bus or star topology and supports data transfer rates of 10 Mbps. The Ethernet specification served as the basis for the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method to handle simultaneous demands. It is one of the most widely implemented LAN standards.

1 TECHNOLOGY BACKROUND

Invented by Robert Metcalfe, during his time at Xerox's Palo Alto Research Center (PARC) the prototype system operated at 3 Mbps and connected 100 computers. Within a number of years Xerox, Intel Corporation and Digital Equipment Corporation had formed an alliance and published a "Blue Book Standard" for the Ethernet, which subsequently formed the basis for the IEEE 802.3 standard.
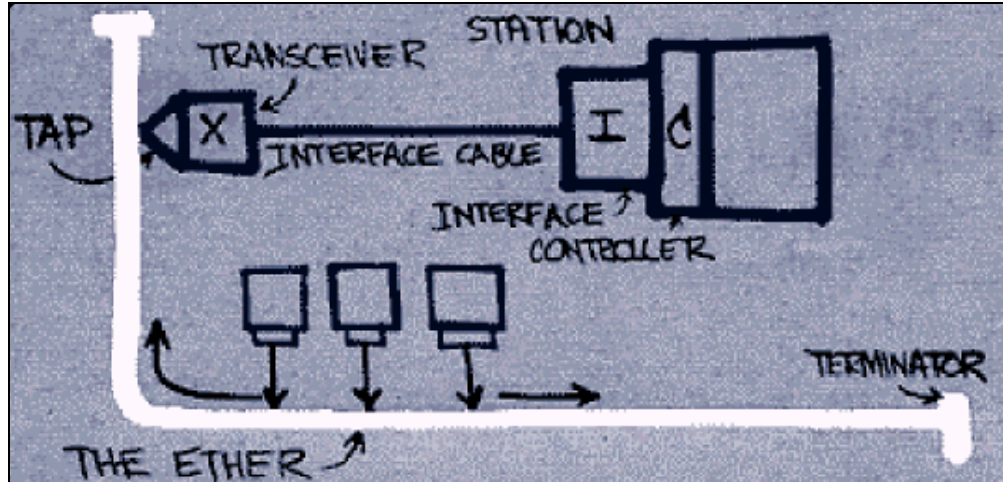

Figure 2.1 Metcalfe's original notes [ii]

Today, the Ethernet is the most commonly employed LAN technology. According to the International Data Corporation (IDC) more than 85% of all network connections are Ethernet.

WHY HAS THE ETHERNET BEEN SO SUCCESSFUL?

Central to the success of the Ethernet, has been its ability to continue to evolve to meet the every growing demand for bandwidth. IT professionals have come to rely on the Ethernet to be simple, easy to use and readily upgradeable. An organization can scale from 10 to 100 (Fast Ethernet) to 1000Mbps (Gigabit Ethernet), assured that any new equipment will be backwards compatible with legacy equipment. This reduces the infrastructure investment that an organization must make.

For a history of the Ethernet visit www.computerhistory.org/.

International Data Corporation - is the world's leading provider of technology intelligence, industry analysis, market and strategic data. IDC provides global research with local content through more than 700 analysts in 43 countries worldwide. IDC is a division of IDG, which was founded in 1964 and has more than 12,000 employees worldwide. IDG media options reach more than 120 million technology buyers in 85 countries; representing 95% of worldwide IT spending.

EVOLUTION OF STANDARDS

The role-played by the Ethernet standards has been to provide a straightforward migration path for companies as their bandwidth requirements have increased.
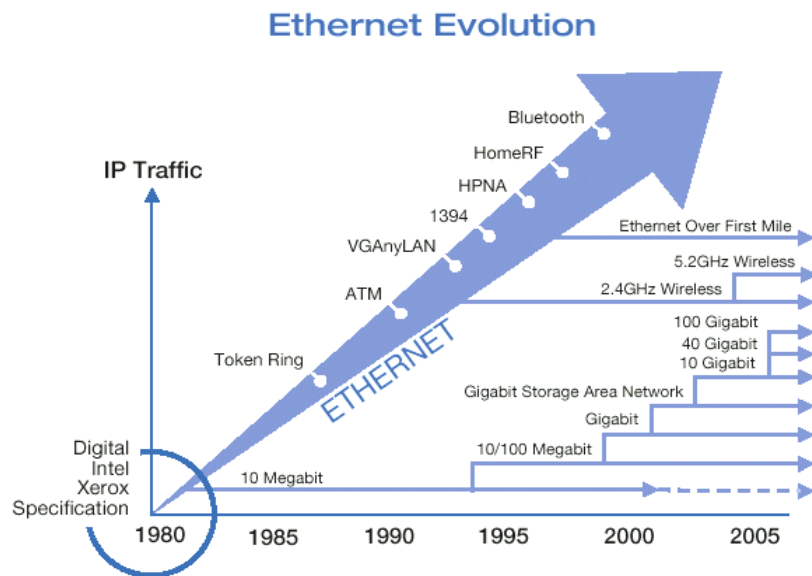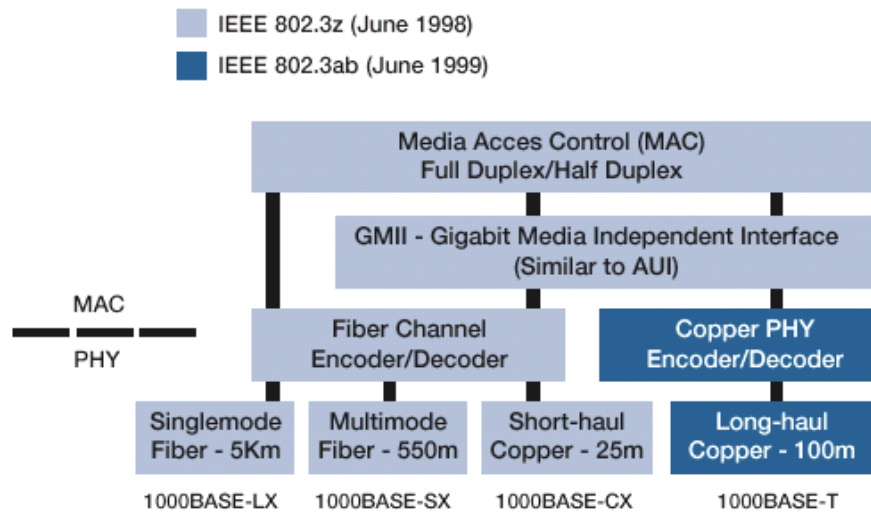

Figure 2.2 [ii]

The Gigabit Ethernet has evolved from the original 10Mbps Ethernet standards and a 10-Gigabit Ethernet standard is being ratified, supported by the IEEE and the 10-Gigabit Ethernet Alliance.

The Gigabit Ethernet Standards

IEEE 802.3z (June 1998)

IEEE 802.3ab (June 1999)

Media Acces Control (MAC)
Full Duplex/Half Duplex

GMII - Gigabit Media Independent Interface
(Similar to AUI)

MAC

PHY

Fiber Channel
Encoder/Decoder

Copper PHY
Encoder/Decoder

Singlemode
Fiber - 5Km

Multimode
Fiber - 550m

Short-haul
Copper - 25m

Long-haul
Copper - 100m

1000BASE-LX   1000BASE-SX   1000BASE-CX   1000BASE-T

Cahners In-Stat Group, 5/01

Figure 2.3 [iii]

In June of 1998, the IEEE adopted a standard for Gigabit Ethernet over fiber optic cabling, IEEE 802.3z, and its implementation was widely supported. A year later, in June of 1999 the IEEE standardized IEEE 802.3ab Gigabit Ethernet over copper as 1000Base-T, allowing Gigabit speed to be transmitted over Cat-5 cable (figure 2.3). As a result companies could rely on an well-understood standards-based approach to improve traffic

1.1 BENEFITS OF GIGABIT

Gigabit Ethernet is 100 times faster than regular 10Mbps Ethernet and 10 times faster than 100Mbps Fast Ethernet. The principal benefits of Gigabit Ethernet include:

• Increased bandwidth for higher performance and elimination of bottlenecks
• Full-duplex capacity, allowing the effective bandwidth to be virtually doubled
• Aggregating bandwidth to multi-Gigabit speeds using Gigabit server adapters and switches
• Quality of Services (QoS) features to help eliminate jittery video or distorted audio
• Low cost of acquisition and ownership
• Full compatibility with the large installed base of Ethernet and Fast Ethernet nodes
• Transferring large amounts of data across a network quickly

2 NETWORK TOPOLOGY

Company X will employ a Bus or Tree topology, all devices on the network connects via one trunk cable. This makes it easy to install and configure and inexpensive. Figure 3.1 shows the proposed topology; geometric arrangement.
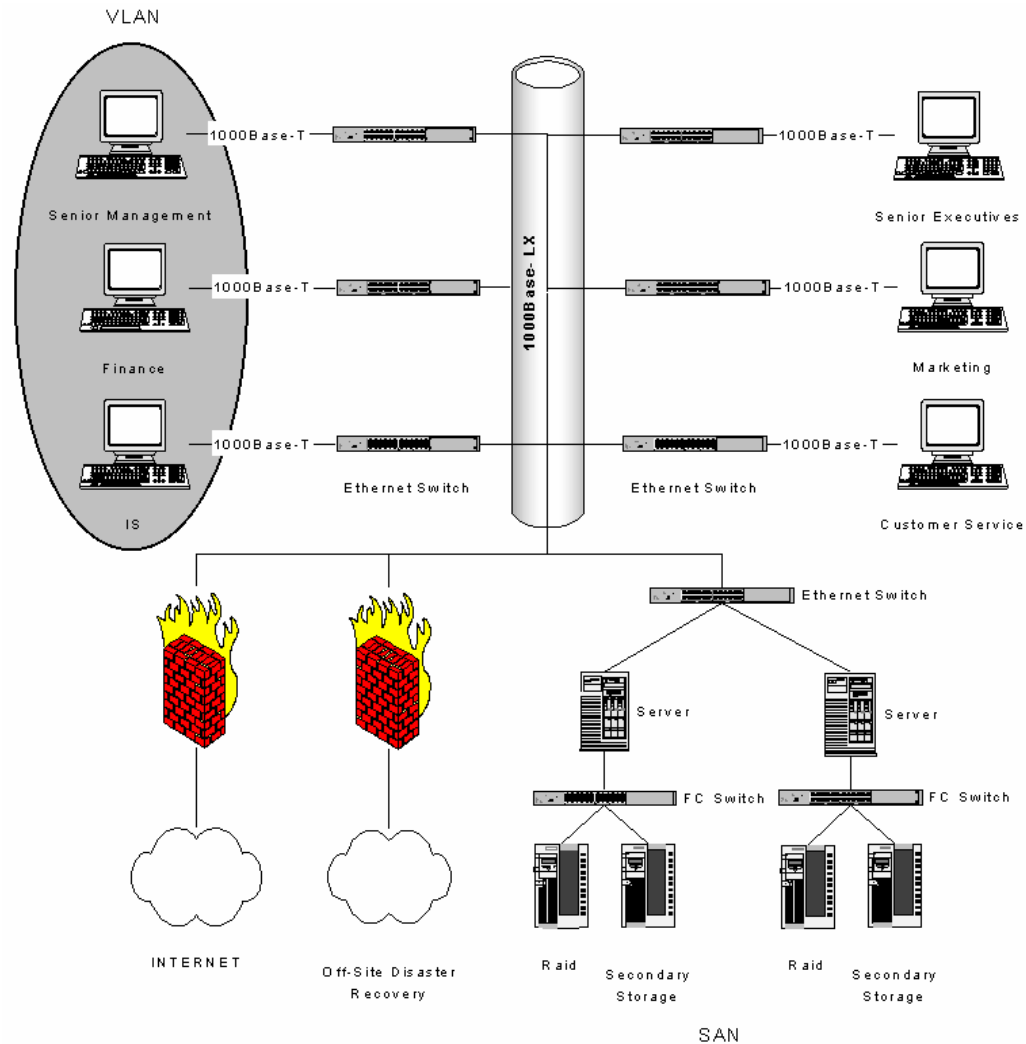
Figure 3.1 Company X's Network Topology

Other issues that should be considered during development of the network topology are the physical limitation of such a network and how to avoid any bottlenecks.

VLAN: Short for virtual LAN, a network of computers that behave as if they are connected to the same wire even though they may actually be physically located on different segments of a LAN. VLANs are configured through software rather than hardware, which makes them extremely flexible. VLANs are often named by color.

NETWORK DIAMETER LIMITATIONS

There are two approaches to calculating the physical limitations of the network topology. The first is a "cookbook" approach, ensuring the 5-4-3 Rule is adhered to.

5-4-3 Rule: divides the network into two types of physical segments: populated (user) segments, and unpopulated (link) segments. User segments have users' systems connected to them. Link segments are used to connect the network's repeaters together. The rule mandates that between any two nodes on the network, there can only be a maximum of five segments, connected through four repeaters, or concentrators, and only three of the five segments may contain user connections. The Ethernet protocol requires that a signal sent out over the LAN reach every part of the network within a specified length of time. The 5-4-3 rule ensures this. Each repeater that a signal goes through adds a small amount of time to the process, so the rule is designed to minimize transmission times of the signals.[iv]

The other approach is more analytical and should be used when the network topology is inconsistent with 5-4-3 model. Firstly, the worst case round trip delay is determined or each segment delay value (SDV) is calculated using the following equation: SDV = Base +[Length(RT delay/meter)]. Next the Interface Gap (IFG) shrinkage is calculated, the IEEE 802.3 standard provides tables for these calculations.

2.1 LOAD BALANCING

An important issue to avoid in network design is server/network bottlenecks, by scaling the backbone and server bandwidth using proven technologies:

Adaptive Load Balancing (ALB): Also known as asymmetric port aggregation—is a method of ensuring consistent high server throughput and transparent backup connections by using multiple network interface cards and balancing the data transmissions across them. As many as four Intel® server adapters, connected to a switch, can be configured to work together as a "team" for an aggregate throughput of up to 400 Mbps with Fast Ethernet adapters or 8 Gbps with Gigabit Ethernet Adapters.
Gigabit EtherChannel (GEC): Provides parallel bandwidth of up to 8 Gbps (4 Gbps full duplex) between a switch and a router, host, or another switch by grouping multiple Gigabit Ethernet interfaces into a single logical transmission path. You can configure Gigabit Ethernet ports into Gigabit EtherChannel groups containing two or four ports, yielding 4 or 8-Gbps bi-directional bandwidth, respectively. Both Fast and Gigabit EtherChannel bundles can be configured as trunk links.

Load balancing can increase server bandwidth up to 8Gbps by automatically balancing traffic across as many as eight network adapters, while technologies such as Intel Link Aggregation, Gigabit EtherChannel and IEEE 802.3ab can increase bandwidth up to 16Gbps.[iii] The application of switching technologies can further increase usable bandwidth but these issues must be considered before deployment. Once the Ethernet is in place any modification can be costly. When sizing a network, particular the backbone it is wise to design in additional capacity, the rule-of-thumb is that on average only 30% of the bandwidth should be occupied.

3 NETWORK ARCHITECTURE (LAYERS AND PROTOCOLS)

Network Architecture is defined as a set of layers and protocols. In this section the elements (hardware and software) that make Company X's network are described.

PHYSICAL LAYER

The Physical Layer transforms data into bits that are sent across the physical media.

Cabling

Gigabit Ethernet supports both copper and fiber cabling and this network employs both. The backbone will consist of 1000Base-LX fiber optic cable. The course text gives a maximum segment length of 5,000 meters, but this is dependent on the fiber type see table 4.1. Either of the 50-micron MMF (Multimode fiber) or 62.5-micron MMF fiber will be sufficient for our network.

Table 4.1

Module	Fiber Type	Maximum Cable Length
1000Base-LX	8-10 micron SMF fiber	10 kilometers (6.2 mi)
1000Base-LX	62.5-micron MMF fiber	550 meters (1805 ft)
1000Base-LX	50-micron MMF fiber	550 meters (1805 ft)

Advantages of 1000Base-LX to Company X's network are:

Performance: 1000Base-LX provides the bandwidth to support the organizational demands.
Security: As fibers do not leak light and are difficult to tap.

Our hosts will be connect to the backbone via 1000Base-T (4 pairs of UTP), this media was chosen as it provides:

Scalability & Flexibility: The backbone will remain fixed but hosts (PC, printers, etc.) must be able to be added and removed with ease to meet the demands of the business.
Simplicity: 1000Base-T is copper and the physical connection from the host(s) to the backbone (fiber) is made though a switch, avoiding complicated and costly splicing/fusing.

KEY GIGABIT OVER COPPER SPECIFICATIONS

Cable configuration – 1000Base-T provides 1Gbps Ethernet signal transmission over four pairs of Category-5 UTP cable. 250Mbps per wire pair multiplied by 4 pairs = 1000Mbps.
Distances – The copper standard covers cabling distances of up to 100 meters, or networks with a diameter of 200 meters (assumes 100 meters in two directions from a switch).

Half duplex and CSMA/CD – Although the standard includes half-duplex operation, few Gigabit products support half duplex to date. Full duplex is preferred to maximize performance. Gigabit Ethernet uses the CSMA/CD protocol only when running in half-duplex mode.

Full duplex and flow control – In full duplex, CSMA/CD collision detection is impractical. Instead, flow control methodology is used to avoid congestion and overloading.

TESTING EXISTING CABLE

Any legacy cabling should be tested for Far-End Crosstalk and Return Loss as per TIA/EIA standard TSB67: Transmission Performance Specifications for Field-Testing of Unshielded Twisted Pair Cabling Systems. If any problems are discovered ANSI/TIA/EIA standard TSB-95 (1998) covers additional testing (recommended) for Gigabit Ethernet (1000BASE-T), which utilize all four pairs for bi-directional transmission and this requires the measuring additional parameters. TSB-95 also defines five simple options for correcting performance:

Option 1 - Change patch cord with Category 5e cord
Option 2 - Change cross connect to an interconnect
Option 3 - Change consolidation point with Category 5e consolidation point
Option 4 - Change work area connector to Category 5e connector
Option 5 - Change interconnect to a Category 5e interconnect

Other factors to be considered are losses at link or splice points, table 4.2. Company X's network will try to avoid as much as possible splicing and connecting fiber.

Link Loss Factor	Estimate of Link Loss Values
Higher-order mode losses	0.5dB
Clock recovery module	1dB
Modal and chromatic dispersion	Dependent on fiber and wavelength
Connector	0.5dB
Splice	0.5dB
Fiber attenuation	1dB/km

table 4.2 Estimating Link Losses [vi]

3.1 DATA LINK LAYER

The Data Link Layer determines access to the network media in terms of frames. Its Media Access Control (MAC) sub-layer is responsible for physical addressing.

3.1.1 Medium Access Control

The Medium Access Control (MAC) for Company X's Gigabit Ethernet is an enhanced version of 802.3, Gigabit Medium-Independent Interface (GMII) has been defined for all medium options except for unshielded twisted-pair (UTP). The GMII defines 8-bit-parallel transmit and receive synchronous data interface. These two signal-encoding schemas, is defined by 8B/10B for optical fibers and shielded copper media and the pulse amplitude modulation 5-PAM is used for UTP.

5-level PAM provides better bandwidth utilization by having each transmitted symbol represents one of five different levels (-2, -1, 0, +1, +2). Since each symbol can represent two bits of information, the symbol rate and therefore also the single bandwidth is reduced by a factor of two. The costs of multilevel signaling includes the need for higher signal-to-noise ratio for a given error rate, the use of multi-bit D/A and A/D converts and the need for better receiver equalization.

The 1000 Mbps specification calls for the same CSMA/CD frame format with two important enhancements:

Carrier Extensions - Carrier extensions appends a set of special symbols to the end of the short MAC frames so that the resulting block is at least 4096 bit-times in duration, up from 512 bit-times imposed at 10 and 100 Mbps. This extension makes the frame length of a transmission longer than the propagation time at 1Gbps.

Frame Bursting - This feature allows for multiple short frames to be transmitted consecutively, up to a limit, without relinquishing control for CSMA/CD between frames. Frame bursting avoids the overhead of carrier extension when a single station has a number of small frames ready to send extension when a single station has numerous small frames ready to send.

The above Data Link protocols can be implemented by the use of SmartBit® GX-1421A, produced by SPIRNET Communication. The GX-1421A is a Media Access Control (MAC) interface for either copper or fiber. GX-1421A utilizes PHY chips that support all IEEE 802.3 standards. For more product information http://www.spirentcom.com/.

3.1.2 Ethernet Frame Format

Currently there are four different Ethernet frame formats, table 4.3. The reason for this variety is historic but regardless of the frame format being used each transmission must be preceded by a Preamble. Because each station has its own oscillating clock the communication stations must have a way of "synch up" their clocks; agreeing on how long one bit is in time. The preamble, which consists of 8 bytes of alternating ones and zeros ending in 11, achieves this.

Table 4.3 The different "flavors" of Ethernet

Frame Type	Novell calls it...	Cisco calls it...
IEEE 802.2	ETHERNET_802.2	LLC
Version II	ETHERNET_II	ARPA
IEEE 802.3 SNAP	ETHERNET_SNAP	SNAP
Novell Proprietary (802.3 Raw)	ETHERNET_802.3	NOVELL

Company X has selected IEEE 802.2 as its frame format, which is detailed below.

6 bytes	6 bytes	2 bytes	1 bytes	1 bytes	1 bytes	43 to 1497 bytes	4 bytes
Des. MAC	Source MAC	Length	DSAP	SSAP	Control	Data	FCS
Data Link Header			Logical Link Control			Data and CRC	

Figure 4.1 Ethernet frame format IEEE 802.2

As can be seen in figure 4.1 802.2 can be subdivided into the Data Link, Logical Link Control and Data and Frame Check Sequence (FCS) headers described below.

THE DATA LINK HEADER

The Data Link Header is composed of the following:

Offset 0-5 - The Destination Address: The first six bytes of an Ethernet frame make up the Destination Address, which specifies where the data frame is being sent. A Destination Address of all ones specifies a Broadcast Message that is read in by all hosts attached to the Ethernet. The IEEE assigns the first three bytes of the Destination Address to vendors. The Destination Address format is identical in all Ethernet implementations.

Offset 6-11 - The Source Address: The next six bytes of an Ethernet frame make up the Source Address, which specifies the originator of the message. Like the Destination Address, the first three bytes specify the vendor of the card. The Source Address format is identical in all implementations of Ethernet.

Offset 12-13 – Length: Bytes 13 and 14 of an Ethernet frame contain the length of the data in the frame; not including the preamble. An Ethernet frame cannot be shorter than 64 bytes total length and no longer than 1518 bytes total length.

LOGICAL LINK CONTROL (LLC) HEADER

Following the Data Link Header is the Logical Link Control Header, which is detailed in the IEEE 802.2 Specification. The purpose of the LLC header is to provide a "hole in the ceiling" of the Data Link Layer. By specifying into which memory buffer the host places the data frame, the LLC header allows the upper layers to know where to find the data.

Offset 15 - The DSAP: The DSAP, or Destination Service Access Point, is a 1 byte field that simply acts as a pointer to a memory buffer in the receiving station. It tells the receiving NIC in which buffer to put this information. This functionality is crucial in situations where users are running multiple protocol stacks.

Offset 16 - The SSAP: The SSAP (Source Service Access Point) is analogous to the DSAP and specifies the Source of the sending process.

Offset 17 - The Control Byte: Following the SAPs is a one-byte control field that specifies the type of LLC frame that this is.

USER DATA AND THE FRAME CHECK SEQUENCE (FCS)

The Data and CRC contain:

Offset 43-1497 - Data: Following the header are 43 to 1,497 bytes of data, generally consisting of upper layer headers such as TCP/IP or IPX and the actual data.

Last 4 Bytes - FCS: The last 4 bytes are the Frame Check Sequence or CRC. When the voltage on the wire returns to zero, the host checks the last 4 bytes it received against a checksum that it generates via a complex polynomial. If the calculated checksum does not match the checksum on the frame, the frame is discarded.

Switches

Switches were chosen for a number of reasons:

- To simulae point-to-point service, reducing network traffic on the backbone/bus by having inter-departmental communication remaining with the department circuit/virtual LAN.
- Point-to-point also improves security; e-mail sent from a member of senior management goes directly to the intended recipient and is not broadcast over the entire network.

    A large number of companies, manufacture switches Cisco, Asanté, Nbase-Xyplex, Allied Telesyn, SMC Networks to name but a few. These manufactures produce a wide range of feature-rich switching equipment providing:

- Flexible media interface.
- Flow Control.
- Comprehensive management software.
- Redundant hot swappable external power supply and other fault tolerant features.

3.2 NETWORK LAYER

As mentioned in the introduction, an Ethernet operates at the Physical and Date Link Layers of the OSI model. Company X's network will provide access to the Internet and the ability to send and receive e-mails, this will require Company X's network to communication with Internet Protocols. What follows is a brief examination of the latest Internet protocol, Internet Protocol Version 6 (IPv6).

IPv6 is the "next generation" protocol designed by the Internet Engineering Task Force (IETF) to replace the current version Internet Protocol, IP Version 4 (IPv4). Most of today's Internet uses IPv4, which is now nearly twenty years old. IPv4 has been remarkably resilient in spite of its age, but it is beginning to have problems. Most importantly, there is a growing shortage of IPv4 addresses, which are needed by all new machines added to the Internet. IPv6 fixes a number of problems by:

Expanded Routing and Addressing Capabilities: IPv6 increases the IP address size from 32 bits to 128 bits, to support more levels of addressing hierarchy and a much greater number of addressable nodes. The scalability of multicast routing is improved by adding a "scope" field to multicast addresses.
Improved Support for Options: Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options and greater flexibility for introducing new options in the future.

6bone: is an IPv6 testbed to assist in the evolution and deployment of IPv6. For more information on 6bone consult http://www.6bone.net/.

The Internet Engineering Task Force (IETF): is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The actual technical work of the IETF is done in its working groups, which are organized by topic into several areas (e.g., routing, transport, security, etc.). Much of the work is handled via mailing lists. The IETF holds meetings three times per year. The IETF working groups are grouped into areas, and managed by Area Directors, or ADs. The ADs are members of the Internet Engineering Steering Group (IESG). Providing architectural oversight is the Internet Architecture Board, (IAB). The IAB also adjudicates appeals when someone complains that the IESG has failed. The IAB and IESG are chartered by the Internet Society (ISOC) for these purposes. The General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB. For more details on the IETF visit http://www.ietf.org/.

Quality-of-Service Capabilities: A new capability is added to enable the labeling of packets belonging to particular traffic "flows" for which the sender requests special handling, such as non-default quality of service or "real- time" service.
Authentication and Privacy Capabilities: IPv6 includes the definition of extensions, which provide support for authentication, data integrity, and confidentiality.

IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period. More information about the IPv6 protocol can be found at http://www.ipv6.org/ and http://www.ipv6forum.com/.

TRANSPORT LAYER

The Transport Layer provides end-to-end connectivity, currently it is estimated that between 80% to 90% of Internet traffic utilizes TCP/IP.

TCP/IP

The TCP/IP is a combination of two protocols, Transmission Control Protocol (TCP) and Internet Protocol (IP) which were developed by a Department of Defense (DoD) research project to connect a number of heterogeneous networks.

INTERNET PROTOCOL

The IP is responsible for moving packet of data from node to node. IP forwards each packet based on a four byte destination address (the IP number). The Internet authorities assign ranges of numbers to different organizations. The organizations assign groups of their numbers to departments. IP operates on gateway machines that move data from department to organization to region and then around the world.

TRANSMISSION CONTROL PROTOCOL

The TCP is a connection-oriented transport protocol, which establishes a virtual connection between a destination and a source. TCP is responsible for verifying the correct delivery of data from client to server. Data can be lost in the intermediate network. TCP adds support to detect errors or lost data and to trigger retransmission until the data is correctly and completely received.
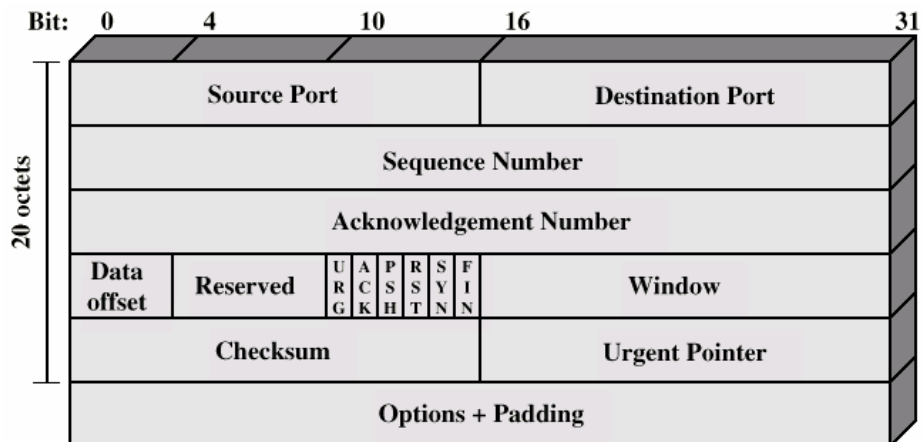
Figure 4.2 TCP Segment Header [vii]

Source Port - 16 bits: The source port number.
Destination Port - 16 bits: The destination pot number.
Sequence Number: 32 bits: The sequence number of the first data octet in this segment (except when SYN is present). If SYN is present the sequence number is the initial sequence number (ISN) and the first data octet is ISN+1.
Acknowledgment Number - 32 bits: If the ACK control bit is set this field contains the value of the next sequence number the sender of the segment is expecting to receive. Once a connection is established this is always sent.
Data Offset - 4 bits: The number of 32 bit words in the TCP Header. This indicates where the data begins. The TCP header (even one including options) is an integral number of 32 bits long.
Reserved - 6 bits: Reserved for future use.
Control/Flag Bits - 6 bits (from left to right):
      URG: Urgent Pointer field significant
      ACK: Acknowledgment field significant
      PSH: Push Function
      RST: Reset the connection
      SYN: Synchronize sequence numbers
      FIN: No more data from sender
Window - 16 bits: The number of data octets beginning with the one indicated in the acknowledgment field which the sender of this segment is willing to accept.
Checksum: 16 bits: The checksum field. The checksum also covers a 96 bit pseudo header conceptually prefixed to the TCP header. This pseudo header contains the Source Address, the Destination Address, the Protocol, and TCP length. This gives the TCP protection against misrouted segments. This information is carried in the Internet Protocol and is transferred across the TCP/Network interface.

Source Address		
Destination address		
Zero	PTCL	TCP Length

Figure 4.3 Pseudo Header

The final layer of Company X's network is the Applicant Layer it provides interfaces between the user's applications and the network through which communication passes.

Telnet

Telnet is a terminal emulation application based on the TELNET protocol. This application is used to connect to remote computers, usually via the telnet port (23). The TELNET protocol is documented in RFC 854. TELNET operates using TCP and depends heavily on option negotiation. Many TELNET options exist; a complete, current list can be found in the Internet Official Standards RFC, currently RFC 2400. A few are listed in table 4.4.

Table 4.4

Protocol	Name	RFC
TOPT-BIN	Binary Transmission	856
TOPT-ECHO	Echo	857
TOPT-SUPP	Suppress Go Ahead	858
TOPT-STAT	Status	959
TOPT-TIM	Timing Mark	860

TELNET SESSION

To establish a Telnet connection the Telnet application must be installed on the clients' PC and contain the name of the server we wish to connect to. The steps in establishing a connection are as follows:

1. Executes the Telnet program, this is normally done by clicking on an icon, that calls the appropriate executable.
1. The Telnet application, using the TELNET protocol which in turn uses TCP to select a port.
1. The Host name is converted to an IP address to set up a connection.
1. The Data Link Layer adds on a header and trailer and sends on the packet to the router.
1. At the Router's Data Link Layer the header and trailer added at 4 (at the client) are stripped.
1. The IP header are read and a decision is made where to send the packet, in this example the server but it could be another router depending on the number of hops required to arrive at the server.
1. The Router adds on a herder and trailer.
1. At the Data Link Layer of the Server the header and trailer added at 7 are stripped.
1. The Network Layer of the Server reads the IP header and it is at this time that a connection is established. Only the client and the server are aware of the connection.

4 CONCLUSIONS

A brief word on cost. The component costs are as follows:

- Server - Sun Fire 4800 Server start at $111,395.00
- Switch - Cisco Catalyst 3550-12G: $6,888.03
- Fiber - 1000Base-LX50-micron MMF fiber $238 per 14ft
- Cable - 1000Base-T (4 pairs of UTP) $134 per 1000ft

The network described in figure 3.1 would cost approximately 1M, not including host devices; PC, printers, scanners, etc. Another element not includes is software, such as protocols and other system management elements (I don't know if they would be included with hardware?).

Giving that Company X is employing an ERP system and other bandwidth-hungry applications there is a real benefit in going with a Gigabit Ethernet. Recent third-party testing conducted by Competitive System Analysis (CSA) measured up to a 47% performance gain with Gigabit Ethernet over traditional Ethernet deployments of commonly user client/server database applications.

THE BOTTOM LINE

Customers who are planning new PC purchases should seek-out systems with Gigabit Ethernet capability. The potential gains, in the form of round-trip transaction throughput to the most common BackOffice services are simply too great to ignore. Combined with 1st tier vendor support and seamless backward compatibility, the performance advantages of today's Gigabit Ethernet solutions fully justify the modest price premium that such forward-looking systems my exact.[viii]

Note to Professor: This is the first time I have tackled this subject matter and I would appreciate any feedback you could provide on this paper.

references

i   http://www.imexresearch.com/
ii   The New Ethernet, Intel Corporation 2001
iii Gigabit Ethernet Technology and Solution, Intel Corporation 2001
iv http://webopedia.com/
vi   Catalyst 5000 Family Supervisor Engine Installation Guide, Cisco
vii Data and Computer Communications, seventh edition, William Stallings, Prentice Hall, 2004
viii Gigabit to the Desktop, Competitive System Analysis, Inc. 2001

WEB SITES REFERENCED

http://www.ansi.org/
http://www.cisco.com/
http://www.cotse.com/
http://www.firewall.cx/home.php
http://www.instat.com/
http://www.intel.com/
http://www.ipv6.org/
http://www.ipv6forum.com/
http://www.iscsistorage.com/
http://www.nextag.com/
http://www.novell.com/
http://www.sun.com/servers/
http://www.telnet.org/
http://www.tiaonline.org/

PAPERS AND JOURNAL REFERENCED

10 Gigabit Ethernet Technology Overview, Intel, 2003
10 Gigabit Ethernet, Cisco, 2000
1000Base-T – Delivering Gigabit Intelligence on Copper Infrastructure, Cisco, 2000
Deploying Gigabit Ethernet to the Desktop: Drivers and Applications, Cisco, 1992 – 2003
EtherChannel Configuration Guidelines and Restrictions, Chapter 11, Cisco, 2002
Ethernet Connectivity Selection Tutorial, ADC Telecommunications, 2003
Gigabit Ethernet 1000Base-T Whitepaper, Gigabit Ethernet Alliance, 1997
Gigabit Ethernet Over Fiber and Copper, Gigabit Ethernet Networking Solutions, Cisco and Intel, 2004
Gigabit Ethernet Technology and Solutions, Intel, 2001
Gigabit Ethernet, SysKonnect, 1999
Gigabit to the desktop, Express Computer - IT Business Weekly, G B Kumar
Introduction to Gigabit Ethernet, Cisco, 2000
Introduction to Industrial Ethernet, Part 4, Volume 1 – Issue 6, George Thomas
ISCSI delivers storage over Ethernet, B Mackin, Network World, September 2002
ISCSI Delivers Storage Over Ethernet, Business Briefing: Data Management & Storage Technology, 2002
The Internet Protocol Journal, September 1999
The New Ethernet, Intel, 2001
Understanding and Configuring EtherChannel, Software Configuration Guide for the Catalyst 4006 Switch, Cisco, Release 12.1(11b) EW