Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Dr. PradeepK. Bhatia**
Lesson Number: 01	

INTRODUCTION

Objective: The objective of the present lesson is to be familiar with concepts, history and terms related to Internet. Internet related organizations and network related command will also be discussed.

Structure

1.1 Growth of Computer Networking

Most historical reviews of the networking imply that networking began with ARPANET(Advance Research Project Agency Network). In a sense, digital transmission of data began when Samuel B. Morse, one of the famous scientiest publicly demonstrated the telegraph in 1844. In 1874, Thomas Edison invented the idea of multiplexing two signals in each direction over a single wire. With higher

speeds and multiplexing, Edison's teletype replaced Morse's manual system; and a few teletype installations still exist today.

The early telegraph systems were, in modern terms, point-to-point links. As the industry grew, switching centers acted as relay stations and paper tape was the medium that the human routers used to relay information from one link to another. Figure 1 illustrates a simple single-layer telegraphic network configuration. Figure 2 shows a more complex multilayered network.

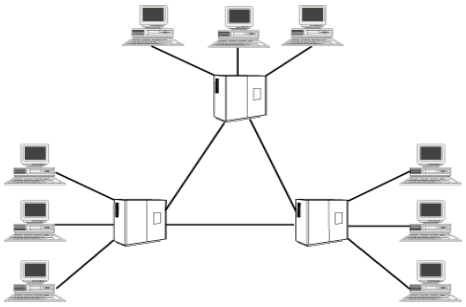

Figure 1 (Single layered telegraphic network)

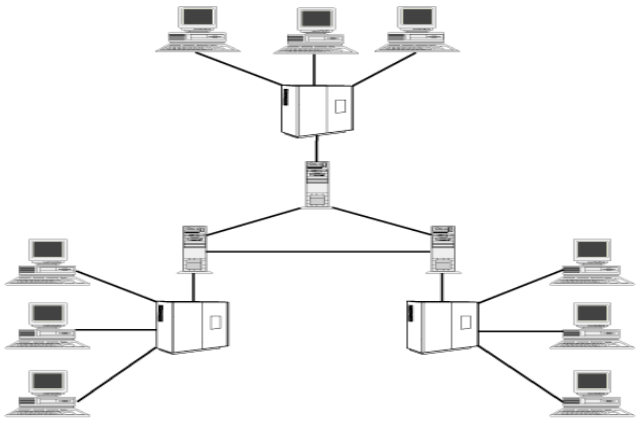

Figure 1.2 (Multilayered Network)

The links of these networks were point-to-point asynchronous serial connections. For a paper tape network, the incoming information was punched on paper tape by high-speed paper tape punches and was then manually loaded on an outgoing paper tape reader.

In 1962, Paul Baran and his colleagues at the Rand Corporation of USA were tackling the problem of how to build a computer network that would survive a nuclear war.

The year 1969 was a year of milestones. Not only did NASA place the first astronauts on the moon but also, and with much less fanfare, Department of Defense's Advanced Research Projects Agency (ARPA) contracted with Bolt, Baranek, and Newman (BBN) to develop a packet-switched network based on Paul Baran's ideas. The initial project linked computers at the University of California at Los Angeles (UCLA), Stanford Research Institute (SRI) in Menlo Park, California, and University of Utah in Salt Lake City, Nevada. On the other side of the continent from the ARPANET action, Brian W. Kernighan and Dennis M. Ritchie brought UNIX to life at Bell Labs (now Lucent Technologies) in Murray Hills, New Jersey.

Even though message switching was well known, the original ARPANET provided only three services: remote login (telnet), file transfer, and remote printing. In 1972, when ARPANET consisted of 37 sites, e-mail joined the ranks of ARPANET services. In October 1972 ARPANET was demonstrated to the public at the International Conference on Computer Communications in Washington, D.C. In the following year, TCP/IP was proposed as a standard for ARPANET.

The amount of military-related traffic continued to increase on ARPANET. In 1975 the Defense Communications Agency (DCA) changed its name to DARPA (Defense Advanced Research Projects Agency) and took control of ARPANET. Many non-government organizations wanted to connect to ARPANET, but DARPA limited private sector connections to defense-related organizations. This policy led to the formation of other networks such as BBN's commercial network Telenet.

The year 1975 marked the beginning of the personal computer industry's rapid growth. In February 1975, Altair announced its microcomputer. In those days when you bought a microcomputer, you received bags of parts that you then assembled. Assembling a computer was a lot of work, for a simple 8KB memory card required over 1,000 solder connections. Only serious electronic hobbyists were able to built computers. From their experiences with the Altair, Paul Allen and Bill Gates founded Microsoft to develop BASIC for the new PC world.

In 1976, four years after the initial public announcement that ARPANET would use packet-switching technology, telephone companies from around the world through the auspices of CCITT (Consultative Committee for International Telegraphy and Telephony) announced the X.25 standard. Although both ARPANET and X.25 used packet switching, there was a crucial difference in the implementations. As the precursor of TCP/IP, the ARPANET protocol was based on the end-to-end principle; that is, only the ends are trusted and the carrier is considered unreliable (the section on TCP/IP later in this book covers this technology in more detail).

On the other hand, the telephone companies preferred a more controllable protocol. They wanted to build packet-switched networks that used a trusted carrier, and they (the phone companies) wanted to control the input of network traffic. Therefore, CCITT based the X.25 protocol on the hop-to-hop principle in which each hop verified that it received the packet correctly. CCITT also reduced the packet size by creating virtual circuits.

In contrast to ARPANET, in which every packet contained enough information to take its own path, with the X.25 protocol the first packet contains the path information and establishes a virtual circuit. After the initial packet, every other packet follows the same virtual circuit. Although this optimizes the flow of traffic over slow links, it means that the connection depends on the continued existence of the virtual circuit.

CCITT regulated input into the network by enabling transmission only when the sender received a credit packet, thereby controlling the overall traffic throughout the

network. Although X.25 is now a dying protocol, it played a very important role in the development of enterprise networks.

Therefore, the end-to-end principle of TCP/IP and the hop-to-hop principle of X.25 represent opposing views of the data transfer process between the source and destination. TCP/IP assumes that the carrier is unreliable and that every packet takes a different route to the destination, and does not worry about the amount of traffic flowing through the various paths to the destination. On the other hand, X.25 corrects errors at every hop to the destination, creates a single virtual path for all packets, and regulates the amount of traffic a device sends to the X.25 network.

The year 1979 was another milestone year for the future of the Network. Computer scientists from all over the world met to establish a research computer network called Usenet. Usenet was a dial-up network using UUCP (UNIX-to-UNIX copy). It offered Usenet News and mail servers. The mail service required a user to enter the entire path to the destination machine using the UUCP bang addressing wherein the names of the different machines were separated by exclamation marks (bangs).

In 1983, LAN (Local Area Network) model of Computer was introduced, in which number of computers with in organization were able to communicate with each other.

Table 1 (Major Landmarks in History of Computer Network)

Year	**Achievements**
1844	Demonstration of Telegraph by Samuel B. Morse
1874	Thomas Edison invented the idea of multiplexing two signals in each direction over a single wire.
1962	Paul Baran and his colleagues at the Rand Corporation were tackling the problem of how to build a computer network that

	would survive a nuclear war.
1969	Department of Defense's Advanced Research Projects Agency (ARPA) contracted with Bolt, Baranek, and Newman (BBN) to develop a packet-switched network based on Paul Baran's ideas
1972	ARPANET consisted of 37 sites, e-mail joined the ranks of ARPANET services
1973	TCP/IP was proposed as a standard for ARPANET.
1975	Defense Communications Agency (DCA) changed name from ARPA to DARPA (Defense Advanced Research Projects Agency) and took control of ARPANET. Microcomputer introduced
1979	USENET was introduced
1983	LAN was introduced

1.2 Complexity in Network Systems

Complexity in Computer Network exists due to following reasons-

- There are many different machines, which are different in both hardware and software but still to be connected on computer network.
- There are many protocols and standards, which suite to different group of people.
- There are increasing requirement of running applications on networks, which are developed for stand- alone PC.
- There are many communication media such as cables, fiber optics, wireless involved in computer networks.
- We wish our communication on computer network to be able to detect and correct errors.

- There is tremendous pressure to increase the speed on Computer network day by day.
- The goals of computer networks are providing not only data and voice transport, but a single integrated network, providing various services.
- Complexity in computer network increases due to requirement of security.
- There are many political, technical and economical complexities attached with networks.

1.3 Growth of Internet

Growth of internet can be well understand from growth of network. Almost as soon as computers were developed, the need to transfer information between machines became apparent. Initially, this was done by writing the information to an intermediate medium (such as magnetic tape or punched cards) and physically carrying the medium to the new machine.

In the early 1960s, computer scientists across the country began exploring ways of directly connecting remote computers and their users. In the mid-to-late' 60s, the United States government began to realize the impact computers would have on education and military research and development. So, the US government decided to fund and experimental network that would allow remote research and development sites to exchange information. This network, funded by the U.S. Advanced Research Projects Agency, was christened the ARPANET.

In the late1980s the National Science Foundation (U.S) funded the development of a network named the NSFNET to connect supercomputer centers all over the US. Many colleges and universities were encouraged to use the NSFNET and the number of sites increased rapidly.

Simultaneously, the Usenet originated in 1979. It allowed users to share information in the form of articles arranged into newsgroups. Usenet was developed as a sort of 'Poor man's ARPANET' for those UNIX users who were not allowed access to the ARPANET.

1.3.1 The Development of the ARPANET

One of the main goals of the ARPANET research was to develop a network whose communication would not be seriously impaired if physical sections of the network were lost. Also, the network needed to allow the addition and removal of new nodes with minimal impact, and allow computers of many different types to communicate easily.

One of the major impacts of the ARPANET research, and the one that led to today's Internet, was the development of the TCP/IP (Transmission Control Protocol/Internet Protocol) network protocol the language that computers connected to the network use to talk to one another. During the 1970's TCP/IP became the standard network protocol for the ARPANET. Also during this time, the US government began encouraging the educational community to take advantage of the ARPANET. The increasing number of users led to the development of many of the services available on to Internet today, including electronic mail (e-mail), files transfer, and remote login.

1.3.2 Host Growth

The growth of the Internet has been absolutely phenomenal, particularly over the last five years. The number of machines connected and the amount of traffic carried has grown tremendously, and the type of organizations connected has changed.

The number of hosts on the Internet has grown from 235 in May of 1982 to approximately 3.2 million hosts in July of 1994. The "edu" domain, which is for

educational and research organizations, has the most hosts (about 850,000). The commercial domain now has almost as many hosts (about 775,000).

The main “product” that the Internet gives you access to is information. Information is usually contained in files on an Internet Host, and can be presented in many different formats depending on what Internet service you are using and whether you are using that service through a GUI or through a terminal command-line interface. An Internet host can be a PC, Macintosh, UNIX-base workstation, or any computer that can speak TCP/IP. If you look at the information available on host through a simple command-line interface, you will probably need to know something about the underlying file structure for that type of host to find the information you are looking for. If you have GUI interface, you will probably not even know what type of host you are connecting to because the GUI interfaces tend to hide all of the low-level information about a host, making the access to information uniform for all machines.

Now that commercial organizations are connecting to the Internet, the number of services available is exploding. The World Wide Web is an Internet service that makes using the Internet as easy a clicking your mouse button. Many organizations are using the World Wide Web to provide access to their products and services. There are businesses that are planning to allow you to peruse their catalogs and order merchandise over the Internet. Most manufacturers have already begun providing product information over the Internet. The potential for information access is almost unlimited.

1.4 Probing the Internet

During the early 1980s, all the interconnected research networks were converted to the TCP/IP protocol, and the ARPANET became the backbone (the physical connection between the major sites) of the new Internet, which comprised all TCP/IP-based networks connected to the ARPANET. This conversion to TCP/IP was

completed by the end of 1983-and the Internet was born. In 1990 ARPANET ceased to exist and the network was turned over to the NSFNET.

When the Internet first came into existence in the early 1980s, there were only 213 registered hosts (computers that provided services) connected to the network. By February of 1986, there were 2308 hosts. Today, the Internet is undergoing tremendous growth, with several million hosts connected world-wide.

Table 2 (Major Landmarks in History of Internet)

Year	**Achievements**
1969	Department of Defense's Advanced Research Projects Agency (ARPA) contracted with Bolt, Baranek, and Newman (BBN) to develop a packet-switched network based on Paul Baran's ideas
1973	TCP/IP was proposed as a standard for ARPANET.
1975	Defense Communications Agency (DCA) changed name from ARPA to DARPA (Defense Advanced Research Projects Agency) and took control of ARPANET. Microcomputer introduced
1979	USENET was introduced
1983	Internet came in to existence
1990	ARPANET ceased to exist and the network was turned over to the NSFNET

1.4.1 Internet Administration

The Internet is not "owned" by anyone, in the usual sense of the word. The computer networks manage the Internet are managed by the Network Information Centre (NIC). The NIC is the Central administrative body for internet. To ensure inter-operability,

every host on internet needs a unique address. The NIC allocates IP network addresses and registers unique gateway and domain name. The NIC also archives the internet protocol specifications called the Request For Comments (RFC).

The backbone in the U.S. is funded by the National Science Foundation (NSF) and is supported from a technical standpoint by the Internet Engineering Task Force (IETF). The IETF is a committee of scientists and experts that works to resolve technical and related support issues for the Internet. There are regional and international segments of the network that have their own funding and administration. But, any network connected to the Internet agrees to the decisions and standards set forth by the Internet Architecture Board (IAB). Anyone who is willing to help may participate in the; process of devising and setting standards.

The reports of the IAB are made public throughout the publication of Request For Comment (RFC) documents. Some of these RFCs document Internet standards, but many of them are meant to introduce new ideas and stimulate discussion about future developments on the Internet. Past and current RFCs can be found at a number of places on the Internet. (A good place to look for them is ftp.internic.net).

1.4.2 Information Infrastructure

Access to a common network will facilitate the concept of telecommuting (working at home, using the network to access information, have video conferences, and so on) and tele-schooling (having students attend classes remotely, using a two-way live video conference, in addition to video broadcasts and online multimedia information and exercises). Companies are already beginning to take customer complaints and inquiries by e-mail, and distribute marketing materials and product updates online. All financial transactions could take place online, with currency becoming almost unnecessary.

Eventually, the Information Superhighway could completely change the structure of our society. The potential for information access through a common network like the

Internet is almost unlimited. In addition to business and educational activities, social forums could allow interaction between millions of people around the world, allowing people to explore other cultures and exchange information about topics of common interest .

1.4.3 Internet-Related Organizations

There are a number of organizations whose members are involved with educating people about the Internet or exploring topics important to the Internet. Some at these organizations are listed here.

- **Corporation for National Research Initiatives**- The Corporation for National Research Initiatives (CNRI) is a non-profit organization that was formed to encourage the co-operation of government, academic, and private industry in the development of a national data network. CNRI is involved in organizing many research projects, including faster transmission lines that will be able to carry live video broadcasts and graphics simulations; and know bots, programs similar to good computer viruses that would be able to search through the internet for information of interest to know bots "owner".
- **Internet Society –** The Internet Society (ISOC) is a non-profit organization that seeks to encourage the use and evolution of the Internet and to provide educational materials for, and a forum for discussion of the Internet Engineering Task Force. It holds and annual meeting that includes workshops and symposia on topics of interest to the membership. In addition, the Internet Society supports organizations involved in network security and Internet educational activities.
- **Computer Professionals for Social Responsibility –** The Computer Professional for Social Responsibility (CPSR) is an organization of people concerned about the ethical use of computers. Originally formed in 1983 by people concerned about the reliability of software developed for military

applications, CPSR members are concerned about many social issues that involve computers.

- **Electronics Frontier Foundation –** The Electronics Frontier Foundation (EFF) is another organization concerned about the social effects of computers on society. They, however, are particularly concerned about the legal rights of computer users. Most current laws do not specifically apply to electronic communications and are sometimes inappropriately applied to activities of computer users. EFF wants to help shape public policy in the emerging area of computer-based communications.

1.4.4 Addresses

Internet addresses are the key to using the Internet. You use mail address to send messages to other users, and you use host addresses (or host names) to retrieve files and connect to hosts that provide Internet services.

1.4.4.1 Host Name

All internet sites are identified by a unique domain. The domain name is made up of several pieces that identify the organization and the domain hierarchy to which it belongs. A host name contains domain name in addition to a name identifying the particular host and any sub-domain it may be associated with at its Internet site.

Host names are found in e-mail addresses and also are used when connecting to Internet hosts to use internet services (such as the World Wide Web) or retrieve files. A host name is made up of several words separated by periods. You can examine these words to find out information about the host. The host name parvinder.50megs.com is used here to illustrate the parts of a host name. The rightmost word, for example, specifies the domain of the machine. In this case, the word “com” means that the machine belongs to a commercial entity-a company of some kind. Some other domains are “edu” for educational institutions, “mil” for

military sites, and “gov” for sites that are part of government. Also, each country that is connected to the Internet has a domain assigned to it; for example “in” is the domain name for India. Examples of real-life institution names (including the domain name) are ibm.com for International Business machines, mit.edu for the Masschusetts Institute of Technology, and nasa.gov for the National Aeronautics and Space Administration.

1.4.4.2 IP Address

Host names are used to access individual hosts on the Internet. The host name is really just a convenient way for people to refer to hosts. The host name represents the IP address (or host address) of the host, which is the address that internet software needs to get information to or from the host. The IP address is the unique number assigned to identify the host on internet. This address is usually represented as four numbers between 1 and 254 separated by periods, for example, 209.102.34.1. Most software automatically translates host name to IP address, so you don’t have to remember which number represent which machine.

1.4.5 Various Services on the Internet

Various services available on internet are-

- Information retrieval services (FTP and Gopher)
- Information search services (Search Engines,WAIS, Archie, Veronica)
- Communication services (E-mail, Telnet, Usenet, IRC)
- Multimedia information services (World Wide Web)

1.5 Interpreting a Ping Response

Ping is a program used to test whether a particular network destination is online, by sending an Internet control message protocol (ICMP) echo request and waiting for a response (Also called packet internet gopher).

A simple way to think of a ping is it is like an echo-sound from a boat in the sea, it measures the time taken (in milliseconds) to get to the server/router and back. The basic ping command syntax is “ping hostname”.For the sake of example, we are going to ping Jolt.co.uk, a popular UK Gaming service which is hosted on the Nildram network. Follow the steps below.

Open a DOS window/Command Prompt window in Windows.

Once this window is open, type “ping www.jolt.co.uk” in to it.

The following should appear-

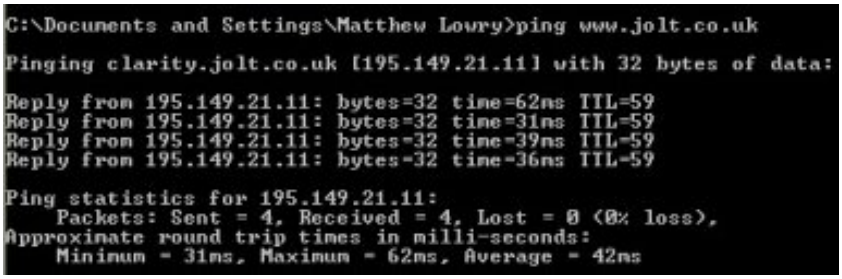

The Ping command can be used on your local network to see whether a particular computer is on-line or not, if reply comes it means computer on- line other- wise not.

TTL reply: Ping sends an ICMP echo packet (with the TTL value set to the host default) to the host listed on the ping command line. Ping expects back an ICMP ‘echo reply’ packet. The millisecond time displayed is the round trip time. The “TTL=245” above says that the incoming ICMP echo reply packet has its TTL field set to 245. Because this value was decremented by one at each hop on the way back, this tells us that visualroute.com is probably setting the initial TTL value to 255.

TTL Expired in Transit: Most computers today initialize the TTL value of outgoing IP Packets 128 or higher. If you ever see a reply above with a “TTL=5” (or some other low TTL number) this tells you that the computer being pinged should most likely have its default TTL value increased. Otherwise, anyone trying to communicate

with the computer that is at a hop count higher than the TTL will not be able to communicate with the computer. For example, if you are 40 hops away from www.xyz.com, and www.xyz.com sets TTL fields in IP packets that it sends out to 32, the IP Packets will not reach you. They will 'expire in transmit' before they reach you.

Discover your TTL: To discover the default TTL value of your computer, 'ping localhost' and examine the TTL reply value. For older Windows machines this value is 32. For newer Windows machines, this value is 128.

1.6 Tracing a Route

In Windows, tracert and in Unix, traceroute are tools used to trace the routes. Tool that shows you the network path between two locations. It shows you the address and how long it takes to get to each hop in the path. When there is a problem with the network, traceroute can often be used to narrow down where the problem is occurring. This form will do a traceroute from this server to another location. Enter the domain name or IP address of the other location in the text entry box.

The basic syntax in Windows is "trcert hostname". For the sake of example, we are going to perform a traceroute to Jolt.co.uk, a popular UK gaming service which is hosted on the Nildram network. Follow the steps below:

Open a DOS/Command Prompt window in Windows.

Once this window is open, type "tracert www.jolt.co.uk" into it

The following should appear:

Discover the path: Tracert sends an ICMP echo packet, but it takes advantage of the fact that most Internet routers will send back an ICMP 'TTL expired in transit' message if the TTL field is ever decremented to zero by a router. Using this knowledge, we can discover the path taken by IP Packets.

How tracert works?: Tracert sends out an ICMP echo packet to the named host, but with a TTL of 1; then with a TTL of 2; then with a TTL of 3 and so on. Tracert will then get 'TTL expired in transit' message back from routers until the destination host computer finally is reached and it responds with the standard ICMP 'echo reply' packet.

Try it yourself: To see this in action yourself, just use the '-i' option of ping, which allows you to set the TTL value of outgoing ping packets. For example, "ping –i 1 visualroute.com" and you will see "Reply from 199.70.3.58: TTL expired in transit" (where the router IP Address returned, 199.70.3.58, is specific to your Internet connection). Then again with "ping –i 2 visualroute.com", and get back "Reply from 199.70.3.49: TTL expired in transit", and so on. Finally at "ping –i 13 visualroute.com" you get "Reply from 192.41.43.189: bytes=32 time=198ms TTL=245", which is the destination host responding.

Round Trip Times: Each millisecond (ms) time in the table is the round-trip time that it took (to send the ICMP packet and to get the ICMP reply packet). The faster (smaller) the times the better. Ms times of 0 mean that the reply was faster than the computers timer of 10 milliseconds, so the time is actually somewhere between 0 and 10 milliseconds.

Packet Loss: Packet loss kills throughput. So, having no packet loss is critical to having a connection to the Internet to responds well. A slower connection with zero packet loss can easily outperform a faster connection with some packet loss. Also, packet loss on the last hop, the destination, is what is most important. Sometimes routers in-between will not send ICMP "TTL expired in transit" messages, causing

what looks to be high packet loss at a particular hop, but all it means is that the particular router is not responding to ICMP echo.

1.7 Self Assessment Questions (SAQA)

Check Yourself

Q1. Write the history of computer network and internet.

Q2. What is role of various internet related organizations?

Q3.Differentiate between host name and IP address.

Q4. Describe the following terms

i) IP ii) ICMP iii) EFF iv) NIC v) ISOC vi) TCP/IP vii) CCITT viii) WWW

Q5 Explain the following commands

i) Ping ii) trcert/traceroute

Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Dr. Pradeep K. Bhatia**
Lesson Number: 02	

MAC and Routing Protocols

Objective: The objective of present lesson is to discuss the concepts related to Media Access Control (MAC) Layer and routing protocols. Different switching techniques and how data is handled in frames are also discussed.

Structure

2.1 Packets

The Request For Comments (RFC) documents frequently use the term packet to mean a stream of binary octets (bytes) of data of some arbitrary length. It is typically used to describe chunks of data created by software, not by hardware. Internet Protocol (IP)

creates packets. The term packet is NOT synonymous with the term frame even though many people make that mistake.

2.2 Frames

The term frame is most frequently used to describe a chunk of data created by network communication hardware at the datalink / physical layer. There are ethernet frames, token- ring frames, FDDI frames etc. A frame is simply a chunk of data, frequently with some soft of pattern of bits at the start (called a header) and bits at the end (called a trailer). Frames are created by hardware protocols that do not have separate control circuits in the physical media to which they are attached. For example, it would be proper to say that Ethernet uses frames.

2.2.1 Datagram

A datagram is a chunk of data, usually containing some sort of message. This term is often used in describing protocols that function at higher levels of the OSI model (network layer and up). It is a less specific term than PDU. For example, it would be proper to say TCP uses datagrams.

2.2.2 Protocol Data Unit

A protocol data unit is a term used in much of the documentation and educational literature for networking technologies. It simply means a chunk of data created and/or labeled by the protocol being discussed. For example, documents describing the operation of spanning tree and routing protocols often refer to protocol data units.

2.2.3 Switching Concepts

Computer networks do not dedicate a single wire to each pair of communicating computers. Instead, the computers share the underlying hardware facilities. Although it is economic to share facilities, it can cause delays when the network is being

heavily used. To avoid long delays, network technologies limit the amount of data that a computer can transmit on each turn. This is called packet switching, and the unit of data that can be transferred is called a packet.

Both LANs and WANs use packet switching. Figure 1 illustrates a shared communication channel.

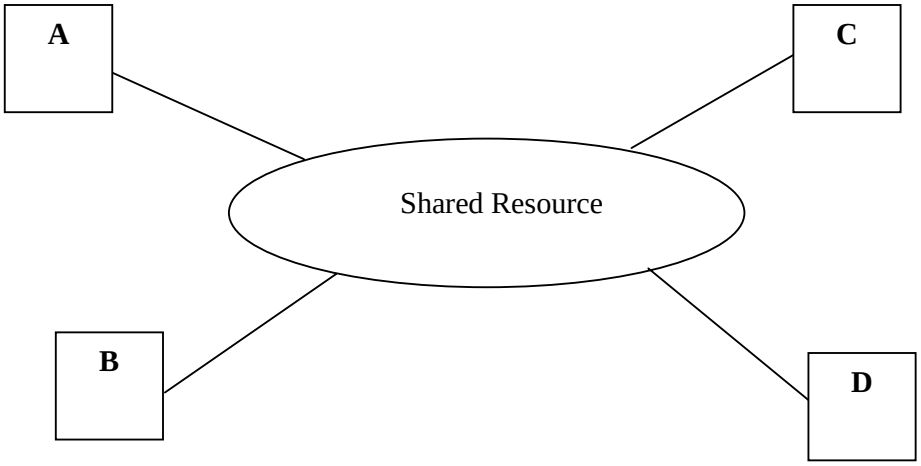

Figure 1 (Shared Communication Channel)

If A was allowed to send to B a 5 MB file (typical large data file) over the communication channel running at 56 kbps (typical long-distance network), then the transfer will take approximately 12 minutes. This is how long C and D wait to use the channel.

In contrast, consider the situation where transfers take place using 1,000 byte packets. A sends its first packet to B - the transfer takes 143 ms. C or D (or even B) can now transmit. There are no long delays.

2.2.4 Time-division Multiplexing

Conceptually, a network that permits sources to take turns to access a communication channel is providing a form of time-division multiplexing, as illustrated in Figure 2.

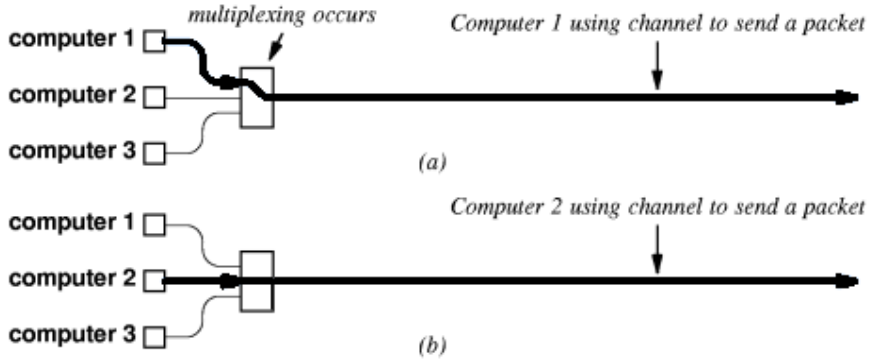

Figure 2(Time-Division Multiplexing)

In (a) computer 1 uses the channel, and then in (b) computer 2 uses the channel. Dividing the data into small packets allows all sources to receive prompt service.

In most packet switching networks, transfers occur quickly. A typical LAN can transfer a thousand packets between two computers every second. To a human, events that require thousandths of a second appear instantaneous. For example, several people can use computers connected to a single network without perceiving any delay. The network hardware handles the sharing of the network automatically. It does not require any computation or co-ordination. A computer generates a packet at any time. When the packet is ready, the computer's interface hardware waits its turn before transmitting the packet.

2.2.5 Data in Frame

The term packet switching is a generic term for the concept of dividing the data into small blocks for transmission. For a specific hardware technology, the term frame is used to describe the specific format to be used.

Suppose one host needs to send a block of data to another host using the RS-232 character transmission mechanism - this does not include a mechanism for the sender to signal the end of a block. The sender and receiver must agree on the mechanism. Figure 3 illustrates a scheme where the ASCII characters soh (start of header) and eot (end of transmission) are used to delimit the frame.

Figure 3

The mechanism is simple and unambiguous. The only disadvantage is the overhead of generating two additional characters per block of data.

2.2.6 Byte Stuffing

The sender and receiver have agreed to use the characters soh and eot to delimit the frame. Therefore these characters cannot appear as data within the block. A technique called byte stuffing resolves this problem by allocating a third character to mark occurrences of reserved characters in the data. The ASCII character esc is usually used, as illustrated in Figure 4.

Character in Data	**Character Sent**
SOH	ESC X
EOT	ESC Y
ESC	ESC Z

Figure 4

For each occurrence of a character in the left column, the sender transmits the two characters in the right column. Figure 5 illustrates a complete data block.

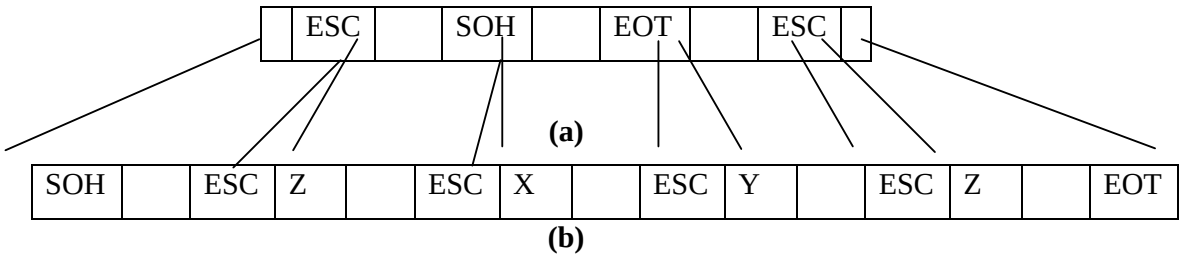

Figure 5

In (a) we have a data block containing the special framing characters. In (b) we have the complete frame after byte stuffing. In most communication systems, the characters x, y, and z in the above figure are soh, eot, and esc respectively.

2.2.7 Transmission Errors

Electro-magnetic interference can introduce unwanted electrical currents into the electronic components or wires used for communication. Severe interference (e.g. lightning) can cause permanent damage to network equipment. More often, interference changes the electrical signal without damaging the equipment. This can cause the receiver to misinterpret one or more bits of data. Lost, changed, or spuriously appearing bits, collectively called transmission errors, account for extra complexity required in computer networks.

2.2.8 Parity Bits

RS-232 circuits can use a parity bit to ensure that each character arrives intact. The sender computes an additional bit based on the seven ASCII data bits and transmits this as an eighth bit. The receiver performs the same computation and verifies that the

parity bits agree. The computation is chosen so that a one-bit alteration can be detected.

Parity can be even or odd; for even (odd) parity the sender sets the parity bit so that the total number of 1 bits is even (odd).

2.2.9 Checksums

Many computer networks send a checksum along with each packet to help the receiver detect errors. To compute a checksum, the sender treats the data as a sequence of binary integers and computes their sum, as illustrated in Figure 6.

H	**e**	**l**	**l**	**o**		**w**	**o**	**r**	**l**	**d**	**.**
48	65	6C	6C	6F	20	77	6F	72	6C	64	2e

4865 + 6C6C + 6F20 + 776F + 726C + 642E + CARRY = 71FC

Figure 6

Each pair of characters is treated as a 16-bit integer. If the sum overflows 16 bits, the carry bits are added to the total. The advantage of such checksums is the size and ease of computation. Addition requires very little computation and the cost of sending an additional 16-bits is negligible.

However, checksums do not detect all common errors, as illustrated in Figure 7.

Data Item in Binary	Checksum Value	Data Item in Binary	Checksum Value
0001	1	0011	**3**
0010	2	0000	**0**
0011	3	0001	**1**
0001	1	0011	**3**
Totals	7		7

Figure 7

A transmission error has reversed the second bit in each of the four data items, yet the checksums are identical.

2.2.10 Cyclic Redundancy Checks

It is possible to detect more errors without increasing the amount of additional information in the packet using a Cyclic Redundancy Check (CRC) technique. Moreover, the hardware necessary to generate them is very simple. The two necessary components are an exclusive or (XOR) unit and a shift register. Figure 8 illustrates the XOR hardware and truth table.

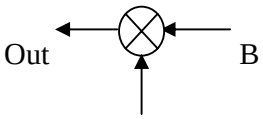

(a)

A	B	Out
0	0	0
0	1	1
1	0	1
1	1	0

(b)

Figure 8

In (a), we have the diagram for the hardware which computes the XOR. In (b), we have the output value for each of the four combinations of input values. The second component, a shift register, holds a fixed number of bits and each time a new bit enters (on the right) the contents of the register shift left and the leftmost bit becomes the output. This is illustrated in Figure 9.

Figure 9

In (a) we have the position before the shift and in (b) the position after the shift. The shift register can be initialized by setting all bits to zero.

Figure 10 illustrates how three XOR units and three shift registers can be combined to form a 16-bit CRC. The hardware is inexpensive and easy to construct.

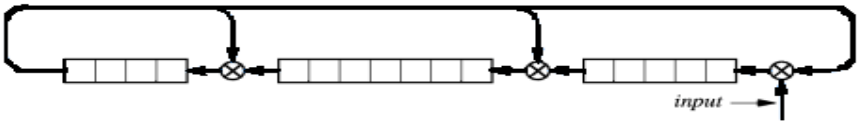

Figure10

Output from the leftmost shift register goes to the three XOR units simultaneously. To compute a CRC, the shift registers are initialized to zero, and the bits of the message are shifted in one at a time. After an entire message has been shifted into the unit, the shift registers contain the 16-bit CRC for the message. The receiver uses identical hardware to calculate a CRC to verify that it agrees with the sender's CRC.

To simplify checking, a small modification is made to the above algorithm. When computing the CRC, the sender appends an additional sixteen zeros to the message. Mathematically, the sixteen zeros cause the resulting CRC to be an inverse. This is useful for the receiver, who computes the CRC over the incoming message and the incoming CRC. If all bits are correctly received, the computed value will be zero.

The mathematics of CRCs are too complex to be covered here. The 16-bit CRC used above will detect all single-bit errors, all double-bit errors, all errors containing an

odd number of bits, and all burst errors of length 16 or less (first and last bits are one), as well as a number of additional errors. Detecting burst errors is important, as electrical interference (lightning, and other sparks) often produces burst errors.

2.3 Error Detection and Correction

2.3.1 Detection

Networks usually handle error detection at the frame level. The sender computes a checksum or CRC and transmits this with each frame. The receiver performs the same calculation and uses the result to validate the incoming frame. A simple frame format is illustrated in Figure 11.

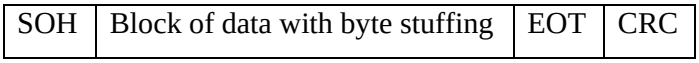

Figure 11

The modified frame format includes a 16-bit CRC. Different standards define whether the CRC is computed over the original message or the stuffed frame.

2.3.2 Error Correction

Error correction can be used on simplex channels, where retransmission cannot be requested, but most often error detection followed by retransmission is preferred because it is more efficient.

A simple method for correcting single bit errors is to transmit each bit three times.

0 -> 000

1 -> 111

If 001, 010, or 100 is received, the original message was 0. If 110, 101, or 011 is received, the original message was 1.

It is possible to do better than this. Hamming codes are used to correct single bit errors (8 data bits require 4 extra bits to be transmitted).

2.4 WAN (Wide Area Network) Technologies

2.4.1 What is a WAN?

A WAN is a data communications network that covers a relatively broad geographic area and that often uses transmission facilities provided by common carriers, such as telephone companies. WAN technologies generally function at the lower three layers of OSI reference model: the physical layer, the data link layer, and the network layer. Figure 12 illustrates the relationship between the common WAN technologies and the OSI model.

Figure 12

2.4.2 Point-to-Point Links

A point-to-point link provides a single, pre-established WAN communications path from the customer premises through a carrier network, such as a telephone company, to a remote network. Point-to-point lines are usually leased from a carrier and thus are often called leased lines. For a point-to-point line, the carrier allocates pairs of wire and facility hardware to your line only. These circuits are generally priced based on bandwidth required and distance between the two connected points. Point-to-point links are generally more expensive than shared services such as Frame Relay. Figure 13 illustrates a typical point-to-point link through a WAN.

Figure 13 (Point to Point Link through WAN)

2.4.3 Circuit Switching

Switched circuits allow data connections that can be initiated when needed and terminated when communication is complete. This works much like a normal telephone line works for voice communication. Integrated Services Digital Network (ISDN) is a good example of circuit switching. When a router has data for a remote site, the switched circuit is initiated with the circuit number of the remote network. In the case of ISDN circuits, the device actually places a call to the telephone number of the remote ISDN circuit. When the two networks are connected and authenticated, they can transfer data. When the data transmission is complete, the call can be terminated. Figure 14 illustrates an example of this type of circuit.

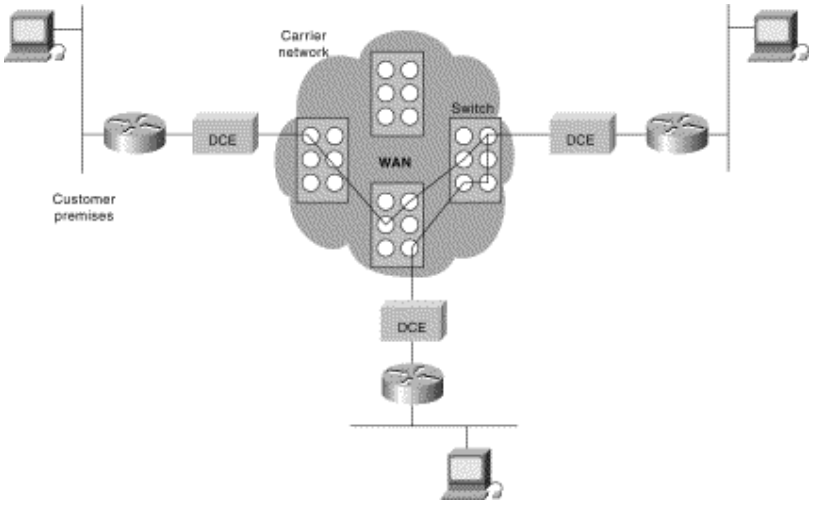

Figure 14 (Circuit Switching)

2.4.4 Packet Switching

Packet switching is a WAN technology in which users share common carrier resources. Because this allows the carrier to make more efficient use of its infrastructure, the cost to the customer is generally much better than with point-to-point lines. In a packet switching setup, networks have connections into the carrier's network, and many customers share the carrier's network. The carrier can then create virtual circuits between customers' sites by which packets of data are delivered from one to the other through the network. The section of the carrier's network that is shared is often referred to as a cloud.

Some examples of packet-switching networks include Asynchronous Transfer Mode (ATM), Frame Relay, Switched Multi-megabit Data Services (SMDS), and X.25. Figure

15 show an example packet-switched circuit.

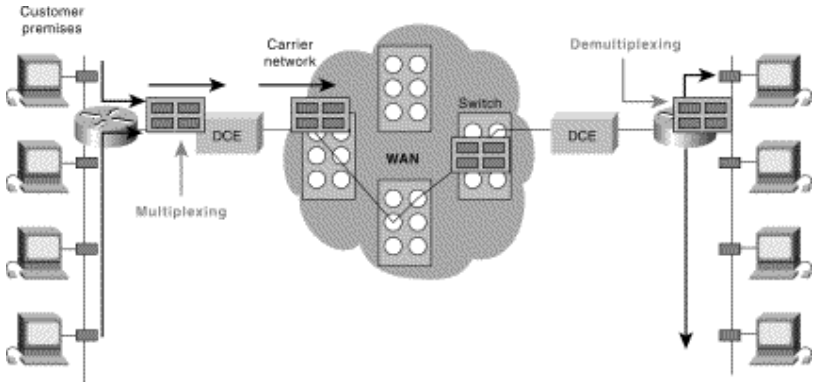

Figure 15 (Packet Switching Circuit)

2.4.5 WAN Virtual Circuits

A virtual circuit is a logical circuit created within a shared network between two network devices. Two types of virtual circuits exist: switched virtual circuits (SVCs) and permanent virtual circuits (PVCs).

SVCs are virtual circuits that are dynamically established on demand and terminated when transmission is complete. Communication over an SVC consists of three phases: circuit establishment, data transfer, and circuit termination. The establishment phase involves creating the virtual circuit between the source and destination devices. Data transfer involves transmitting data between the devices over the virtual circuit, and the circuit termination phase involves tearing down the virtual circuit between the source and destination devices. SVCs are used in situations in which data transmission between devices is sporadic, largely because SVCs increase bandwidth used due to the circuit establishment and termination phases, but they decrease the cost associated with constant virtual circuit availability.

PVC is a permanently established virtual circuit that consists of one mode: data transfer. PVCs are used in situations in which data transfer between devices is constant. PVCs decrease the bandwidth use associated with the establishment and

termination of virtual circuits, but they increase costs due to constant virtual circuit availability. PVCs are generally configured by the service provider when an order is placed for service.

2.4.6 WAN Devices

WANs use numerous types of devices that are specific to WAN environments. WAN switches, access servers, modems, CSU/DSUs, and ISDN terminal adapters are discussed in the following sections. Other devices found in WAN environments that are used in WAN implementations include routers, ATM switches, and multiplexers.

2.4.7 WAN Switch

A WAN switch is a multi-port internetworking device used in carrier networks. These devices typically switch such traffic as Frame Relay, X.25, and SMDS, and operate at the data link layer of the OSI reference model. Figure 16 illustrates two routers at remote ends of a WAN that are connected by WAN switches.

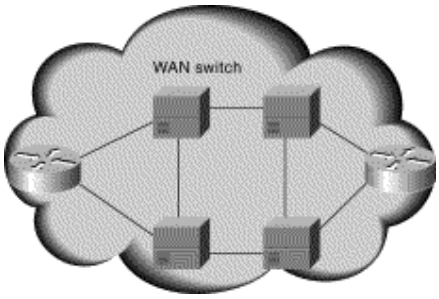

Figure 16

2.4.8 Access Server

An access server acts as a concentration point for dial-in and dial-out connections. Figure 17 illustrates an access server concentrating dial-out connections into a WAN.

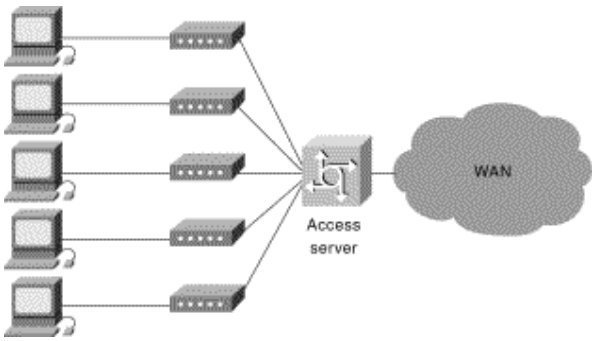

Figure 17

2.4.9 Modem

A modem is a device that interprets digital and analog signals, enabling data to be transmitted over voice-grade telephone lines. At the source, digital signals are converted to a form suitable for transmission over analog communication facilities. At the destination, these analog signals are returned to their digital form. Figure 18 illustrates a simple modem-to-modem connection through a WAN.

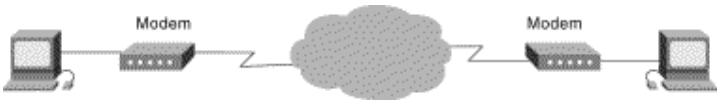

Figure 18

2.4.10 CSU/DSU

A channel service unit/digital service unit (CSU/DSU) is a digital-interface device used to connect a router to a digital circuit like a T1. The CSU/DSU also provides signal timing for communication between these devices. Figure 19 illustrates the placement of the CSU/DSU in a WAN implementation.

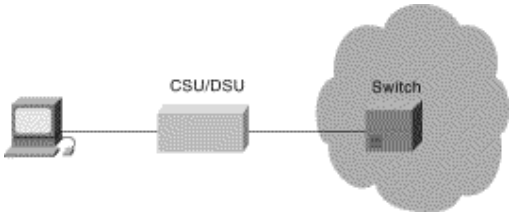

Figure 19

2.4.11 ISDN Terminal Adapter

An ISDN terminal adapter is a device used to connect ISDN Basic Rate Interface (BRI) connections to other interfaces, such as EIA/TIA-232 on a router. A terminal adapter is essentially an ISDN modem, although it is called a terminal adapter because it does not actually convert analog to digital signals. Figure 20 illustrates the placement of the terminal adapter in an ISDN environment.

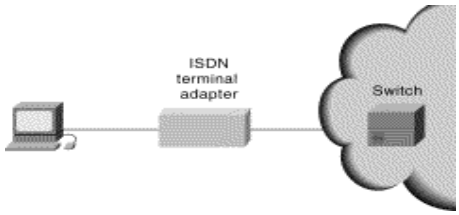

Figure 20

2.5 Routing

Routing is the act of moving information across an inter-network from a source to a destination. Along the way, at least one intermediate node typically is encountered. Routing is often contrasted with bridging, which might seem to accomplish precisely the same thing to the casual observer. The primary difference between the two is that bridging occurs at Layer 2 (the link layer) of the OSI reference model, whereas routing occurs at Layer 3 (the network layer). This distinction provides routing and

bridging with different information to use in the process of moving information from source to destination, so the two functions accomplish their tasks in different ways.

The topic of routing has been covered in computer science literature for more than two decades, but routing achieved commercial popularity as late as the mid-1980s. The primary reason for this time lag is that networks in the 1970s were simple, homogeneous environments. Only recently large-scale inter-networking becomes popular.

2.5.1 Routing Components

Routing involves two basic activities: determining optimal routing paths and transporting information groups (typically called packets) through an inter-network. In the context of the routing process, the latter of these is referred to as packet switching. Although packet switching is relatively straightforward, path determination can be very complex.

2.5.2 Path Determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which contain route information. Route information varies depending on the routing algorithm used.

Routing algorithms fill routing tables with a variety of information. Destination/next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing the "next hop" on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with a next hop. Figure 21 depicts a sample destination/next hop routing table.

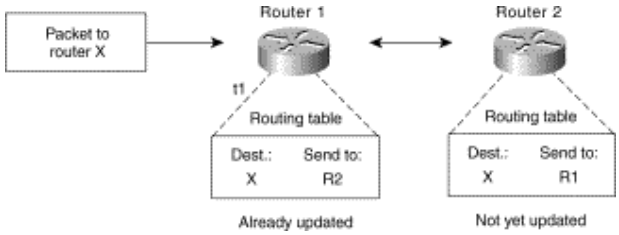

Figure 21

Routing tables also can contain other information, such as data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used. A variety of common metrics will be introduced and described later in this chapter.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender's links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

2.5.3 Routing Algorithm Types

Routing algorithms can be classified by type. Key differentiators include these:

- Static versus dynamic
- Single-path versus multi-path
- Flat versus hierarchical
- Host-intelligent versus router-intelligent
- Intra-domain versus inter-domain
- Link-state versus distance vector

2.5.3.1 Static versus Dynamic

Static routing algorithms are hardly algorithms at all, but are table mappings established by the network administrator before the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, they generally are considered unsuitable for today's large, constantly changing networks. Most of the dominant routing algorithms today are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a router to which all un-routable packets are sent), for example, can be designated to act as a repository for all un-routable packets, ensuring that all messages are at least handled in some way.

2.5.3.2 Single-Path versus Multi-path

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multi-path algorithms permit traffic multiplexing over multiple lines. The advantages of multi-path algorithms are obvious: They can provide substantially better throughput and reliability. This is generally called load sharing.

2.5.3.3 Flat versus Hierarchical

Some routing algorithms operate in a flat space, while others use routing hierarchies. In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from nonbackbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more non-backbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas. In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intra-domain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

2.5.3.4 Host-Intelligent versus Router-Intelligent

Some routing algorithms assume that the source end node will determine the entire route. This is usually referred to as source routing. In source-routing systems, routers merely act as store-and-forward devices, mindlessly sending the packet to the next stop.

Other algorithms assume that hosts know nothing about routes. In these algorithms, routers determine the path through the inter-network based on their own calculations.

In the first system, the hosts have the routing intelligence. In the latter system, routers have the routing intelligence.

2.5.3.5 Intra-domain versus Inter-domain

Some routing algorithms work only within domains; others work within and between domains. The nature of these two algorithm types is different. It stands to reason, therefore, that an optimal intra-domain-routing algorithm would not necessarily be an optimal inter-domain routing algorithm.

2.5.3.6 Link-State versus Distance Vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the inter-network. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors.

Because they converge more quickly, link-state algorithms are somewhat less prone to routing loops than distance vector algorithms. On the other hand, link-state algorithms require more CPU power and memory than distance vector algorithms. Link-state algorithms, therefore, can be more expensive to implement and support. Link-state protocols are generally more scalable than distance vector protocols.

2.5.4 Network Protocols

Routed protocols are transported by routing protocols across an inter-network. In general, routed protocols in this context also are referred to as network protocols.

These network protocols perform a variety of functions required for communication between user applications in source and destination devices, and these functions can differ widely among protocol suites. Network protocols occur at the upper five layers of the OSI reference model: the network layer, the transport layer, the session layer, the presentation layer, and the application layer.

Confusion about the terms routed protocol and routing protocol is common. Routed protocols are protocols that are routed over an inter-network. Examples of such protocols are the Internet Protocol (IP), DECnet, AppleTalk, Novell NetWare, OSI, Banyan VINES, and Xerox Network System (XNS). Routing protocols, on the other hand, are protocols that implement routing algorithms. Put simply, routing protocols are used by intermediate systems to build tables used in determining path selection of routed protocols. Examples of these protocols include Interior Gateway Routing Protocol (IGRP), Enhanced Interior Gateway Routing Protocol (Enhanced IGRP), Open Shortest Path First (OSPF), Exterior Gateway Protocol (EGP), Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Routing Information Protocol (RIP).

2.6 Network Ownership

Network ownership can be classified in two categories as Private or Public.

2.6.1 Private Network

Features of private networks are-

– Owned by individual or corporation

– Restricted to owner's use

– Typically used by large corporations

2.6.2 Public

Features of public networks are-

– Owned by a common carrier

– Individuals or corporations can subscribe

– ‘‘Public’’ refers to availability, not data.

2.6.3 Advantages and Disadvantages

Private

– One has complete control on it.

– Installation and operation costs are high.

Public

– No need for staff to install / operate network.

– It is having disadvantage of Dependency on carrier

– It require subscription fee.

2.6.4 Public Network Connections

Public network connections can be –

2.6.4.1 One connection per subscriber

– Typical for small corporation or individual.

– Communicate with another subscriber.

2.6.4.2 Multiple connections per subscriber

– Typical for large, multi-site corporation

– Communicate among multiple sites as well as with another subscriber

2.6.5 Virtual Private Network

Virtual Private Network is a service provided over public network. This interconnects multiple sites of Single Corporation. It acts like private network in which no packets sent to other subscribers and no packets received from other subscribers. Data is encrypted for security point of view.

2.7 Service Paradigm

Fundamental characteristic of a network is that it is understood by hardware and visible to applications.

Two basic types of networks according to this characteristic are-

– Connectionless

– Connection-oriented

2.7.1 Connectionless (CL)

In Connectionless network, Sender forms packet to be sent, places address of intended recipient in packet and transfers packet to network for delivery. Network uses destination address to forward packet and delivers to recipient.

Characteristics of Connectionless Networks are-

- Packet contains identification of destination.
- Each packet handled independently.
- No setup required before transmitting data.
- No cleanup required after sending data.
- We can think it as postcards.

2.7.2 Connection-Oriented (CO)

In Connection Oriented network, Sender requests connection to receiver. Then it waits for network to form connection. It leaves connection in place while sending data

and terminates connection when no longer needed. Network receives connection request, forms path to specified destination and informs sender. Network transfers data across connection and removes connection when sender requests. We can think it as telephone calls.

2.7.3 Comparison of Connection Oriented and Connection Less

Connection Oriented

– More intelligence in network

– Can reserve bandwidth

– Connection setup overhead

– State in packet switches

– Well-suited to real-time applications

Connection Less

– Less overhead

– Permits asynchronous use

– Allows broadcast/multicast

Examples of Service Paradigm

Technology	**CO**	**CL**	**Used for LAN**	**Used for WAN**
Ethernet	No	Yes	Yes	No
Wi-Fi	No	Yes	Yes	No
Frame Relay	Yes	No	No	Yes
SMDS	No	Yes	No	Yes
ATM	Yes	No	Yes	Yes
Local Talk	No	Yes	Yes	No

2.8 Performance

Two Primary Performance Measures are:

- Delay
- Throughput

2.8.1 Delay

Delay is time required for one bit to travel through the network. It can be classified in to three types according to causes.

– Propagation delay

– Switching delay

– Queuing delay

Intuition: ‘‘length’’ of the pipe

2.8.2 Throughput

Throughput is number of bits per second that can be transmitted. It is measured with reference to capacity.

Intuition: ‘‘width’’ of the pipe.

2.8.3 Components of Delay

Fixed (nearly constant)

– Propagation delay

– Switching delay

Variable

– Queuing delay

– Depends on throughput

2.8.4 Relationship between delay and throughput

When network is idle, queuing delay is zero. As load on network increases queuing delay rises. Here load is defined as ratio of throughput to capacity, called utilization.

2.8.5 Relationship between delay and utilization

Let us, define:

– D0 to be the propagation and switching delay.

– U to be the utilization.

– D to be the total delay.

Then $D = D_0 / (1-U)$.

High utilization is known as congestion. Practical consequence is - if any network that operates with a utilization approaching 100% of capacity is doomed.

2.8.6 Delay-Throughput Product

– Measured in bits

– Gives quantity of data ‘‘in transit’’

2.8 Self Assessment Questions (SAQA)

Q1. Explain different components of Frames?

Q2. Compare connection oriented and connection less network services

Q3.What is routing? Explain different routing algorithms

Q4. Differentiate circuit switching and packet switching

Q5. Explain different WAN technologies.

Q6. Write short notes on –

i) Byte Stuffing ii) Error correction and detection iii) Time division Multiplexing

Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Rishipal Singh**
Lesson Number: 03	

TCP/IP Protocols

Objective: The objective of this lesson is to understand the role of TCP/IP in working of Internet. This lesson describes the different ways to connect a client computer to internet, Internetworking, architecture and protocols such as TCP and IP in detail. The future of IP is also discussed.

Structure

3.1 Internet Working Concepts

By connecting local, regional, national and international networks, the Internet forms the world's largest network of networks. Computers connected to the Internet work together to transfer information around the world using servers and clients. A server is a computer that manages the resources on a network and provides a centralized

storage for programs and data. The server, also called a host, "serves" files and services out to clients, computers that access the contents of the server. The speed at which data can be transferred between the sender and receiver in a network is called the data transfer rate. Transfer rates are expressed in bits per second (bps). A bit, short for binary digit is the smallest unit of computerized data. A bit has a value of either 1 or 0 that the computer interprets as "on" or "off" respectively. When calculating Internet-access speeds, it is important to recognize the difference between bits and bytes. There are eight bits in a byte. The small "b" stands for bits, and the big "B" stands for bytes. Transfer speeds are often shown in kilobytes per second (KB/s), and connect speeds are usually quoted in kilobits per second (kbps). For example, if the Web browser is downloading a file at 100 KB/s over a cable modem connection, that is equal to a speed of 800 kbps. Bandwidth is the width of the communication channel, the total volume of traffic that can be transferred across a given transmission line in a given period of time, usually measured in bits per second.

The most popular ways in world to connect a client computer to the Internet are: a dial-up connection using a telephone line or an Integrated Services Digital Network (ISDN), a Digital Subscriber Line (DSL), a cable TV connection and a satellite connection.

3.1.1 Dial-up Connection

A dial-up connection uses the analog telephone line for establishing a temporary communication. Computer's digital signals must be converted to analog signals before they are transmitted over standard telephone lines. This conversion is performed by a modem, a device that modulates (changes into an analog signal) and demodulates (converts an analog signal into a digital signal). Both the sending and receiving ends of a communication channel must have a modem for data transmission to occur. Using a dial-up line to transmit data is similar to using the telephone to make a call. The client computer modem dials the preprogrammed phone number for a user's

Internet Service Provider (ISP) and connects to one of the ISP's modems. After the ISP has verified the user's account, a connection is established and data can be transmitted. When either modem hangs up, the communication ends. The advantage of a dial-up line is that it costs no more than a local telephone call. Computers at any two locations can establish a connection using modems and a telephone network, to include wireless modems and cellular telephone connections. The limitation of a connection using the ordinary telephone line is a low speed, 28 kbps. There are dedicated telephone lines that can transmit data at 56 kbps. Most 56 kbps modems connect at a speed less than 46 kbps because of the limitations of analog phone lines and telephone company switches.

3.1.2 ISDN

ISDNs are special digital telephone lines that can be used to dial into the Internet at speeds ranging from 64 to 256 kbps. These types of connections are not available everywhere. Telephone companies have to install special ISDN digital switching equipment. ISDNs require use of a special "digital modem" that sends and receives digital signals over ISDN lines. With an ISDN, the telephone line is divided into three channels (BRI - Basic Rate Interface), two-64 kbps B (bearer) channels that send data and one 16 kbps D (data) channel that sends routing information. This type of access is commonly referred to as 2B+D. To use the ISDN access to the Internet, an ISP has to offer the ISDN access. ISDN lines cost more than normal phone lines, so the telephone rates are usually higher.

3.1.3 Cable TV Connection

Currently most households with cable TV have the option for cable modem Internet access. The cable modem offers a high-speed link at low cost for unlimited, "always connected" access. The connection speeds range from 128 kbps up to 10 mbps (megabits per second). A cable modem is a device that connects to the existing TV

cable feed and to an Ethernet network card in the user's PC (also called an NIC - Network Interface Card). The cable network is designed to support the highest speeds in the "downstream" direction, which is from the Internet to the client computer. This downstream speed affects the performance of downloading Web pages and software. The "upstream" bandwidth for data sent from a user's computer to the Internet is typically less, in the range of 200 kbps to 2 mbps. The benefit of the cable modem for Internet access is that, unlike DSL, its performance doesn't depend on distance from the central cable office. However, with the cable TV network, the computer is put on a Local Area Network (LAN) with other users in the neighborhood and like with any LAN, the performance degrades as usage increases. A more disturbing issue is that of network security. One of the main purposes of a LAN is to allow file sharing among the computers on the LAN. This LAN feature doesn't work well with cable Internet access, as most users do not want neighbors accessing their files. Turning the sharing option off can prevent file sharing. Also, installing the firewall hardware or software may protect from hackers.

3.1.4 DSL (Digital Subscriber Line)

DSL service is a high-speed data service that works over POTS (Plain Old Telephone Service) copper telephone lines and is typically offered by telephone companies without costly installation of a higher-grade cable. DSL uses a different part of the frequency spectrum than analog voice signals, so it can work in conjunction with a standard analog telephone service, providing separate voice and data "channels" on the same line. ADSL (Asymmetric DSL) is the type of DSL that provides different bandwidths in the upstream and downstream directions, giving the user a much bigger "pipe" in the downstream direction. ADSL can support downstream bandwidths of up to 8 mbps and upstream bandwidths of 1.5 mbps. For comparison, a T-1 connection also provides 1.5 mbps. This scheme works well for the typical Internet user;

upstream communication is usually small (link requests) compared to downstream communication (Web pages with graphics).

SDSL (Symmetric DSL) offers the same bandwidth capability in both directions. Besides higher bandwidth, some of the advantages of ADSL access from telephone companies are that there are no per-minute charges and the user gets an "always-on" connection for a monthly fee. Most modern computers can be easily equipped to connect to a DSL service. This is accomplished by connecting an ADSL modem to an Ethernet network card in the PC. The downside of DSL includes strict distance limitation that DSL circuits can operate within. As the connection's length increases, the signal quality decreases and the connection speed goes down. DSL services that provide greater that 1.5 mbps require shorter distances to the central office compared to a cable modem that can be located far away from the service provider.

The limit for ADSL service is 18,000 feet (5,460 meters), though for speed and quality of service reasons many ADSL providers place a lower limit on the distances for the service. At the extremes of the distance limits, ADSL customers may see speeds far below the promised maximums, while customers near the central office have the potential for very high speeds. Unlike cable modem technology, DSL provides a point-to-point connection to ISP. DSL proponents claim this technology is both more secure and less prone to local traffic fluctuations than its cable rival. By not sharing a LAN segment with other users, the systems are not as open to intrusion or susceptible to performance degradations related to local traffic.

3.1.5 Satellite Connection

Getting the Internet feed from a satellite is really not all that different from getting TV signals from one. In both cases data is being sent from the satellite to a user's equipment and then translated and decoded. One major limitation of satellite technology is that it can only send data from the satellite to a user's receiver—not the

other way. To get around this problem, a separate ISP connection is needed to send data to the Internet, typically over an analog modem. This connection works in conjunction with the satellite feed. As information is requested via the modem line, data are sent back via the satellite. Since most Internet users need high bandwidth from the Web, downstream (typically Web pages and file downloads), and less bandwidth going to the Web, upstream (typically link requests), this scenario of sending upstream data over a standard modem line and downstream data over the high-bandwidth satellite feed has been effective. The newest satellite technology allows for two-way communications and higher upstream bandwidths. A satellite return channel can be added for traffic bound for the Internet. The upload speeds through this satellite return channel may peak at 128 kbps. Download speeds with this system are up to 400 kbps. Satellite technology has one strong advantage over cable modems and DSL: accessibility. For many it is today's only high-speed option. It can reach areas that are otherwise difficult to establish contact with. The infrastructure exists to provide 400 kbps downstream bandwidth to almost anyone with a 21" satellite dish. It is eight times faster than fastest analog telephone modems and three times faster than ISDN. However, it is not as fast as cable modems or DSL services, which both can provide more than megabits of bandwidth. Also, cable and DSL access methods are cheaper. Equipment required for satellite connection includes installation of a mini-dish satellite receiver and a satellite modem. Like cable modem systems, satellite provides a "shared bandwidth" pipe. This means that download performance may vary depending upon other users of the satellite transponder. Another potential problem can be associated with severe weather. In severe snowstorms and heavy rain, users may experience signal fade.

The general rule about the Internet connection is the faster, the better. The bandwidth and transfer rate determine how quickly pictures, sounds, animation and video clips will be downloaded. Since multimedia and interactivity make the Internet such an exciting tool for information sharing, the speed is the key. Dial-up access provides an

easy and inexpensive way for users to connect to the Internet, however, it is a slow-speed technology and most users are no longer satisfied with dial-up or ISDN connections. Fortunately, the broadband access, we once dreamed of, is now possible with TV cable, DSL and satellite links.

3.2 Architecture and Protocols

In this section, we start by introducing TCP/IP and describing its basic properties such as internetworking, protocol layering and routing. We then discuss each of the specific protocols in more detail.

3.2.1 Architectural Model

The TCP/IP protocol suite is named for two of its most important protocols: Transmission Control Protocol (TCP) and Internet Protocol (IP). Another name for it is the Internet Protocol Suite, and this is the phrase used in official Internet standards documents.

3.2.2 Internetworking

The first design goal of TCP/IP was to build an interconnection of networks that provided universal communication services: an inter-network, or internet. Each physical network has its own technology-dependent communication interface, in the form of a programming interface that provides basic communication functions (primitives). Communication services are provided by software that runs between the physical network and the user applications and that provides a common interface for these applications, independent of the underlying physical network. The architecture of the physical networks is hidden from the user.

The second aim is to interconnect different physical networks to form what appears to the user to be one large network. Such a set of interconnected networks is called an inter-network or an internet.

To be able to interconnect two networks, we need a computer that is attached to both networks and that can forward packets from one network to the other; such a machine is called a router. The term IP router is also used because the routing function is part of the IP layer of the TCP/IP protocol suite. Figure 1 shows examples of two internet.

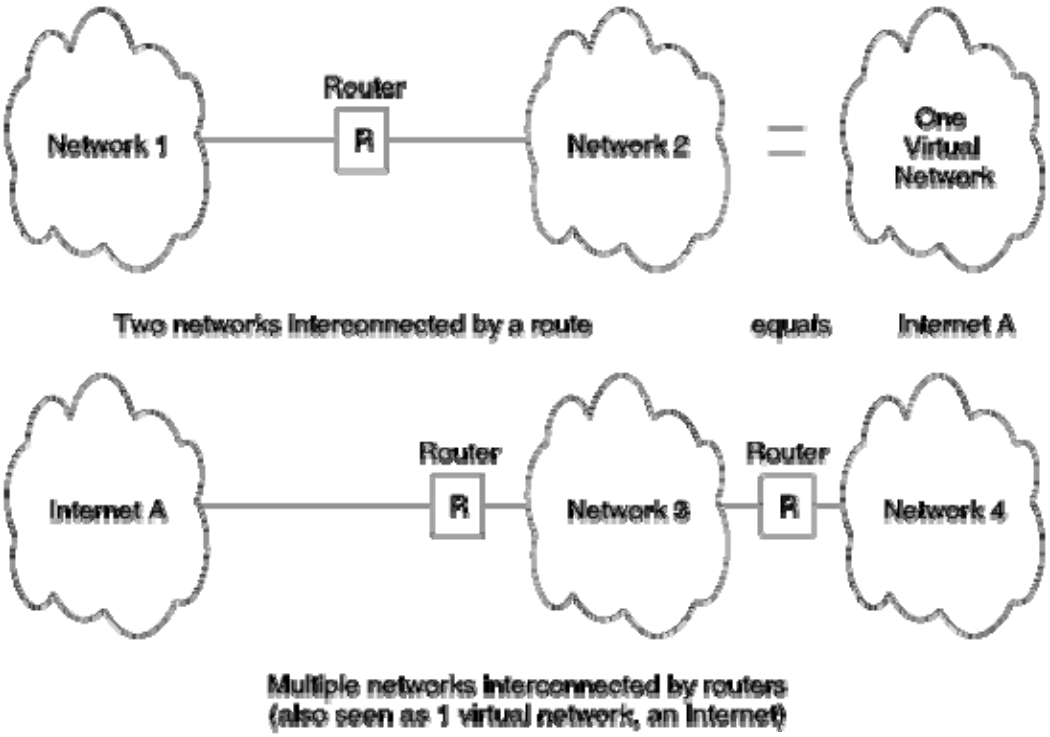

Figure 1

The basic properties of a router are:

- From the network standpoint, a router is a normal host.
- From the user standpoint, routers are invisible. The user sees only one large inter-network.

To be able to identify a host on the inter-network, each host is assigned an address, the IP address. When a host has multiple network adapters, each adapter has a separate IP address. The IP address consists of two parts:

IP address = <network number><host number>

The network number part of the IP address is assigned by a central authority and is unique throughout the Internet. The authority for assigning the host number part of the IP address resides with the organization which controls the network identified by the network number. The addressing scheme is described in detail in Addressing.

3.2.3 Internet Architecture

The TCP/IP protocol suite has evolved over a time period of some 25 years. We will describe the most important aspects of the protocol suite in this and the following chapters.

3.2.4 Layered Protocols

TCP/IP, like most networking software, is modeled in layers. This layered representation leads to the term protocol stack which is synonymous with protocol suite. It can be used for situating (but not for comparing functionally) the TCP/IP protocol suite against others, such as SNA and Open System Interconnection (OSI). Functional comparisons cannot easily be extracted from this, as there are basic differences in the layered models used by the different protocol suites.

The Internet protocols are modeled in four layers as shown in architecture model below:

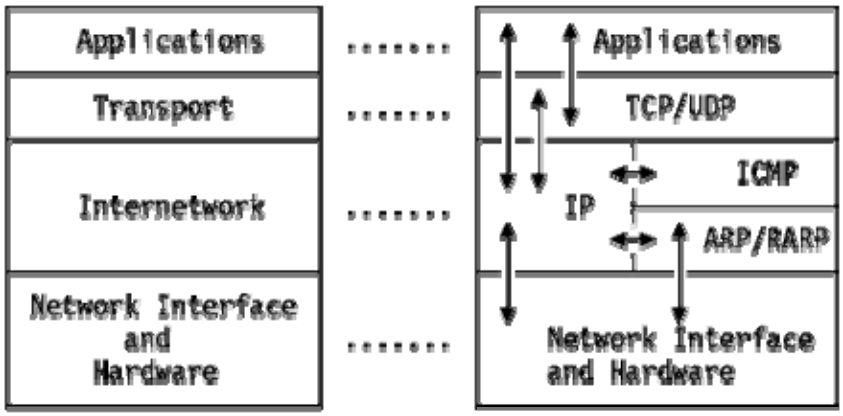

Figure 2

3.2.5 Application

Application is a user process cooperating with another process on the same or a different host. Examples are TELNET (a protocol for remote terminal connections), FTP (File Transfer Protocol) and SMTP (Simple Mail Transfer Protocol).

Transport provides the end-to-end data transfer. Example protocols are TCP (connection-oriented) and UDP (connectionless). Both are discussed in detail in Transmission Control Protocol (TCP) and User Datagram Protocol (UDP)

Inter-network also called the internet layer or the network layer, the inter-network provides the “virtual network” image of internet (that is, this layer shields the higher levels from the typical network architecture below it). Internet Protocol (IP) is the most important protocol in this layer. It is a connectionless protocol which doesn't assume reliability from the lower layers. IP does not provide reliability, flow control or error recovery. These functions must be provided at a higher level, either at the Transport layer by using TCP as the transport protocol, or at the Application layer if UDP is used as the transport protocol. IP is discussed in detail in Internet Protocol (IP). A message unit in an IP network is called an IP datagram. This is the basic unit of information transmitted across TCP/IP networks. It is described in IP Datagram but

we shall refer to it in this section to show how the different TCP/IP layers relate to an internet.

Network Interface also called the link layer or the data-link layer, the network interface layer is the interface to the actual network hardware. This interface may or may not provide reliable delivery, and may be packet or stream oriented. In fact, TCP/IP does not specify any protocol here, but can use almost any network interface available, which illustrates the flexibility of the IP layer. Examples are IEEE 802.2, X.25 (which is reliable in itself), ATM, FDDI, Packet Radio Networks (such as the AlohaNet) and even SNA.

The actual interactions between the layers are shown by the arrows in Figure 2. A more detailed "layering model" is shown in Figure 3.

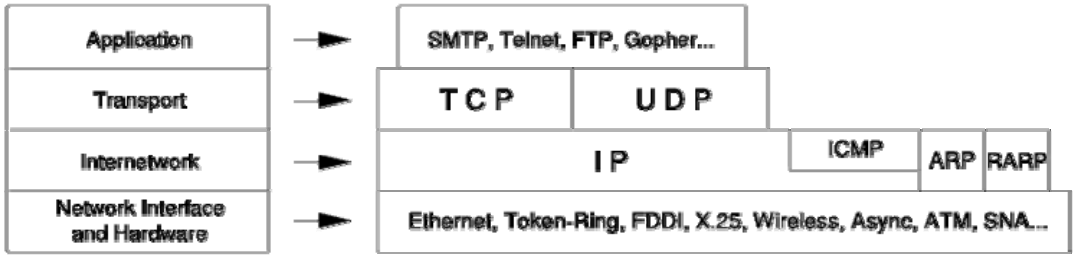

Figure 3

3.2.6 Bridges, Routers and Gateways

Forming an inter-network by interconnecting multiple networks is done by routers. It is important to distinguish between a router, a bridge and a gateway.

Bridge

It interconnects LAN segments at the Network Interface layer level and forwards frames between them. A bridge performs the function of a MAC relay, and is independent of any higher layer protocol (including the Logical Link protocol). It provides MAC layer protocol conversion, if required. Examples of bridges are:

A PS/2 running the IBM Token-Ring Network Bridge program

The IBM 8229 LAN bridge

A bridge can be said to be transparent to IP. That is, when a host sends an IP datagram to another host on a network connected by a bridge, it sends the datagram directly to the host and the datagram “crosses” the bridge without the sending host being aware of it.

Router

Router interconnects networks at the inter-network layer level and routes packets between them. The router must understand the addressing structure associated with the networking protocols it supports and take decisions on whether, or how, to forward packets. Routers are able to select the best transmission paths and optimal packet sizes. The basic routing function is implemented in the IP layer of the TCP/IP protocol stack. Therefore any host or workstation running TCP/IP may be used as a router. However, dedicated routers such as the IBM 6611 Network Processor provide much more sophisticated routing than the minimum function implemented by IP. Because IP provides this basic routing function, the term “IP router”, is often used. Other, older, terms for router are “IP gateway”, “Internet gateway” and “gateway”. The term gateway is now normally used for connections at a higher level than the router level.

A router can be said to be visible to IP. That is, when a host sends an IP datagram to another host on a network connected by a router, it sends the datagram to the router and not directly to the target host.

Gateway

Gateway interconnects networks at higher levels than bridges or routers. A gateway usually supports address mapping from one network to another, and may also provide transformation of the data between the environments to support end-to-end application connectivity. Gateways typically limit the interconnectivity of two networks to a subset of the application protocols supported on either one. For example, a VM host running TCP/IP may be used as an SMTP/RSCS mail gateway.

Note: The term "gateway", when used in this sense, is not synonymous with "IP gateway".

A gateway can be said to be opaque to IP. That is, a host cannot send an IP datagram through a gateway: it can only send it to a gateway. The higher-level protocol information carried by the datagrams is then passed on by the gateway using whatever networking architecture is used on the other side of the gateway.

Closely related to routers and gateways is the concept of a firewall or firewall gateway which is used to restrict access from the Internet to a network or a group of networks controlled by an organization for security reasons.

3.2.7 IP Routing

Incoming datagrams will be checked to see if the local host is the IP destination host:

If answer is yes, then datagram is passed to the higher-level protocols. Otherwise the datagram is for a different host. The action depends on the value of the ipforwarding flag.

If value of this flag is true, the datagram is treated as an outgoing datagram and is routed to the next hop according to the algorithm described below. In case value of this flag is false, the datagram is discarded.

In the internet protocol, outgoing IP datagrams pass through the IP routing algorithm which determines where to send the datagram according to the destination IP address.

- If the host has an entry in its IP routing table, which matches the destination IP address, the datagram is sent to the address in the entry.
- If the network number of the destination IP address is the same as the network number for one of the host's network adapters (that is, the destination and the host are on the same network) the datagram is sent to the physical address of the host matching the destination IP address.
- Otherwise, the datagram is sent to a default router. This base algorithm, needed on all IP implementations, is sufficient to perform the base routing function.

As noted above, a TCP/IP host has basic router functionality included in the IP protocol. Such a router is adequate for simple routing, but not for complex networks.

The IP routing mechanism combined with the “layered” view of the TCP/IP protocol stack, is represented in Figure 4. This shows an IP datagram, going from one IP address (network number X, host number A) to another (network number Y, host number B), through two physical networks. Note that at the intermediate router, only the lower part of the TCP/IP protocol stack (the inter-network and the network interface layers) are involved.

3.2.8 Addressing

Internet addresses can be symbolic or numeric. The symbolic form is easier to read, for example: myname@ibm.com. The numeric form is a 32-bit unsigned binary value which is usually expressed in a dotted decimal format. For example, 9.167.5.8 is a valid Internet address. The numeric form is used by the IP software. The mapping between the two is done by the Domain Name System discussed in Domain Name System (DNS). We shall first look at the numeric form, which is called the IP address.

3.3 IP Addresses

The standards for IP addresses are described in RFC 1166 -- Internet Numbers.

To be able to identify a host on the internet, each host is assigned an address, the IP address, or Internet Address. When the host is attached to more than one network, it is called multi-homed and it has one IP address for each network interface. The IP address consists of a pair of numbers:

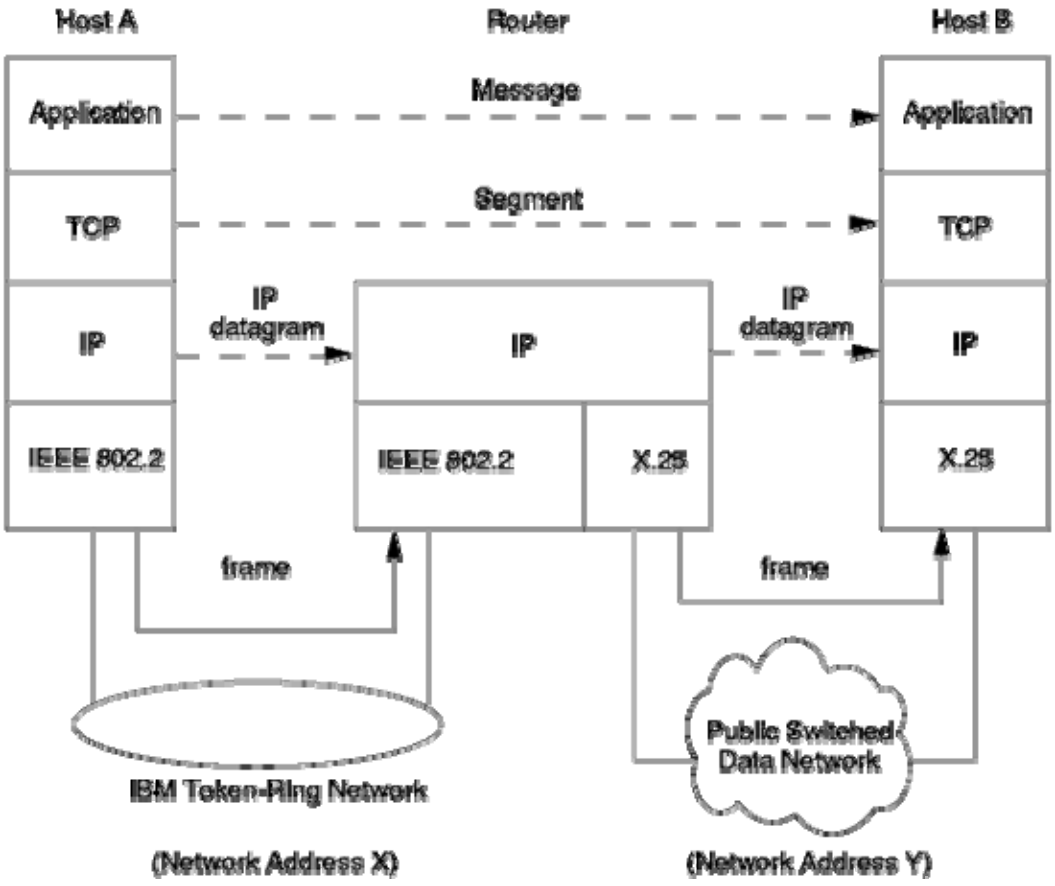

Figure 4 (IP Routing Mechanism)

IP address = <network number><host number>

The network number part of the IP address is centrally administered by the Internet Network Information Center (the InterNIC) and is unique throughout the Internet. (1)
IP addresses are 32-bit numbers usually represented in a dotted decimal form (as the decimal representation of four 8-bit values concatenated with dots). For example 128.2.7.9 is an IP address with 128.2 being the network number and 7.9 being the host number. The rules used to divide an IP address into its network and host parts are explained below.
The binary format of the IP address 128.2.7.9 is:

10000000 00000010 00000111 00001001

IP addresses are used by the IP protocol (see Internet Protocol (IP)) to uniquely identify a host on the internet. IP datagrams (the basic data packets exchanged between hosts) are transmitted by some physical network attached to the host and each IP datagram contains a source IP address and a destination IP address. To send a datagram to a certain IP destination, the target IP address must be translated or mapped to a physical address. This may require transmissions on the network to find out the destination's physical network address (for example, on LANs the Address Resolution Protocol, is used to translate IP addresses to physical MAC addresses).
The first bits of the IP address specify how the rest of the address should be separated into its network and host part.
The terms network address and netID are sometimes used instead of network number, but the formal term, used in RFC 1166, is network number. Similarly, the terms host address and hostID are sometimes used instead of host number.

There are five classes of IP addresses. These are shown in Figure 5

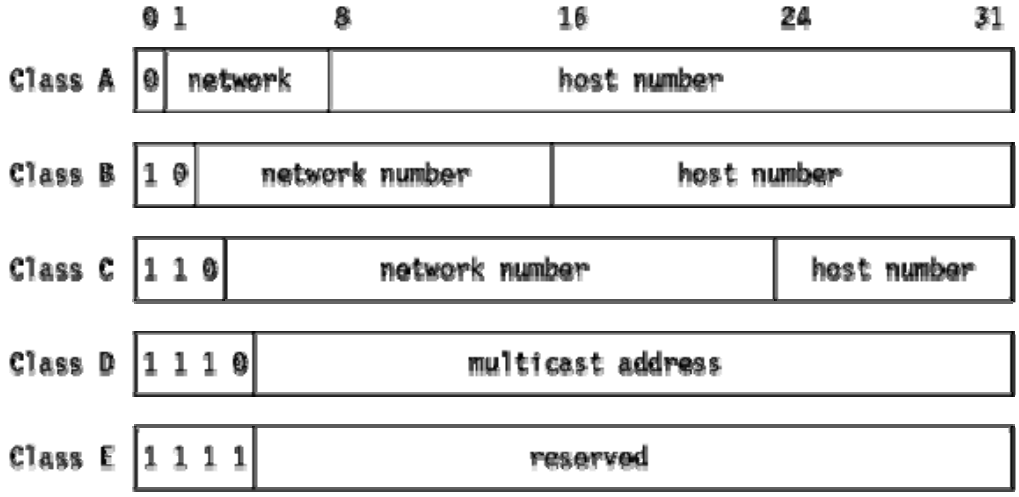

Figure 5

Note: Two numbers out of each of the class A, class B and class C network numbers, and two host numbers out of every network are pre-assigned: the “all bits 0” number and the “all bits 1” number. These are discussed below in Special IP Addresses.

Class A addresses use 7 bits for the network number giving 126 possible networks (we shall see below that out of every group of network and host numbers, two have a special meaning). The remaining 24 bits are used for the host number, so each networks can have up to 2(superscript 24)-2 2 to the power 24 minus 2 (16,777,214) hosts.

Class B addresses use 14 bits for the network number, and 16 bits for the host number giving 16382 networks each with a maximum of 65534 hosts.

Class C addresses use 21 bits for the network number and 8 for the host number giving 2,097,150 networks each with up to 254 hosts.

Class D addresses are reserved for multicasting, which is used to address groups of hosts in a limited area. See Multicasting for more information on multicasting.

Class E addresses are reserved for future use.

It is clear that a class A address will only be assigned to networks with a huge number of hosts, and that class C addresses are suitable for networks with a small number of hosts. However, this means that medium-sized networks (those with more than 254 hosts or where there is an expectation that there may be more than 254 hosts in the future) must use Class B addresses. The number of small- to medium-sized networks has been growing very rapidly in the last few years and it was feared that, if this growth had been allowed to continue unabated, all of the available Class B network addresses would have been used up till. This is termed the IP Address Exhaustion problem.

One point to note about the split of an IP address into two parts is that this split also splits the responsibility for selecting the IP address into two parts. The network number is assigned by the InterNIC, and the host number by the authority which controls the network. As we shall see in the next section, the host number can be further subdivided: this division is controlled by the authority which owns the network, and not by the InterNIC.

3.3.1 Subnets

Due to the explosive growth of the Internet, the use of assigned IP addresses became too inflexible to allow easy changes to local network configurations. These changes might occur when:

- A new physical network is installed at a location.
- Growth of the number of hosts requires splitting the local network into two or more separate networks.

To avoid having to request additional IP network addresses in these cases, the concept of subnets was introduced.

The host number part of the IP address is sub-divided again into a network number and a host number. This second network is termed a sub-network or subnet. The main network now consists of a number of subnets and the IP address is interpreted as:
<network number><subnet number><host number>
The combination of the subnet number and the host number is often termed the “local address” or the “local part”. “Sub-netting” is implemented in a way that is transparent to remote networks. A host within a network which has subnets is aware of the sub-netting but a host in a different network is not; it still regards the local part of the IP address as a host number.
The division of the local part of the IP addresses into subnet number and host number parts can be chosen freely by the local administrator; any bits in the local part can be used to form the subnet accomplished. The division is done using a subnet mask which is a 32 bit number. Zero bits in the subnet mask indicate bit positions ascribed to the host number, and ones indicate bit positions ascribed to the subnet number. The bit positions in the subnet mask belonging to the network number are set to ones but are not used. Subnet masks are usually written in dotted decimal form, like IP addresses.
The special treatment of “all bits zero” and “all bits one” applies to each of the three parts of a sub-netted IP address just as it does to both parts of an IP address which has not been sub-netted. See Special IP Addresses. For example, a sub-netted Class B network, which has a 16-bit local part, could use one of the following schemes:
The first byte is the subnet number, the second the host number. This gives us 254 (256 minus 2 with the values 0 and 255 being reserved) possible subnets, each having up to 254 hosts. The subnet mask is 255.255.255.0.

The first 12 bits 15 are used for the subnet number and the last four for the host number. This gives us 4094 possible subnets (4096 minus 2) but only 14 hosts per

subnet (16 minus 2). The subnet mask is 255.255.255.240. There are many other possibilities.

While the administrator is completely free to assign the subnet part of the local address in any legal fashion, the objective is to assign a number of bits to the subnet number and the remainder to the local address. Therefore, it is normal to use a contiguous block of bits at the beginning of the local address part for the subnet number because this makes the addresses more readable (this is particularly true when the subnet occupies 8 or 16 bits). With this approach, either of the subnet masks above are “good” masks, but masks like

255.255.252.252 and 255.255.255.15 are not.

3.3.2 Types of Subnetting

There are two types of subnetting: static and variable length. Variable length is the more flexible of the two. Which type of subnetting is available depends upon the routing protocol being used; native IP routing supports only static subnetting, as does the widely used RIP protocol. However, RIP Version 2 supports variable length subnetting as well.

3.3.3 Static Subnetting

Static subnetting means that all subnets in the subnetted network use the same subnet mask. This is simple to implement and easy to maintain, but it implies wasted address space for small networks. For example, a network of four hosts that uses a subnet mask of 255.255.255.0 wastes 250 IP addresses. It also makes the network more difficult to reorganize with a new subnet mask. Currently, almost every host and router supports static subnetting.

3.3.4 Variable Length Subnetting

When variable length subnetting is used, the subnets that make up the network may use different subnet masks. A small subnet with only a few hosts needs a subnet mask that accommodates only these few hosts. A subnet with many hosts attached may need a different subnet mask to accommodate the large number of hosts. The possibility to assign subnet masks according to the needs of the individual subnets will help conserve network addresses. Also, a subnet can be split into two parts by adding another bit to the subnet mask. Other subnets in the network are unaffected by the change. Not every host and router supports variable length subnetting.

Only networks of the size needed will be allocated and routing problems will be solved by isolating networks with routers that support variable subnetting. A host that does not support this kind of subnetting would have to route to a router that supports variable subnetting.

3.3.5 Mixing Static and Variable Length Subnetting

At first sight, it appears that the presence of a host which only supports static subnetting would prevent variable length subnetting from being used anywhere in the network. Fortunately this is not the case. Provided that the routers between subnets with different subnet masks are using variable length subnetting, the routing protocols employed are able to hide the difference between subnet masks from the hosts in a subnet. Hosts can continue to use basic IP routing and offload all of the complexities of the subnetting to dedicated routers.

3.3.6 Special IP Addresses

As noted above, any component of an IP address with a value “all bits 0” or all “all bits 1” has a special meaning.

All bits 0 stands for “this”: “this” host (IP address with <host number>=0) or “this” network (IP address with <network number>=0) and is only used when the real value is not known. This form is only used in source addresses when the host is trying to determine its IP addresses from a remote server. The host may know include its host number if known, but not its subnet or network number.

All bits 1 stands for “all”: “all” networks or “all” hosts. For example, 128.2.255.255 (A class B address with a host number of 255.255) means all hosts on network 128.2. These are used in broadcast messages, as described below.

There is another address of special importance: the “all bits 1” class A network number 127 is reserved for the loop-back address. Anything sent to an address with 127 as the value of the high order byte, for example 127.0.0.1, must not be routed via a network but must be routed directly from the IP implementation's output driver to its input driver.

Internet Protocol (IP)

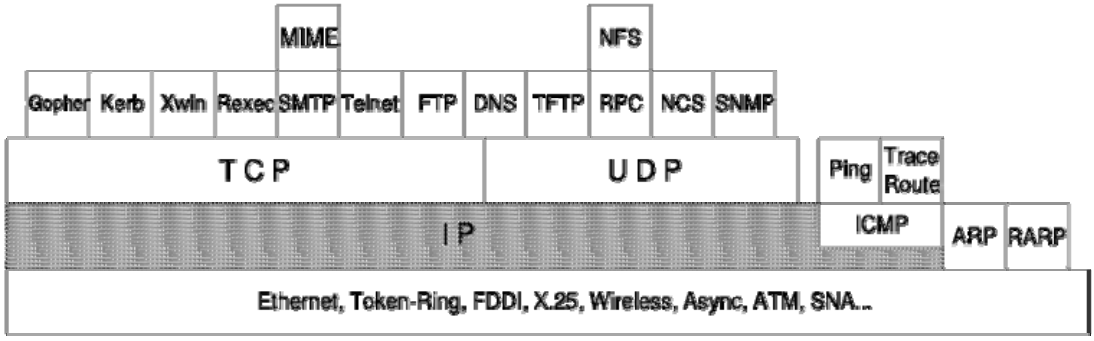

Figure 6 (IP Protocol)

IP is a standard protocol with STD number 5 which also includes ICMP (Internet Control Message Protocol) and IGMP (Internet Group Management Protocol). Its status is required.

The current IP specification can be found in RFCs 791, 950, 919 and 922, with updates in RFC 1349.

IP is the protocol that hides the underlying physical network by creating a virtual network view. It is an unreliable, best-effort connectionless packet delivery protocol.

It adds no reliability, flow control or error recovery to the underlying network interface protocol. Packets (datagrams) sent by IP may be lost, out of order, or even duplicated, and IP will not handle these situations. It is up to higher layers to provide these facilities.

IP also assumes little from the underlying network mechanisms, only that the datagrams will “probably” (best-effort) be transported to the addressed host.

3.4 IP Datagrams and Datagram Forwarding

The Internet datagram (IP datagram) is the base transfer packet in the Internet protocol suite. It has a header containing information for IP, and data that is relevant only to the higher level protocols. Figure 7 shows the base IP datagram.

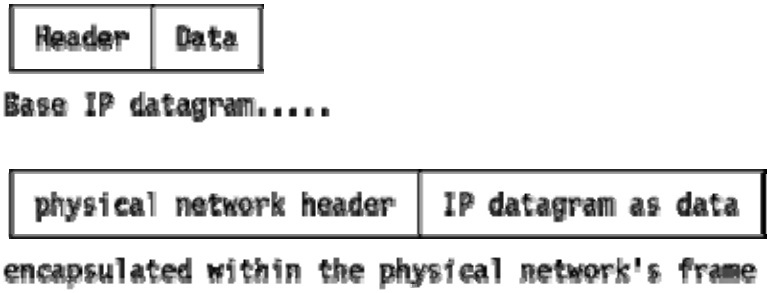

Figure 7

The IP datagram is encapsulated in the underlying network's frame, which usually has a maximum length or frame limitation, depending on the hardware used. For Ethernet, this will typically be 1500 bytes. Instead of limiting the IP datagram length to some maximum size, IP can deal with fragmentation and re-assembly of its datagrams. In

particular, the IP standard does not impose a maximum size, but states that all sub-networks should be able to handle datagrams of at least 576 bytes.

Fragments of a datagram all have a header, basically copied from the original datagram, and data following it. They are treated as normal IP datagrams while being transported to their destination. Note, however, that if one of the fragments gets lost, the complete datagram is considered lost since IP does not provide any acknowledgment mechanism, so the remaining fragments will simply be discarded by the destination host.

IP Datagram Format

Figure 8 shows the format of IP Datagram. It is a minimum of 20 bytes long:

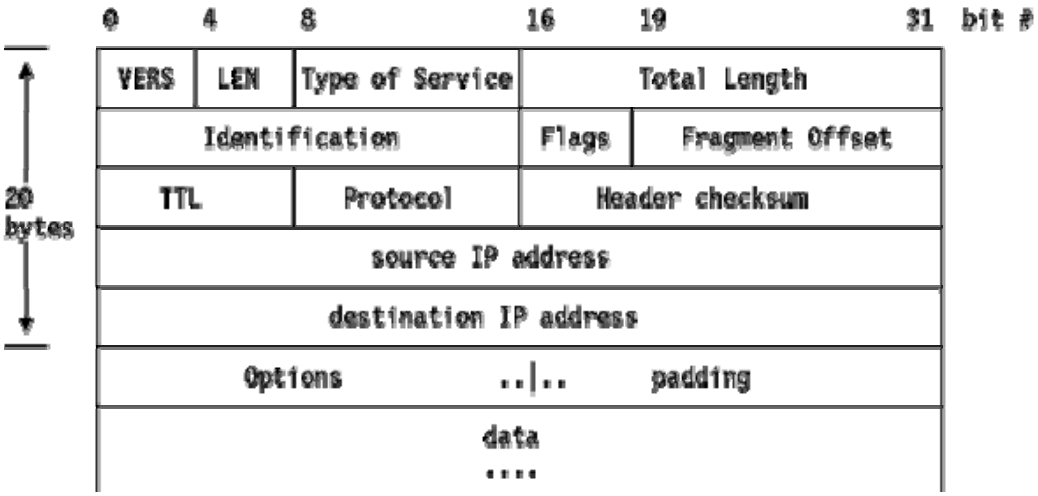

Figure 8 (Format of IP Datagram)

In the figure 8 :

VERS -The version of the IP protocol. Current is 4 and future is 6.

LEN - The length of the IP header counted in 32-bit quantities. This does not include the data field.

Type of Service - The type of service is an indication of the quality of service requested for this IP datagram.

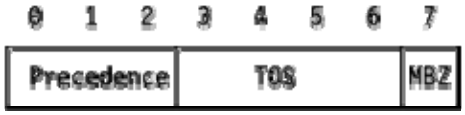

Where:

Precedence is a measure of the nature and priority of this datagram:

000 - Routine

001 - Priority

010 - Immediate

011 - Flash

100 - Flash override

101 - Critical

110 – Inter-network control

111 - Network control

TOS Specifies the type of service value:

1000 - Minimize delay

0100 - Maximize throughput

0010 - Maximize reliability

0001 - Minimize monetary cost

0000 - Normal service

MBZ is reserved for future use ("must be zero" unless participating in an Internet protocol experiment which makes use of this bit).A detailed description of the type of service can be found in the RFC 1349.

Total Length - The total length of the datagram, header and data, specified in bytes.

Identification - A unique number assigned by the sender to aid in reassembling a fragmented datagram. Fragments of a datagram will have the same identification number.

Flags - Various control flags:

Where first cell is reserved, must be zero

DF - Don't Fragment: 0 means allow fragmentation, 1 means do not allow fragmentation.

MF - More Fragments: 0 means that this is the last fragment of this datagram, 1 means that this is not the last fragment.

Fragment Offset - Used with fragmented datagrams, to aid in reassembly of the full datagram. The value is the number of 64-bit pieces (header bytes are not counted) that are contained in earlier fragments. In the first (or only) fragment, this value is always zero.

Time to Live - Specifies the time (in seconds) this datagram is allowed to travel. Each router where this datagram passes is supposed to subtract from this field its processing time for this datagram. Actually a router is able to process a datagram in less than 1 second; thus it will subtract one from this field, and the TTL becomes a hop-count metric rather than a time metric. When the value reaches zero, it is assumed that this datagram has been traveling in a closed loop and it is discarded. The initial value should be set by the higher-level protocol which creates the datagram.

Protocol - Indicates the higher-level protocol to which IP should deliver the data in this datagram. Some important values are:

0 - Reserved

1- Internet Control Message Protocol (ICMP)

2 - Internet Group Management Protocol (IGMP)

3 - Gateway-to-Gateway Protocol (GGP)

4 - IP (IP encapsulation)

5 - Stream

6 - Transmission Control (TCP)

8 -Exterior Gateway Protocol (EGP)

9 - Private Interior Routing Protocol

17 - User Datagram (UDP)

89 -Open Shortest Path First

The full list can be found in STD 2 - Assigned Internet Numbers.

Header Checksum - Is a checksum on the header only. It does not include the data. The checksum is calculated as the 16-bit one's complement of the one's complement sum of all 16-bit words in the header. For the purpose of this calculation, the checksum field is assumed to be zero. If the header checksum does not match the contents, the datagram is discarded because at least one bit in the header is corrupt, and the datagram may even have arrived at the wrong destination.

Source IP Address - The 32-bit IP address of the host sending this datagram.

Destination IP Address - The 32-bit IP address of the destination host for this datagram.

Options - Variable length. An IP implementation is not required to be capable of generating options in the datagrams it creates, but all IP implementations are required to be able to process datagrams containing options. The Options field is variable in length. There may be zero or more options. There are two option formats. The format for each is dependent on the value of the option number found in the first byte.

Length - counts the length (in bytes) of the option, including the type and length fields.

Option data - contains data relevant to the option.

Padding - If an option is used, the datagram is padded with all-zero bytes up to the next 32-bit boundary.

Data - The data contained in the datagram is passed to a higher-level protocol, as specified in the protocol field.

3.5 IP Encapsulation

One of the most important concepts in inter-protocol operation is that of encapsulation. Most data originates within the higher layers of the OSI model. The protocols at these layers pass the data down to lower layers for transmission, usually in the form of discrete messages. Upon receipt, each lower-level protocol takes the entire contents of the message received and encapsulates it into its own message format, adding a header and possibly a footer that contain important control information. Encapsulation is explained in general terms in a separate topic.

A good analogy for how encapsulation works is a comparison to sending a letter enclosed in an envelope. You might write a letter and put it in a white envelope with a name and address, but if you gave it to a courier for overnight delivery, they would take that envelope and put it in a larger delivery envelope.

Due to the prominence of TCP/IP, the Internet Protocol is one of the most important places where data encapsulation occurs on a modern network. Data is passed to IP typically from one of the two main transport layer protocols: TCP or UDP. This data is already in the form of a TCP or UDP message with TCP or UDP headers. This is then encapsulated into the body of an IP message, usually called an IP datagram or IP packet. Encapsulation and formatting of an IP datagram is also sometimes called packaging—again, the implied comparison to an envelope is obvious.

3.6 Fragmentation

When an IP datagram travels from one host to another, it can cross different physical networks. Physical networks have a maximum frame size, called the Maximum

Transmission Unit (MTU), which limits the length of a datagram that can be placed in one physical frame. Therefore, a scheme has been put in place to fragment long IP datagrams into smaller ones, and to reassemble them at the destination host. IP requires that each link has an MTU of at least 68 bytes, so if any network provides a lower value than this, fragmentation and re-assembly must be implemented in the network interface layer in a way that is transparent to IP. 68 is the sum of the maximum IP header length of 60 bytes and the minimum possible length of data in a non-final fragment (8 bytes). IP implementations are not required to handle unfragmented datagrams larger than 576 bytes, but most implementations will handle larger values, typically slightly more than 8192 bytes or higher, and rarely less than 1500.

An unfragmented datagram has all-zero fragmentation information. That is, the more fragments flag bit is zero and the fragment offset is zero. When fragmentation is to be done, the following steps are performed:

- The DF flag bit is checked to see if fragmentation is allowed. If the bit is set, the datagram will be discarded and an error will be returned to the originator using ICMP.
- Based on the MTU value, the data field is split into two or more parts. All newly created data portions must have a length which is a multiple of 8 bytes, with the exception of the last data portion.
- All data portions are placed in IP datagrams. The header of these datagrams are copies of the original one, with some modifications:
    - The more fragments flag bit is set in all fragments except the last.
    - The fragment offset field in each is set to the location this data portion occupied in the original datagram, relative to the beginning of the original unfragmented datagram. The offset is measured in 8-byte units.

- If options were included in the original datagram, the high order bit of the option type byte determines whether or not they will be copied to all fragment datagrams or just to the first one. For instance, source route options have to be copied in all fragments and therefore they have this bit set.
- The header length field is of the new datagram is set.
- The total length field of the new datagram is set.
- The header checksum field is re-calculated.

- Each of these fragmented datagrams is now forwarded as a normal IP datagram. IP handles each fragment independently, that is, the fragments may traverse different routers to the intended destination, and they may be subject to further fragmentation if they pass through networks that have smaller MTUs.

At the destination host, the data has to be reassembled into one datagram. The identification field of the datagram was set by the sending host to a unique number (for the source host, within the limits imposed by the use of a 16-bit number). As fragmentation doesn't alter this field, incoming fragments at the receiving side can be identified, if this ID field is used together with the Source and Destination IP addresses in the datagram. The Protocol field is also be checked for this identification. In order to reassemble the fragments, the receiving host allocates a buffer in storage as soon as the first fragment arrives. A timer routine is then started. When the timer timeouts and not all of the fragments have been received, the datagram is discarded. The initial value of this timer is called the IP datagram time-to-live (TTL) value. It is implementation dependent, and some implementations allow it to be configured; for example AIX Version 3.2 provides an ipfragttl option with a default value of 60 seconds.

When subsequent fragments of the datagram arrive, before the timer expires, the data is simply copied into the buffer storage, at the location indicated by the fragment offset field. As soon as all fragments have arrived, the complete original unfragmented datagram is restored, and processing continues, just as for unfragmented datagrams.

Note: IP does not provide the reassembly timer. It will treat each and every datagram, fragmented or not, the same way, that is, as individual messages. It is up to the higher layer to implement a timeout and to look after any missing fragments. The higher layer could be TCP for a connection-oriented transport network or the application for connectionless transport networks based upon UDP and IP.

The netstat command may be used on some TCP/IP hosts to list details of fragmentation that is occurring. An example of this is the netstat -i command in TCP/IP for OS/2.

3.7 The Future IP (IPv6)

One of the newest major standards on the horizon is IPv6. Although IPv6 has not officially become a standard, it is worth some overview. It is very possible that this information will change as we move closer to IPv6 as a standard, so you should use this as a guide into IPv6, not the definitive information.

A number of books are now being published that cover in detail this emerging standard. All the RFCs available on the Internet have the raw details on how this standard is developing. However, these documents are difficult to interpret at first glance and require some commitment to going through any number of RFCs pertaining to many subjects all related to IPv6 development.

Internet Protocol Version 4 is the most popular protocol in use today, although there are some questions about its capability to serve the Internet community much longer. IPv4 was finished in the 1970s and has started to show its age. The main issue

surrounding IPv6 is addressing—or, the lack of addressing—because many experts believe that we are nearly out of the four billion addresses available in IPv4. Although this seems like a very large number of addresses, multiple large blocks are given to government agencies and large organizations. IPv6 could be the solution to many problems, but it is still not fully developed and is not a standard—yet.

Many of the finest developers and engineering minds have been working on IPv6 since the early 1990s. Hundreds of RFCs have been written and have detailed some major areas, including expanded addressing, simplified header format, flow labeling, authentication, and privacy.

Expanded addressing moves us from 32-bit address to a 128-bit addressing method. It also provides newer unicast and broadcasting methods, injects hexadecimal into the IP address, and moves from using "." to using ":" as delimiters. Figure 9 shows the IPv6 packet header format.

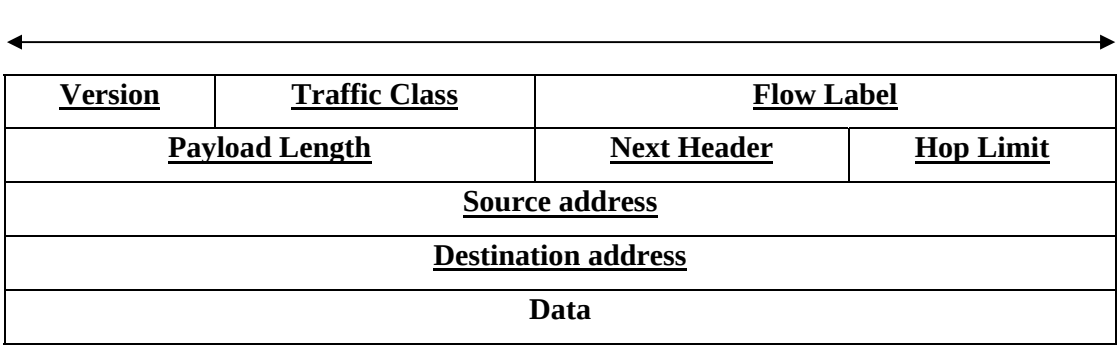

Figure 9 (IPv6 Packet Header Format)

The simplified header is 40 bits long and the format consists of Version, Class, Flow Label, Payload Length, Next Header, Hop Limit, Source Address, Destination Address, Data, and Payload fields.

Version (4 bits) - IPv6 version number.

Traffic Class (8 bits.) - Internet traffic priority delivery value.

Flow Label (20 bits) - Used for specifying special router handling from source to destination(s) for a sequence of packets.

Payload Length (16 bits unsigned) - Specifies the length of the data in the packet. When cleared to zero, the option is a hop-by-hop Jumbo payload.

Next Header (8 bits) - Specifies the next encapsulated protocol. The values are compatible with those specified for the IPv4 protocol field.

Hop Limit (8 bits unsigned) - For each router that forwards the packet, the hop limit is decremented by 1. When the hop limit field reaches zero, the packet is discarded. This replaces the TTL field in the IPv4 header that was originally intended to be used as a time based hop limit.

Source address (16 bytes) - The IPv6 address of the sending node.

Destination address (16 bytes) - The IPv6 address of the destination node.

Some of the benefits of IPv6 seem obvious: greater addressing space, QoS(Quality of Service), and better routing performance and services. However, a number of barriers must be overcome before the implementation of IPv6. The biggest question for most of us will be what the business need is for moving from current IPv4 to IPv6. The killer app has not appeared yet, but it may be closer than we think. The second consideration is the cost—it may not have much to do with hardware replacement cost. All the larger routers have upgradable OSs IOS; the only necessity is the commitment to upgrading IOS. More likely to do with training and support of minor IP devices such as printers and network faxes, they will support the new address space. IPv6 has schemes to support old and new, however, so this may not even be a barrier. The last issue to consider is training: This will need to happen sooner or later because we all need to start thinking about 128-bit addressing based on MAC addresses in HEX. This involves all new ways of addressing and will be an uncomfortable change for many people.

This conclusion may seem negative, but the greater good will overpower all the up-front issues. The issue is not whether you will have to move to IPv6, but when! We all

need IPv6; the increased address space is needed for the growth of IP appliances. IP-ready cars are already shipping today. This requires mobility, which is addressed in IPv6.

3.8 TCP Reliable Transport Service

What was the most impressive aspect of the ENIAC computer? It was reliable in spite of the fact that its components (vacuum tubes) were unreliable. Internet pioneers faced a similar challenge. To provide a reliable transport service over unreliable (conceptually) technology. IP datagram service is unreliable in that datagrams may be Lost (usually dropped due to congestion), duplicated, and delivered out of order. TCP provides reliable service nevertheless.

The Service TCP Provides to Applications

- TCP is connection-oriented.
- Each TCP connection has exactly two endpoints.
- TCP is completely reliable.
- TCP allows full duplex communication.
- TCP provides a stream interface.
- TCP uses a reliable connection startup and a graceful shutdown.

3.8.1 TCP Connections

TCP connections are either virtual or an ordered pair of endpoints. An endpoint is an ordered pair containing an IP address and a TCP port number. This information is kept by the OS. Use netstat command on window OS to view this. Achieving reliability is a major challenge due to following reasons-

- One computer reboots during a file transfer

- It is rebooted and a second connection is established
- Messages from the two sessions get mixed up
- Packets from the first session must be rejected

3.8.2 The Window Principle

A simple transport protocol might use the following principle: send a packet and then wait for an acknowledgment from the receiver before sending the next packet. If the ACK is not received within a certain amount of time, retransmit the packet. This principle is described in Figure 10.

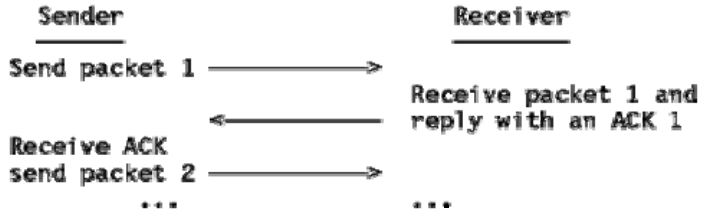

Figure 10 (Window Principle)

While this mechanism ensures reliability, it only uses a part of the available network bandwidth.

Consider now a protocol where the sender groups its packets to be transmitted as shown in Figure 11.

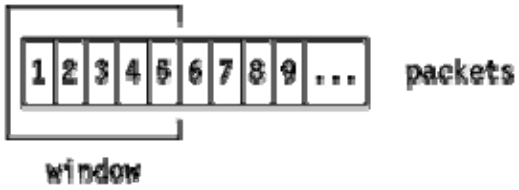

Figure 11 (Example of Window Principle)

It uses the following rules:

- The sender may send all packets within the window without receiving an ACK, but must start a timeout timer for each of them.
- The receiver must acknowledge each packet received, indicating the sequence number of the last well-received packet.
- The sender slides the window on each ACK received.

In our example, the sender may transmit packets 1 to 5 without waiting for any acknowledgment:

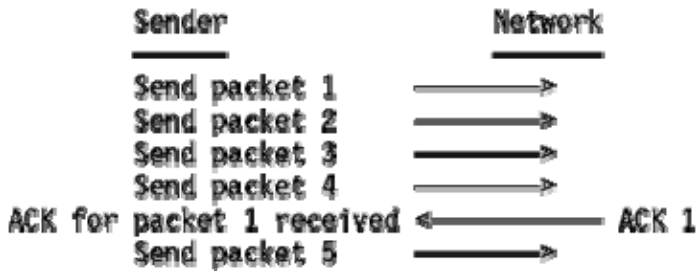

Figure 12 (Example of Rules between Sender and Receiver)

At the moment the sender receives the ACK 1 (acknowledgment for packet 1), it may slide its window to exclude packet 1:

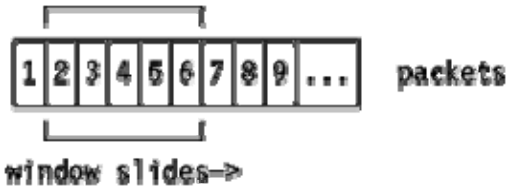

Figure 13 (Sliding Window)

At this point, the sender may also transmit packet 6.

Imagine some special cases:

- Packet 2 gets lost: the sender will not receive an ACK 2, so its window will remain in the position 1 (as last picture above). In fact, as the receiver did not receive packet 2, it will acknowledge packets 3, 4 and 5 with an ACK 1, since packet 1 was the last one received “in sequence”. At the sender's side, eventually a timeout will occur for packet 2 and it will be retransmitted. Note that reception of this packet by the receiver will generate an ACK 5, since it has now successfully received all packets 1 to 5, and the sender's window will slide four positions upon receiving this ACK 5.
- Packet 2 did arrive, but the acknowledgment gets lost: the sender does not receive ACK 2, but will receive ACK 3. ACK 3 is an acknowledgment for all packets up to 3 (including packet 2) and the sender may now slide his window to packet 4.

Conclusion:

This window mechanism ensures:

- Reliable transmission.
- Better use of the network bandwidth (better throughput).
- Flow-control, as the receiver may delay replying to a packet with an acknowledgment, knowing its free buffers available and the window-size of the communication.

3.8.3 The Window Principle Applied to TCP

The above window principle is used in TCP, but with a few differences:

- As TCP provides a byte-stream connection, sequence numbers are assigned to each byte in the stream. TCP divides this contiguous byte stream into TCP segments to transmit them. The window principle is used at the byte level; that is, the segments sent and ACKs received will carry byte-sequence numbers

and the window size is expressed as a number of bytes, rather than a number of packets.

- The window size is determined by the receiver, when the connection is established, and is variable during the data transfer. Each ACK message will include the window-size that the receiver is ready to deal with at that particular time.

The sender's data stream can now be seen as:

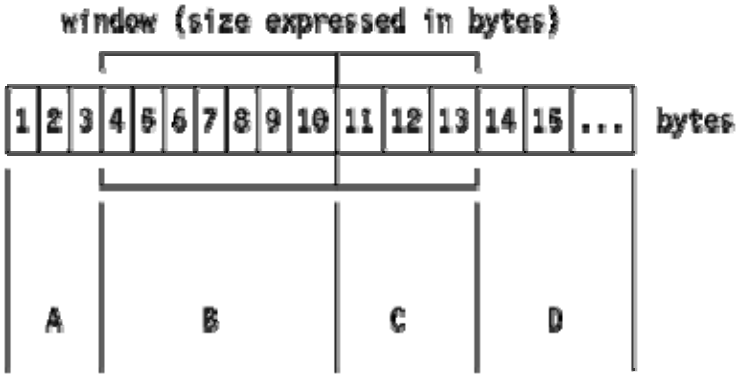

Figure 14 (Sender's Data Stream)

Where:

A - Bytes that are transmitted and have been acknowledged.

B - Bytes that are sent but not yet acknowledged.

C - Bytes that may be sent without waiting for any acknowledgment.

D - Bytes that may not yet be sent.

Remember that TCP will block bytes into segments, and a TCP segment only carries the sequence number of the first byte in the segment.

3.8.4 TCP Segment Format

TCP segment format is shown in figure 15. Where:

Source Port - The 16-bit source port number, used by the receiver to reply.

Destination Port - The 16-bit destination port number.

Sequence Number - The sequence number of the first data byte in this segment. If the SYN control bit is set, the sequence number is the initial sequence number (n) and the first data byte is n+1.

Acknowledgment Number - If the ACK control bit is set, this field contains the value of the next sequence number that the receiver is expecting to receive.

Data Offset - The number of 32-bit words in the TCP header. It indicates where the data begins.

Reserved - Six bits reserved for future use; must be zero.

URG - Indicates that the urgent pointer field is significant in this segment.

ACK - Indicates that the acknowledgment field is significant in this segment.

PSH - Push function.

RST - Resets the connection.

SYN - Synchronizes the sequence numbers.

FIN - No more data from sender.

Window - Used in ACK segments. It specifies the number of data bytes beginning with the one indicated in the acknowledgment number field which the receiver (= the sender of this segment) is willing to accept.

Checksum - The 16-bit one's complement of the one's complement sum of all 16-bit words in a pseudo-header, the TCP header and the TCP data. While computing the checksum, the checksum field itself is considered zero.

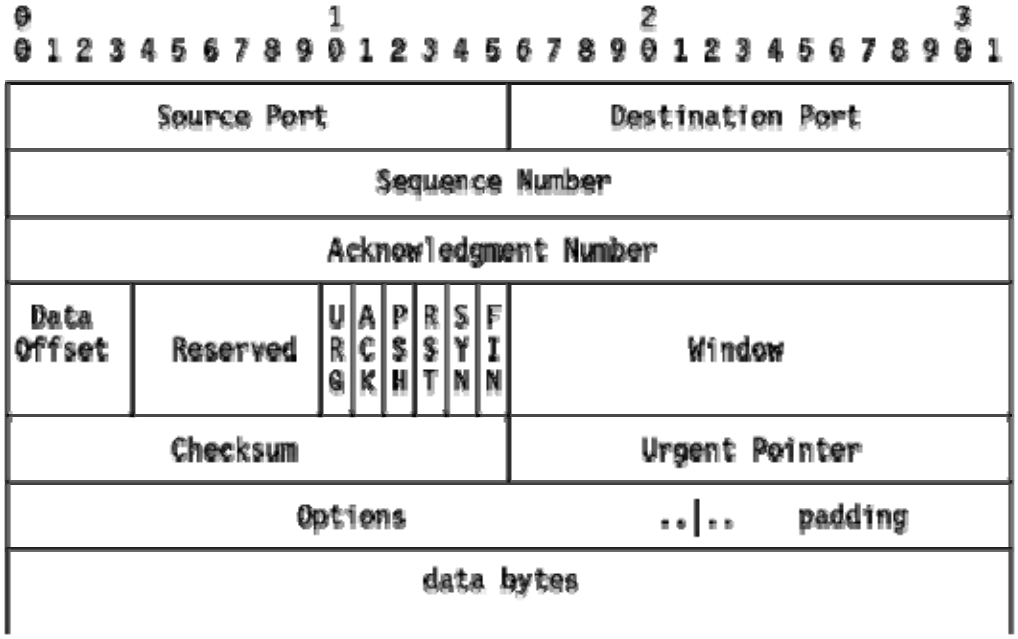

Figure 15 (TCP Segment Format)

The pseudo-header is the same as that used by UDP for calculating the checksum. It is a pseudo-IP-header, only used for the checksum calculation, with the format shown in Figure 16:

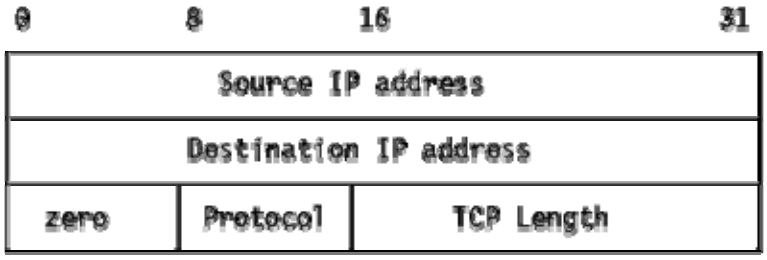

Figure 16 (Pseudo IP Header)

Urgent Pointer - Points to the first data octet following the urgent data. Only significant when the URG control bit is set.

Options - Just as in the case of IP datagram options, options can be either:

- A single byte containing the option number, or
- A variable length option in the following format:

Figure17 shows IP Datagram Option - Variable length option.

Figure 17 (IP Datagram Option)

There are currently only three options defined:

Kind	Length	Meaning
----	------	-------
0	-	End of option list.
1	-	No-Operation.
2	4	Maximum Segment Size.

Figure 18 shows Maximum Segment Size Option.

Figure 18 (Maximum Segment Size Option)

This option is only used during the establishment of the connection (SYN control bit set) and is sent from the side that is to receive data to indicate the maximum segment length it can handle. If this option is not used, any segment size is allowed.

Padding - All zero bytes used to fill up the TCP header to a total length that is a multiple of 32 bits.

3.8.5 Acknowledgments and Retransmissions

TCP sends data in variable length segments. Sequence numbers are based on a byte count. Acknowledgments specify the sequence number of the next byte that the receiver expects to receive.

Now suppose that a segment gets lost or corrupted. In this case, the receiver will acknowledge all further well-received segments with an acknowledgment referring to the first byte of the missing packet. The sender will stop transmitting when it has sent all the bytes in the window. Eventually, a timeout will occur and the missing segment will be retransmitted.

Suppose a window size of 1500 bytes, and segments of 500 bytes. Figure 19 describes the Acknowledgment and Retransmission Process.

A problem now arises, since the sender does know that segment 2 is lost or corrupted, but doesn't know anything about segments 3 and 4. The sender should at least retransmit segment 2, but it could also retransmit segments 3 and 4 (since they are within the current window). It is possible that:

Segment 3 has been received, and for segment 4 we don't know: it could be received, but ACK didn't reach us yet, or it could be lost also.

Segment 3 was lost, and we received the ACK 1500 upon the reception of segment 4.

Each TCP implementation is free to react to a timeout as the implementers wish. It could retransmit only segment 2, but in case 2 above, we will be waiting again until segment 3 times out. In this case, we lose all of the throughput advantages of the

window mechanism. Or TCP might immediately resend all of the segments in the current window.

Whatever the choice, maximal throughput is lost. This is because the ACK does not contain a second acknowledgment sequence number indicating the actual frame received.

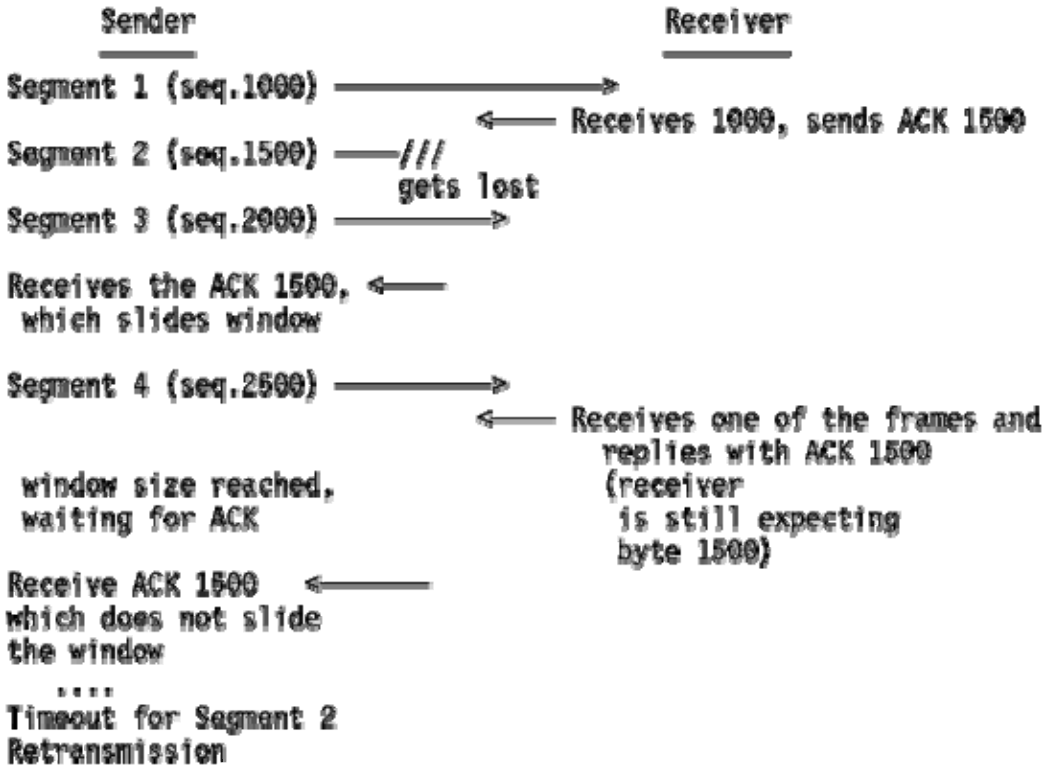

Figure 19(Example of Acknowledgement and Retransmission)

3.8.6 Variable timeout intervals

Each TCP should implement an algorithm to adapt the timeout values to be used for the round trip time of the segments. To do this, TCP records the time at which a segment was sent, and the time at which the ACK is received. A weighted average is

calculated over several of these round trip times, to be used as a timeout value for the next segment(s) to be sent.

This is an important feature, since delays may be variable on an internet, depending on multiple factors, such as the load of an intermediate low-speed network or the saturation of an intermediate IP gateway.

3.8.7 Establishing a TCP Connection

Before any data can be transferred, a connection has to be established between the two processes. One of the processes (usually the server) issues a passive OPEN call, the other an active OPEN call. The passive OPEN call remains dormant until another process tries to connect to it by an active OPEN.

On the network, three TCP segments are exchanged:

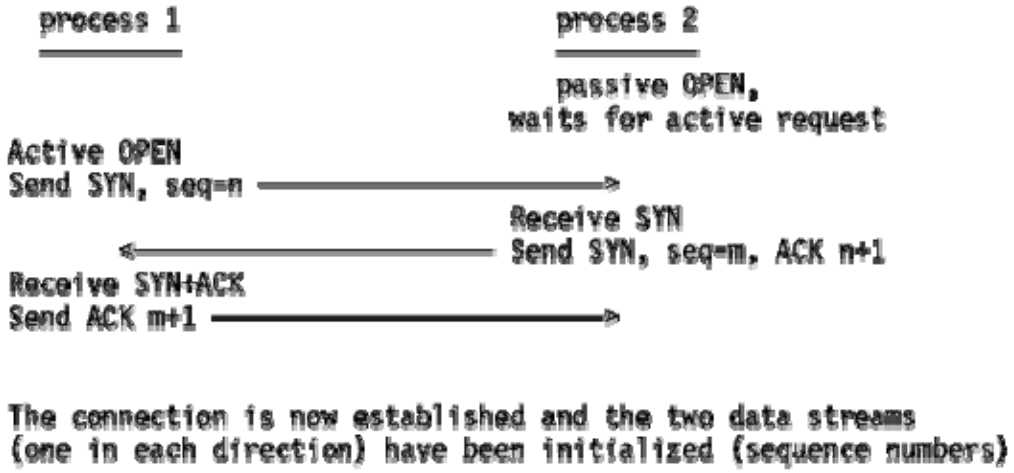

Figure 20 (Establishing a TCP Connection)

This whole process is known as **three-way handshake**. Note that the exchanged TCP segments include the initial sequence numbers from both sides, to be used on subsequent data transfers.

Closing the connection is done implicitly by sending a TCP segment with the FIN bit (no more data) set. As the connection is full-duplex (that is, we have two independent data streams, one in each direction), the FIN segment only closes the data transfer in one direction. The other process will now send the remaining data it still has to transmit and also ends with a TCP segment where the FIN bit is set. The connection is deleted (status information on both sides) once the data stream is closed in both directions.

3.8.8 TCP Segments Carried by IP Datagrams

TCP segments are transported in IP datagrams with the following parameter settings:

Type of Service = 00000000
that is: Precedence = routine
Delay = normal
Throughput = normal
Time to Live = 00111100 (1 minute)

3.9 Self Assessment Questions (SAQA)

Q1. Explain different type of internet connections.

Q2. Write short notes on-

i) Internetworking ii) Layered protocols iii) IP Encapsulation

iv) Fragmentation

Q3. Differentiate Bridges, Routers and Gateways

Q4. Explain IP addressing in detail

Q5. Explain the format of IP Header.

Q6. Explain different type of subnetting.

Q7. What is the future of IP?

Q8. Explain TCP segment format in detail.

Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Rishipal Singh**
Lesson Number: 04	

Internet Applications

Objective: Objective of this lesson is to discuss the popular Internet Applications..

Structure

4.1 Naming with Domain Name System

The Domain Name System (DNS) is a distributed database providing a hierarchical naming system for identifying hosts on the Internet. DNS was developed to solve the problems that arose when the number of hosts on the Internet grew dramatically in the early 1980s. The DNS database is a tree structure called the domain name space. Each domain (node in the tree structure) is named and can contain sub-domains. The domain name identifies the domain's position in the database in relation to its parent domain. A period (.) separates each part of the names for the network nodes of the DNS domain. For example, the DNS domain name csu.edu specifies the csu sub-domain whose parent is the edu domain; csu.com specifies the csu sub-domain whose

parent is the com domain. Following figure illustrates the parent-child relationships of DNS domains.

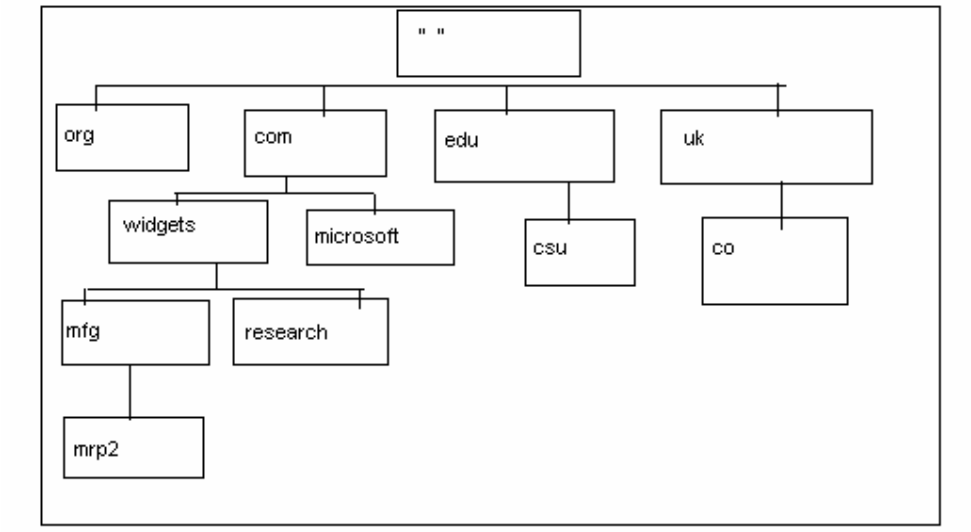

Figure 1 (Parent child relationships of DNS Domain)

As shown in figure 1 the portion of DNS database, here root node of the DNS database is null (unnamed).It is referenced in DNS names with a trailing period(.). For example, in the name : "parvinder.50megs.com.", it is the period after the com that denotes the DNS root node.

4.1.1 Top-Level Domains

The root and top-level domains of the DNS database are managed by the InterNIC. The top-level domain names are divided into three main areas:

- Organizational domains (3-character names)
- Geographical domains (2-character country codes found in ISO 3166)
- The in-addr.arpa.domain (a special domain used for address-to-name mappings)

Organizational domain names were originally used in the United States, but as the Internet began to grow internationally, it became obvious that an organizational division was inadequate for a global entity. Geographical domain names were then introduced. Even though a ".us" country domain exists, domain names in the United States are still predominantly organizational. As shown in following Table, there are currently seven organizational domains.

4.1.2 The DNS Organizational Domains

DNS domain name abbreviation	*Type of organization or institution*
Com	Commercial
Edu	Educational
Gov	Government
Org	Noncommercial
Net	Networking
Mil	Military
Int	International

4.1.3 Delegation

Responsibility for managing the DNS name space below the top level is delegated to other organizations by the InterNIC. These organizations further subdivide the name space and delegate responsibility down. This decentralized administrative model allows DNS to be autonomously managed at the levels that make the most sense for each organization involved.

4.1.4 Zones

The administrative unit for DNS is the zone. A zone is a sub-tree of the DNS database that is administered as a single separate entity. It can consist of a single domain or a domain with sub-domains. The lower-level sub-domains of a zone can also be split

into separate zone(s). Following figure 2 illustrates the relationship between DNS domains and zones.

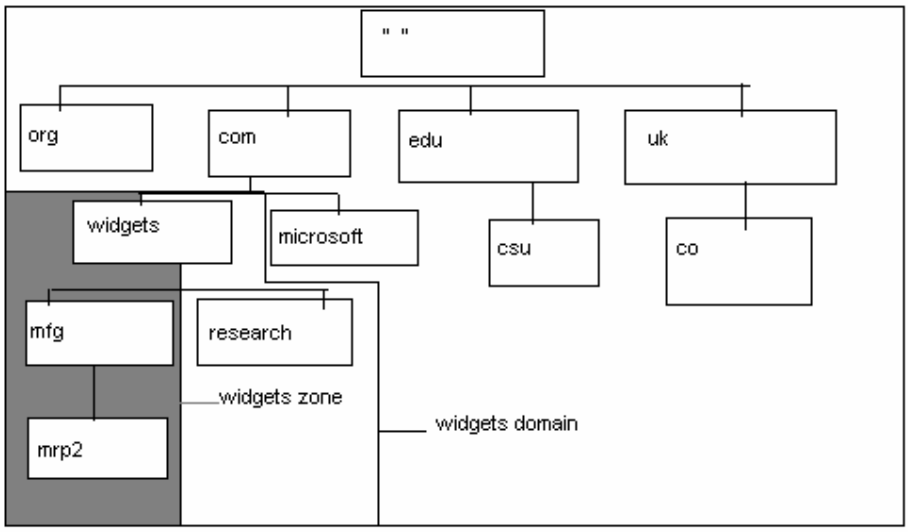

Figure 2 (Relationship b/w DNS domain and zones)

4.1.5 Fully Qualified Domain Names

With the exception of the root, each node in the DNS database has a name *(label)* of up to 63 characters. Each sub-domain must have a unique name within its parent domain. This ensures name uniqueness throughout the DNS name space. Following the path from the bottom of the DNS tree to the root forms DNS domain names. The node names are concatenated, and a period (.) separates each part. Such names are known as *fully qualified domain names* (FQDN). Here's an example of one:

mrp2.widgets.mfg.universal.co.uk.

Note In practice, most DNS host entries appear no lower than the fifth level of the DNS tree, with three or four being more typical.

4.1.6 Name Servers and Resolvers

DNS uses a client-server model, where the DNS servers *(name servers)* contain information about a portion of the DNS database *(zone)* and make this information available to clients *(resolvers).* A resolver queries a name server for information about the DNS name space. This name server can, in turn, query other name servers as it tries to respond to the query from the resolver.

A DNS zone administrator sets up one or more name servers for the zone:

- A *primary master* name server. A primary server contains the master copy of the database files with resource records for all sub-domains and hosts in the zone.
- A s*econdary master* name server. A secondary server receives a replicated copy of the database files from the primary server. When the zone structure changes, the primary master database files are modified and copied to the secondary masters. The secondary master files are never touched.
- A *caching-only* name server. Unlike a primary or secondary server, a caching-only server is not associated with any specific DNS zone(s) and contains no database files. A caching-only server starts with no knowledge of the DNS domain structure and must rely on other name servers for this information. Each time a caching-only server queries a name server and receives al1 answer, it stores the information in its cache. When additional queries come in for this information, the caching-only server answers them directly from cache. Over time, the cache will grow to include the information most often requested.

Although the DNS software does not require them, secondary servers are a good idea for the following reasons:

- **Load balance.** Secondary servers ease the load on the primary server. This can be significant in a busy network where name server queries can reach volumes of 20,000 per hour and beyond.
- **Fault tolerance.** Secondary servers allow DNS name resolution to continue when the primary server is unavailable.
- **Reduced network traffic.** Secondary servers placed in close proximity to client computers reduce inter-network traffic across routers.

4.1.7 Name Resolution

The key task for DNS is to present friendly names for users and then resolve those names to IP addresses, as required by the inter-network. Name resolution is provided through DNS by the name servers, which interpret the information in an FQDN(Fully Qualified Domain names) to find its specific address. As illustrated in following figure, the process begins when a resolver passes a query to its local name server. If the local name server does not have the data requested in the query, it queries other name servers on behalf of the resolver. In the worst-case scenario, the local name server starts at the top of the DNS tree with one of the *root name servers* and works its way down until the requested data is found.

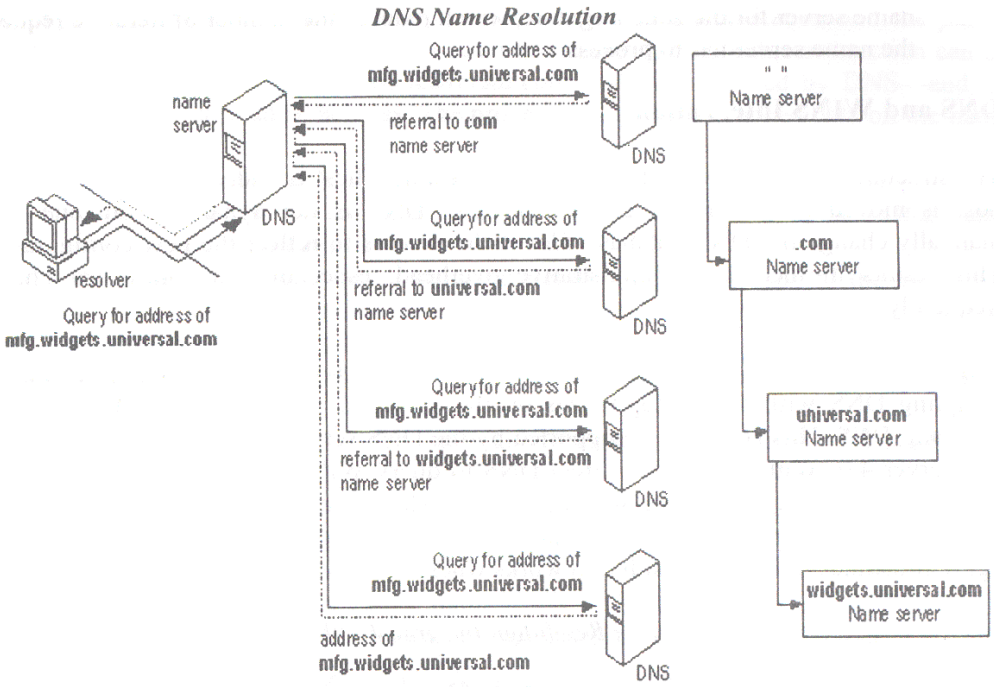

Figure 3

DNS name resolution consists of three key concepts: *recursion*, *iteration*, and *caching.*

- **Recursion**

  A resolver typically passes a *recursive resolution request* to its local name server. A recursive resolution request tells the name server that the resolver expects a complete answer to the query, not just a pointer to another name server. Recursive resolution effectively puts the workload onto the name server and allows the resolver to be small and simple.

- **Iteration**

  If the local name server cannot fully resolve the query, it enlists the aid of other DNS name servers throughout the DNS name space. A well-behaved

local name server keeps the burden of processing on itself and passes only *iterative resolution* requests to other name servers. An iterative resolution request tells the name server that the requester expects the best answer the name server can provide without help from others. If the name server has the requested data, it returns it; otherwise it returns pointers to name servers that are more likely to have the answer. However, if a primary master name server is unable to resolve a request for data that should be in its zone, it returns an error to the requester.

- **Caching**

  As local name servers process recursive requests, they discover a lot of information about the DNS domain name space. To speed the performance of DNS and ease the burden on both the inter-network and the other name servers, local name servers temporarily keep this information in a local cache. Whenever a resolver request arrives, the local name server checks both its static Information and the cache for an answer. Even if the answer is not cached, the identity of the name server for the zone might be, which reduces the number of iterative requests the name server has to process.

4.2 Electronic Mail Representation and Transfer

Electronic mail (e-mail) is one of the most widely used services on the Internet. E-mail is easy to send, read, reply to, and manage. E-mail is fast and convenient. For these reasons, e-mail has grown from a simple service offered to researchers for communicating ideas and results into a complex, talented messaging system.

- It is much easier to write an e-mail message than to write a formal paper letter or note. Many studies have shown that recipients are much more likely reply

to an e-mail message than a written request, primarily because of the ease of formulating the response.

- E-mail is very fast, the mail usually reaches its destination in a matter of seconds
- E-mail is also economic. It is much cheaper to send an email message than a letter or to make a long distance telephone call.
- You can send letters, notes, files, data or reports using the same technique
- It is asynchronous, i.e. you don't have to be present when your mail arrives. You can receive and read it at any convenient time

4.2.1 Dissecting an E-mail Message -Mail addresses

An e-mail address is the address that you use to send a message to someone and usually has the form

Login-id@domain name

The domain name is the computer system that handles the e-mail for the user. For eg. sashaya@yahoo.com

where sashya is the username or login id and yahoo.com is the domain name of the e-mail site.

4.2.2 Parts

One piece of e-mail will have the following parts:

1. Headers
2. Message Body
3. Signature

The headers are the pieces of information that tell the e-mail system a number of things about an e-mail. Each of these has a specific name and purpose. Most of these are generated by e-mail program you use. A list of common headers is :

Date	:	When the e-mail was sent
From	:	E-mail address of the sender
To	:	E-mail address of the recipient
Subject	:	Subject of the e-mail

The Message body is the content of the e-mail- what you send and what you receive. The signature isn't a signed name but a sequence of lines usually giving some information of the person who's sent the mail. It is optional.

4.2.3 Setting up e-mail

1. Before you can exchange e-mail you need to tell the application how to make the appropriate connection to the server computer handling various protocols. Your mail server uses SMTP (Simple Mail Transfer Protocol) and POP3 protocols. The following steps guide you to setting up the basic items for using e-mail in a TCP/IP environment.

2. After specifying the names of these servers as preference items in the Servers panel, software lets you send and receive email and newsgroup postings. Here are some of the preferences you can set.

3. In the Servers panel, specify your mail server names in the Outgoing Mail (SMTP) Server field and the Incoming Mail (POP3) Server field. You should specify local mail servers, if available. Often, the same server name is appropriate for both fields and the name can be as simple as mail. You should also enter your email id (the part of your email address to the left of the @ symbol; not the entire address and without the @ symbol) in the POP User Name field.

4. In the Servers panel, specify your news server name in the News (NNTP) Server field. YOU should specify a local news server, if available. Often, the name can be as simple as news.

5. In the Identity panel, specify your name, email address (the entire address including the @ symbol), and your reply-to address (if you want mail replies sent to a different address than your email address) in the designated fields. This and other information provided in the panel establishes your identity to those who receive your mail and news messages.

6. If you have an Organization panel, specify whether you want mail and news messages threaded and the method messages are sorted. By default, news messages are threaded and mail messages are not. When messages are threaded, replies are displayed adjacent to original message and other replies to the original message. When messages are unthreaded (the box unchecked) , replies are displayed according to sorted order without regard to the position of the original message or other replies.

4.2.4 Composing a Message, Queuing and Sending

To display the Message Composition window, choose File | New Mail Message or File | Mail Document in your software.

The Various fields displayed in the Message composition screen are:

- The **From** field shows your email address.
- The **Reply To** field contains the email address where you want replies to your email sent.
- The **Mail To** contains the email address where you want your message sent.
- The **Mail Cc** field contains the email address where you want a copy of your message sent.
- The **Mail Bcc** field contains the email address where you want a blind copy of your message sent. A blind copy does not display the address of the copy recipient.

- The **File Cc** field contains the location where you want to store a copy of the message you're sending. By default, messages are stored in the folder specified in the Mail and News | Servers panel.
- The **Subject** field contains a description of your email or posting. If you're sending mail this field is preset with the name of the current page.
- The **Attachment** field shows the page name or file name of any attachments you've designated.

Enter a message or include the text of the current page in the large message field. If you're sending mail, this field is preset with the current page's URL. If you have specified a text file containing your signature, the signature is appended.

Follow these steps:

- You need to know the Internet mail address of where to send your e-mail. Internet addresses typically contain a user name followed by @ symbol (pronounced “at”), followed by mail server location name. For example my e-mail address is parvinder23@rediffmail.com. Type this name in **Mail To** field to send me an e-mail.
- When you finished go to **Subject** Field.
- Type *trial message* and go to message body area.
- Type in the following message ‘This is my trial message’ and then type your name. Your e-mail message is now ready to send.

4.2.5 Queuing & Sending Mail

The Mail window offers options that let you send mail immediately or defer delivery. Typically, you'll want to send messages immediately if you're connected to the network. However, to reduce connection time, you can compose messages off-line and defer sending mail until the next time you're connected.

When the Mail window's Options | Immediate Delivery menu item is checked, the Message Composition window offers a Send Now toolbar button and menu item (Send button on Windows and UNIX). This allows messages you've written to be sent over the network immediately. Press the Send Now button or choose File | Send Now to transmit the message and any attachments to the recipient.

When the Mail window's Options | Deferred Delivery menu item is checked, the Message Composition window offers a Send Later toolbar button and menu item (Send button on Windows; Later button on UNIX). This allows messages you've written to be stored in' your Outbox folder (deferred) until you explicitly specify that the contents of the Outbox be sent. Press the Send Later button or choose File | Send Later to store the current message in the Outbox folder on your disk for transmission at a later time. Choose File | Send Mail in Outbox to distribute deferred messages.

4.2.6 Attaching Files

Pressing the Attach button produces a dialog box that lets you send email with a file attachment. An attachment is a separate document sent along with the email message. The dialog box buttons let you select a page URL or a file. You can choose to send a page as is (embedded with the HTML instructions that format Internet pages) or converted to plain text. After completing the dialog box (you can list multiple attachments), the attachment is ready to be sent along with your message.

4.2.7 Checking for mail

1. Depending on the e-mail software you are using, choose the Get Mail icon or menu option to check for e-mail message which may be lying in the server mailbox.
2. Your software may prompt you for the password. This is the same password you use to log in to the internet account
3. Enter your password and click OK

4. The application will check for the mail. A message will indicate the new mail and place it in your Inbox

5. Click OK

6. A part of the window will now display the new messages. Pressing Enter on the highlighted Message will display the text of the message.

4.2.8 Reading Mail

1. Click on Inbox

2. A part of the window will now display all messages lying in the Inbox

3. Locate the message you want to read and double- click on it.

4. The message will appear in another window. The size of this window may be changed if desired.

5. On the upper left of the screen will be the headers containing senders name and address, date, time and subject of message.

6. Double click on control menu to close message

4.2.9 Replying to Mail

Commonly, you'll send mail by replying to other mail. To do this:

1. Double click on the message in the inbox that you want to reply to

2. Choose the reply icon, or the reply to menu option. By default, the original text appears as quoted text (each line proceeded by the> symbol), though you can turn this feature off in the Composition panel.

3. Press [Ctrl + END], and enter.

4. Type in the reply to the original message.

5. Click on the queue option or the send menu option

6. Double click on the control menu to close message

4.2.10 Forwarding Mail

When you receive some mail that you want to share with a friend or a colleague, do the following:

1. Highlight the message in the Inbox list
2. Choose Message | Forward from the menu. The original message will now appear in the bottom part of the screen, each line offset by a greater than (>) sign.
3. Type the address you wish to forward the message to in the To Box
4. Click on the send button. Your message will now get sent or queued, which ever has been configured by you.

4.2.11 Printing your mail

To print a mail message on your local printer:

1.Highlight the message
2. Choose File | Print or click on the Print icon
3. Choose settings in the printer dialog box and press 'OK'

4.2.12 Deleting

To delete a message:

1. Highlight a message
2. Choose Message | Delete from the menu or Click on the delete icon
3. Open the Trash folder by clicking on it
4. The Trash folder will now have all the message you have marked for deletion. Choose File | Empty Trash folder to remove all Messages

4.2.13 Other basic features in the Mail Window

This section specifies other features which generally exist in e-mail software. Your software mayor may not have all these options.

Searching in messages: Choose the Edit | Find menu item to search for text in the message header or content panes When the first occurrence of matching text is found, the message is selected and displayed in the message field. Choose Edit | Find Again for additional occurrences.

Email addresses: Choose Message | Add to Address Book to insert the address of the sender of the selected message into your Address Book.

Marking messages: The Message menu also contains items that let you mark messages as read or unread, and flagged or unflagged. When you wish to mark or flag multiple selections, the menu item is more convenient than clicking on the small icons in the pane.

Navigation: The Go menu contains items for navigating among adjacent messages, unread messages, and flagged messages.

Viewing messages: The Options menu contains items that help determine the content of the messagc heading and message panes. You can specify that the message heading pane Show All Messages or Show Only Unread Messages. You can have header information displayed with in each message by checking the Show All Header item. Use the Document encoding menu item, if you wish to select alternative character sets. To preserve your menu changes for subsequent sessions, choose the Option | Save option menu item.

4.3 File Transfer

One of the reasons for creating the Internet was that people could exchange ideas and results of their work. FTP, short for File Transfer Protocol is used for transferring files over the Internet between a remote computer and your computer. FTP provides for access control and negotiation of file parameters. FTP provides for access control and negotiation of file parameters. FTP is most effective when you know the exact location -file name, directory name and Internet name of the remote computer where

the file is located. FTP allows you to connect with an FTP server, browse through its directories and files, and then retrieve a copy of any file useful to you.

4.3.1 Connecting to an Anonymous FTP site

An anonymous FTP site is one that allows anyone to connect and download files without having an account on that machine. You access an anonymous FTP site by giving the command ftp followed by the domain name or address of the remote system. Once ftp makes the connection to the remote site, you will be prompted for the username. Enter the name anonymous -that's where the term anonymous comes from. On being prompted for the password, enter your e-mail address as your password though it is not mandatory. Some systems may prompt you to enter guest as the username.

The ftp program running on your computer is called the client. The remote system is also running an ftp program called the server. The commands typed by you are passed by your program, the client, to the remote system. At the remote site the server program receives the commands, interprets them, and sends responses to your client.

4.3.2 Example of an FTP session

As an example we use anonymous ftp to access the Internet site nic.merit.edu, a site having a variety of information related to the Internet. Follow the steps:

1. Type ftp nic.merit.edu at the command prompt and press enter.
2. On being prompted for the username enter anonymous
3. On being prompted for the password enter your e-mail address
4. You will now see the ftp > prompt. At this prompt enter get readme.txt. The get command allows you to retrieve the file specified.
5. The screen now shows you whether the command you issued was completed successfully.
6. To end the session, enter quit.

7. To end the session is provided below:

The command will look like this example-

ftp nic.merit.edu

Connected to nic.merit.edu

220 nic.merit.edu FTP server (SunOS 4.1) ready

Name (nic.merit.edu:mozart) : anonymous

Guest login ok, send identity as password

Password:

Guest login ok, access restriction apply

ftp> get readme.txt

PORT command successful.

ASCII data connection for readme.txt (192.65.245.76., 4619) (16622 bytes)

ASCII data transfer complete.

16979 bytes received in 6.16 seconds (2.69 Kbytes/s)

ftp> quit

4.3.3 Working with files and directories

Sometimes you need to move between directories to access the file that you want. The commands for working with directories are:

dir	displays the contents of the current directory
dir<dir_name>	displays the contents of the current directory
cd<dir_name>	changes the directory to the named one
pwd	Prints the working directory name .
CTRL +S	Pause screen display.
CTRL +Q	Resume screen display.
|more	Show screen display page by page

4.3.4 Common FTP Commands

To get a list of available commands on your ftp server, type **?** at your **ftp>** prompt. To get help on a particular commands type

ftp> **help<commands_name>**

A list of the commonly available commands is given below:

debug	mget	pwd	status
dir	mkdir	quit	struct
account	disconnect	mls	quote
append	form	mode	recv
ascii	get	mput	rstatus
bell	hash	type	rhelp
binary	help	rename	user
bye	image	restart	verbose
case	mdir	put	size
cd	cdup	chmod	close
delete			

Most of these commands are regular UNIX comands available at the ftp UNIX server. However all commands are generally not available for public logins.

4.3.5 Waking with different file types

There are two general file types in the ftp context. ASCII file types are those that have only plain text, the type that are printable. Binary are non text type files like executable programs, compressed files, word processor files etc. To transfer these two different file types the following commands are used:

1. binary

Typing this command changes the mode of file transfer. After using this command you may now transfer any non text binary file type. If the ftp site is a UNIX system, this mode is set automatically for binary file types.

2. ASCII

This sets the file transfer mode to text file type. After using this command you may now transfer only text files.

Because of disk space shortage most programs available on ftp servers are compressed. To retrieve these files set the mode to **binary** before issuing the command **get**. Most sites use the popular disk compression program called **PKZip**. Files compressed using **PKZip** may easily be decompressed using PKUnzip.exe. However some sites may not use this compression program. In such a case, the compression and decompression software are provided at the site itself, and may have to be retrieved along with the useful information files.

Most ftp sites also have an information file called the **FAQ**, or the frequently asked questions. It may be named faq, faq.txt etc. For a first time user, this file is very informative and provides answers to commonly asked questions.

4.4 Remote File Access

The Internet is a system of networked computers. Networks allow users from one computer to access information or run programs on another computer on the network. In the Internet environment this is possible through a service called the **Telnet**. (This service allows a user at a remote site to access facilities, software or data at another site (remote site). While using telnet it is as though you are directly connected to the remote site. As with most UNIX systems it is important to have a valid username and password for logging onto these remote sites, though most sites allow anyone to log onto their systems.

4.4.1 Using Telnet to access a remote site

The command Telnet allows you to log in to a remote computer. In essence, this creates a virtual terminal session on the remote system as if you are connected to that system. To access a remote site, there are basically three methods:

1. Type telnet followed by a domain name.
2. Type telnet followed by the IP address.
3. Type telnet followed by the domain name and port number.

The port number is used by some sites to allow the user to by-pass the login and password requirement of most UNIX systems.

Typing any of the above will generally establish a connection with the remote site, as soon as the server responds. For eg. :

Telnet delcon.udel.edu

4.4.2 Telnet Commands

Once connected, type CTRL +] to send a command to the client. You will now see the prompt, telnet>. Type **?** to get a list of available commands. These are generally the same but may differ from server to server.

4.4.3 Examples of Telnet Sessions

1. Type Telnet followed by the name or address of the remote site. Press enter. As an example try to connect to “Shrsys.hslc.org”, this is the site of the Health Sciences Information Network. It contains holdings of various Health Science Libraries.

2. When you see Username enter SAL and press enter.
3. Press CTRL+] to give a command to the client
4. When you see the Telnet > prompt type Quit and press enter to end the session.

Some sites provide menu services on login. These menus guide the user to the available services on the site. Other sites may provide other Internet Services on login. For eg., telnet to consultant.micro.umn.edu and on the login prompt, enter Gopher. This will guide you to more than one gopher server and thereon to the information that you require.

4.4.4 What's available on Telnet?

Telnet gives you public access to some valuable sources of business information including ways to search library catalogs and other databases of business information. One such example is the Case Online Information system (COLIS) of business case study material from the Harvard Business School, Darden Business School and other notable sites from both in the US and abroad. To connect here, try site: ecch.babson.edu and enter COLIS as username.

With so many interesting sites, it is fortunate that there are some information guides to tell you what's accessible.

1. "Special Internet Connections" compiled and maintained by Scott Yanoff (e-mail address **Yanoff@alpha2.csd.uwm.edu**) is a list of sites that offer information and services on the Internet. To receive such a list, send an e-mail to **inetlist@aug3.augsburg.edu**. or check the newsgroup **alt.internet.services**.

2. "Information Sources :The Internet and Computer Mediated Communication" compiled and maintained by John December (e-mail **decemj@rpi.edu**) is another such list. To get this list connect to URL: http://www.rpi.edu/internet/Guides/decemj/icmc/internet-cmc.html

4.5 Self Assessment Questions (SAQA)

Q1. Explain parent child relationship of DNS

Q2. Explain DNS name resolution system

Q3. Explain some popular options available in E-mail system

Q4. How file is transferred using FTP?

Q5. Differentiate between FTP and Telnet.

Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Dr. Dharminder Kumar**
Lesson Number: 05	

Web Development

Objective: The objective of this lesson is to discuss the origin of web pages, important definitions related to WWW, Web Browsers and CGI technologies.

Structure:

5.1 Introduction

The World Wide Web has grown rapidly since its introduction to the world in the early 1990s. Having gained the attention of millions of people in many segments of society, the Web is now a frequent subject of (and increasingly the delivery mechanism for) mass media reports. Advertisements and editorial content in many publications such as *Newsweek, Time, Times of India,* and *The Hindustan Times* now routinely use the Web's identifying scheme, the *universal resource locator* (URL), as a means for directing the user for further information. Although Web literacy and its use is just in its nascent stage among a minority of the world's population today, the Web's capability to create a global audience for information has been recognized. Many organizations now use the Web to deliver information-ranging from government organizations such as the U.S. White House and the European

Community to major corporations like Boeing and CBS. Small businesses, organizations, and individuals all over the world also use the Web for communication, information, and interaction.

This intense interest in the Web is a result of the potential it offers for communication. Using the Web, individuals or organizations can instantaneously and continuously present hypermedia-text images, movies, and sound-to a global audience. Today, many people use the Web's potential to serve information from tens of thousands of Web servers around the world to millions of users about subjects on just about every pursuit imaginable. This vast range of content includes informal home pages that individuals create, as well as systems of information for major institutions and corporations. With such a burgeoning of information content and variation in quality and value, Web users are taxed in their ability to make choices about what information to experience. For Web developers, this information environment demands excellent, effective content development in order to rise above the information clutter. With so much information on the Web, only that which truly meets user needs well can survive and flourish.

Today the World Wide Web is a more complex system for communication than when it was introduced almost a decade ago. Although technically still based on the system of hypertext that Tim Berners-Lee and others developed at the European Laboratory for Particle Physics in Geneva, Switzerland in the late 1980s, the Web today is more diverse technologically and more diffused within society and culture.

The range of technologies a Web developer can choose from is now more varied than ever. Besides an array of techniques and tools to shape meaning with *HyperText Markup Language* (HTML), developers also can use many technologies to add new kinds of multimedia and interactive content to on-line services. New kinds of software

to observe Web content are being developed, and the competition for being the provider of Internet software has risen to the highest priority in the personal computer industry.

5.2 World Wide Web Pages

Whereas the view of the Web in 1989 was a text-based browser deployed on an internal network, today the Web is a global medium that encompasses many software and communications systems across many networks. Although sound, video, and other multimedia effects were possible before 1995-1996, systems emerged now that provided high-quality solutions to some of the problems of distributing multimedia on global networks. Notably, RealAudio (http://www.realaudio.com/) emerged as an outstanding solution to provide audio on demand over the Internet. Instead of a *click and wait* cycle of sound retrieval, RealAudio provides a streaming solution; users can listen to an audio file as it is downloaded instead of having to wait for the whole file to download. This system was a boon to the sound industry; ABC Radio News, CBC Radio, the National Public Radio of the United States, and dozens of radio stations world wide suddenly could have the Internet as a supplement to the airwaves as a broadcast medium.

Multimedia developers also gained new possibilities for networked communication with Macromedia's Shockwave product (http://www.macromedia.com/). Shockwave provides a set of plugins for its existing Macromedia authorware and enables developers to create multimedia content. Users can view innovative multimedia presentations by using the Shockwave viewers given away for free on the Macromedia site.

The method of distributing RealAudio or Shockwave content to users who obtain freely available viewer software remains a common place model for distributing new

media on the Web. *Virtual Reality Modeling Language* (VRML) (http://www.vrml.org/) is yet another major technology that emerged now as a component of a Web developer's choices for expression. VRML opens up the possibility for three-dimensional world-the creation of a cyberspace more in line with the visions of science fiction writers of the past. Combined with Java, the 3-D worlds of VRML can have behavior-the shapes and structures can respond to a user's presence and input.

5.2.1 Origins of the Web

Certainly, the idea of presenting information in a nonlinear fashion did not start with the twentieth century. The Talmud, an important document in the Jewish faith, includes commentaries and opinions of the first five books of the Bible. The Talmud's organization contains commentaries and commentaries on commentaries that extend from central paragraphs in the middle of the page. Footnotes, which are used in traditional paper texts, also have a relational, nonsequential quality that is similar to the spirit of hypertext.

Hypertext, as implemented on the Web, however, has its origins in the start of the electronic computer age, when ideas about associative linking could be married with the possibilities of automated storage-and-retrieval systems.

Vannevar Bush described a system for associatively linking information in his July 1945 article in *The Atlantic Monthly*, "As We May Think" (this article is available on the Web at http://www.isg.sfu.ca/~duchier/misc/vbush/). Bush called his system a *memex* (memory extension) and proposed it as a tool to help the human mind cope with information. Having observed that previous inventions had expanded human abilities for dealing with the physical world, Bush wanted his memex to expand human knowledge in a way that took advantage of the associative nature of human

thought. Bush's design for the memex involved technologies for recording information on film and mechanical systems for manipulation. Although the memex was never built, Bush's article defined, in detail, many concepts of associative linking and an information system to capture these in a design.

Ideas about information systems design as well as working computer systems emerged in the decades after Bush's article. In 1962, Doug Englebart began a project called *Augment* at the Stanford Research Institute. Augment's goal was to unite and cross-reference the written material of many researchers into a shared document. One portion of the project, *oN-Line System* (NLS), included several hypertext features.

In 1965, Ted Nelson coined the term *hypertext* to describe text that is not constrained to be sequential. Hypertext, as described by Nelson, links documents to form a web of relationships that draws on the possibilities for extending and augmenting the meaning of a "flat" piece of text with links to other texts. Hypertext therefore is more than just footnotes that serve as commentary or further information in a text. Instead, hypertext extends the structure of ideas by making "chunks" of ideas available for inclusion in many parts of multiple texts.

Nelson also coined the term *hypermedia*, which is hypertext not constrained to be text. Hypermedia can include multimedia pictures, graphics, sound, and movies. In 1967, Nelson proposed a global hypermedia system, *Xanadu*, which would link all world literature with provisions for automatically paying royalties to authors. Although Xanadu has never been completed, a Xanadu group did convene in 1979, and the project was bought by Autodesk, Inc. in 1988 and developed until its cancellation in 1992. Afterward, Nelson re-obtained the Xanadu trademark and, as of 1994, was working to develop the project further.

Also in 1967, a working hypertext system called *Hypertext Editing System* was operational at Brown University. Andries van Dam lead a team that developed the system, which later was used for documentation during the Apollo space missions at the Houston Manned Spacecraft Center. By 1985, another hypertext system came out of Brown University, called *Intermedia*, which included bi-directional links and the possibility for different views of hypertext, including a single-node overview and an entire hypertext structure view called a *Web view*.

Also in 1985, Xerox *Palo Alto Research Center* (PARC) (http://www.parc.xerox.com/) introduced a LISP-based system called *NoteCards*. Each node in NoteCards could contain any amount of information, but there were many types of specialized cards (50) for special data structures.

Hypertext's stature as an important approach to information organization in industry and academia was marked in 1987, when the Association for Computing Machinery (http://www.acm.org/) held its first conference on hypertext at the University of North Carolina. This was the same year that Apple Computer Corporation (http://www.apple.com/) introduced its HyperCard system. Bundled free with each Macintosh computer sold, HyperCard quickly became popular. Users organized the cards and stacks in HyperCard and took advantage of the possibilities for ordering the cards in various ways in the stack.

Vannevar Bush's, Ted Nelson's, and others' ideas about information systems showed up in another project in the late 1980s. In March 1989, Tim Berners-Lee, a researcher at the *Conseil Europeen pour la Recherche Nucleaire* (CERN) European Laboratory for Particle Physics in Geneva, Switzerland, proposed a hypertext system to enable efficient information sharing for members of the high-energy physics community. Berners-Lee had a background in text processing, real-time software, and

communications and had previously developed a hypertext system that he called *Enquire* in 1980 (at that time, he had been unaware of Nelson's term, *hypertext*). Berners-Lee's 1989 proposal, called *HyperText and CERN*, circulated for comments. The important components on the proposal follow:

- A user interface that would be consistent across all platforms and that would allow users to access information from many different computers
- A scheme for this interface to access a variety of document types and information protocols
- A provision for universal access, which would allow any user on the network to access any information

By late 1990, an operating prototype of the WWW ran on a NeXT computer, and a line-mode user interface (called www) was completed. The essential pieces of the Web were in place, although not widely available for network use.

In March 1991, the www interface was used on a network, and by May of that year, it was made available on central CERN machines. The CERN team spread the word about its system throughout the rest of 1991, announcing the availability of the files in the Usenet newsgroup `alt.hypertext` on August 19, 1991 and to the high-energy physics community through its newsletter in December 1991. In October of 1991, a gateway from the Web to *Wide-Area Information Server* (WAIS) software was completed.

During 1992, the Web continued to develop, and interest in it grew. On January 15th, the www interface became publicly available from CERN, and the CERN team demonstrated the Web to researchers internationally throughout the rest of the year. By the start of 1993, there were 50 known Web servers in existence, and the first

graphical interface (called *clients* or *browsers*) for the X Window System and the Macintosh became available in January.

Until 1993, most of the development of Web technologies came out of CERN in Switzerland. In early 1993, however, a young undergraduate at the University of Illinois at Urbana-Champaign named Marc Andreessen shifted attention to the United States. Working on a project for the *National Center for Supercomputing Applications* (NCSA), Andreessen led a team that developed an X Window System browser for the Web called *Mosaic*. Mosaic was released in alpha version in February 1993 and was among the first crop of graphical interfaces to the Web.

Mosaic-with its fresh look and graphical interface presenting the Web using a point-and-click design-fueled great interest in the Web. Berners-Lee continued promoting the Web itself, presenting a seminar at CERN in February 1993 outlining the Web's components and architecture.

Communication using the Web continued to increase throughout 1993. Data communication traffic from Web servers grew from 0.1 percent of the U.S. *National Science Foundation Network* (NSFNet) backbone traffic in March to 1.0 percent of the backbone traffic in September. Although not a complete measure of Web traffic throughout the world, the NSFNet backbone measurements give a sample of Web use. In September, NCSA released the first (1.0) operational versions of Mosaic for the X Window System, Macintosh, and Microsoft Windows platforms. By October, there were 500 known Web servers (versus 50 at the year's start). During Mecklermedia's Internet World in New York City in 1993, John Markoff, writing on the front page of the business section of *The New York Times*, hailed Mosaic as the "killer app [application]" of the Internet. The Web ended 1993 with 2.2 percent of the NSFNet backbone traffic for the month of December.

In 1994, more commercial players got into the Web game. Companies announced commercial versions of Web browser software, including Spry, Inc. Marc Andreessen and colleagues left NCSA in March to form, with Jim Clark (former chairman of Silicon Graphics), a company that later became known as Netscape Communications Corporation. By May 1994, interest in the Web was so intense that the first international conference on the WWW, held in Geneva, overflowed with attendees. By June 1994, there were 1,500 known (public) Web servers.

By mid-1994, it was clear to the original developers at CERN that the stable development of the Web should fall under the guidance of an international organization. In July, the *Massachusetts Institute of Technology* (MIT) and CERN announced the World Wide Web Organization (which later became known as the World Wide Web Consortium, or W3C). Today, the W3C (http://www.w3.org/hypertext/WWW/Consortium/) guides the technical development and standards for the evolution of the Web. The W3C is a consortium of universities and private industries, run by the *Laboratory for Computer Science* (LCS) at MIT collaborating with CERN (http://www.cern.ch/) and *Institut National de Recherche en Informatique et en Automatique* (INRIA), a French research institute in computer science (http://www.inria.fr/). The Web ended 1994 with 16 percent of the NSFNet backbone traffic for the month of December, beating out Telnet and Gopher traffic in terms of bytes transferred.

In 1995, the Web's development was marked by rapid commercialization and technical change. Netscape Communication's browser, called Netscape Navigator (nicknamed *Mozilla*) continued to include more extensions of the HTML, and issues of security for commercial cash transactions garnered much attention. By May 1995, there were more than 15,000 known public Web servers-a tenfold increase over the number from a year before. Many companies had joined the W3C by 1995, including

AT&T, Digital Equipment Corporation, Enterprise Integration Technologies, FTP Software, Hummingbird Communication, IBM, MCI, NCSA, Netscape Communications, Novell, Open Market, O'Reilly & Associates, Spyglass, and Sun Microsystems.

In May 1995, Sun Microsystems introduced its Java language in its initial form to the Internet community. The language transformed the on-line world for the rest of the year, challenging long-held strategies of many companies. By December 1995, even Microsoft recognized the Internet as a key player in personal and business communications. Microsoft announced its intent to license the Java language.

5.2.2 Important Definitions related to World Wide Web

WWW (World Wide Web)

The WWW is a hyper-text information and communications system popularly used on the Internet computer network with data communications operating according to a client/server model. Web clients (browsers) can access multi-protocol and hypermedia information (possibly by using helper applications with the browser) by using an addressing scheme.

Hypertext (and Hypermedia)

When you use the WWW, the documents that you find will be hypertext documents. Hypertext is text that contains links to other text. This allows you to quickly access other related text from the text you are currently reading. The linked text might be within the document that you are currently reading, or it might be somewhere halfway around the world.

HTML

HTML, (hypertext mark-up language) is used when writing a document that is to be displayed through the WWW. HTML is a fairly simple set of commands that describes how a document is structured. . This type of mark-up language allows you to define the part of the document, but not the formatting, so the browser that you run when reading the document can format it to best suit your display.

Links

One of the defining features of any hypertext document is Links (also known as hyper-Links); Links imply references to other documents. But they are n't just stated references like "see page 2 for more information." They are actual live links, where you can activate the link and cause whatever it references to appear on your screen. When someone writes a hypertext document, he or she can insert links to other documents that have information relevant to the text document.

URL

A URL (Uniform Resource Locater) is a complete description of an item, containing the location of the item that you want to retrieve. The location of the item can range from a file on your local disk to a file on an Internet site halfway around the world.

The URL is not limited to describing the location of WWW files. Many browsers can access a number of different Internet services, including anonymous FTP, gopher, WAIS, UseNet news, and Telnet.

A typical URL would look like this:

http://www.gju.ernet.in

Other protocols that WWW clients can use to retrieve documents are-

Protocol	**Use**
Gopher	Starts a Gopher session.
FTP	Starts an FTP session.

File	Get a file on your local disk.
Wais	Accesses a WAI server.
Telnet	Starts a Telnet session.

HTTP

Another of the goals of the WWW project was to have documents that were easy to retrieve, no matter where they resided. After it was decided to use hypertext as the standard format for WWW documents, a protocol that allowed these hypertext documents to be retrieved quickly was developed. This protocol is HTTP, the Hypertext Transfer Protocol. HTTP is a fairly simple communications protocol that takes advantage of the fact that the documents it retrieves contain information about future links the user may reference (unlike FTP or Gopher, where information about the next possible links must be transmitted via the protocol).

Home Pages

All WWW users can set up their own home page, where they can set up links to sites that they use frequently. This is generally the first screen at any site which then has links to all other sites of information. For example, if a company wishes to advertise its services, it offers, it can make a Home Page where it would describe the general features of its company. It would also offer a sort of Index to the services it offers and each item in this index could be a link to further pages of information.

5.3 Browsing

5.3.1 Web Browser

Following the good success of the World Wide Web, the number of organizations developing browsers grew explosively. Browsers are available for use on almost any

computer operating system, including Amiga, DOS, Macintosh, NeXT, RISC, Windows 95, Windows NT, Windows 2000, UNIX and Linux.

A *browser* is a software application that enables you to access the World Wide Web. You can think of a browser as your window to the Web. Change your browser and you get a whole new view of the Web. When you use Lynx, your window to the Web has only text. Text-only browsers are the original browsers for the Web. Although it may be hard to believe, several text-only browsers are still being developed.

When you use ncSA Mosaic, your window to the Web has text and graphics. Browsers that enable you to view Web documents containing text and graphics are the second generation of browsers. These browsers are largely responsible for the phenomenal success of the Web.

When you use Internet Explorer, your window has text, graphics, and live animation. Browsers that enable you to view Web documents containing text, graphics, and inline multimedia are the third generation of browsers. These browsers are driving the Web's transition to an extremely visual medium that rivals the television for information content and entertainment value.

Although the Web is increasingly commercial, you can still find freeware and shareware browsers. One reason to keep abreast of browser developments is that somewhere in the myriad of options is the gem that may one day replace Internet Explorer as top dog. Before the Microsoft entered in to browser market, Netscape Navigator was at top position. If you have heard about the Netscape Navigator, you may not think this is possible, but keep in mind that until Netscape Navigator came along a browser called Mosaic was king.

5.3.2 Internet Explorer

The most exciting browser today is Microsoft's Internet Explorer. Already Internet Explorer has gone through six versions, and every version has introduced hot new features to the Web. Internet Explorer is the premier browser that directly supports internal multimedia. With direct support for internal multimedia, users are freed from the hassles of installing and configuring helper applications to view the multimedia, and Web publishers have greater freedom to include multimedia in their publications.

Internet Explorer supports Netscape extensions to HTML. Internet Explorer also supports the SSL secure transfer protocol and more secure, transfer protocol called STT. Extensions unique to Internet Explorer include:

- Scrolling marquees
- Dynamic sources to create inline motion video
- Documents with sound tracks

Internet Explorer extensions are powerful multimedia solutions for your advanced Web publishing needs. However, only browsers capable of handling Internet Explorer extensions can use these features. Currently, Internet Explorer extensions support video in Microsoft AVI format and sound in WAV, AU, and MIDI formats. If you plan to incorporate sound and video into your Web publications, you should seriously consider using Internet Explorer extensions in addition to hypertext references to the multimedia files.

When Internet Explorer was introduced, it featured the most complete support for the HTML table model standard first proposed in HTML 3.0 and was the only browser to support the cascading style sheets standard.

Internet Explorer features complete support for HTML and has the broadest support for the latest Internet technologies and standards. It supports TrueType fonts, ActiveX, VBScript, Java, VRML, and Active VRML, and features enhancements for frames.

Internet Explorer is available for Macintosh, Windows 95, and Windows NT, and Windows 2000. International versions of Internet Explorer are available for over a dozen languages. For more information about Internet Explorer, visit Microsoft's Internet Explorer page at this URL:

http://www.microsoft.com/ie/

5.3.3 Using a browser software: Internet Explorer

Understanding pages and frames

Now, you're probably comfortable with the idea that information on the Internet is presented on pages you see on the screen. Even the navigational concepts are pretty easy:

- You start with a home page.
- You click on highlighted words (colored or underlined) in a page to bring another page of related information to your screen.
- You click on arrow buttons to go back (or forward) to a page you have previously seen.

Plus, you can go directly to pages that interest you by choosing menu items:

- History toolbar display pages you have viewed before.
- Favorites items in the Favorites menu display pages you have designated as worthy of easy access.

- Help items in the Help menu display pages that help you use Internet Explorer and Internet features.

Ideally, the act of finding pages becomes secondary to what you really care about: the page's content. Like pages of a magazine, you'll want to flip from one screen page to another, sometimes to continue with the same article and other times to begin a new article. But you can't hold screen pages in your hands like you can a magazine. Screen pages are rarely uniform in length and, displayed one page at a time, don't provide intuitive feedback on where the information begins and ends.

So even though Internet pages bring information to you rather gloriously, there is something distinctly uncomfortable about content that continues over numerous links to pages of varying lengths. Anyone who has witnessed a slide show of a neighbor's family vacation can identify with the queasy sensation of boundlessness.

Readers of electronic pages need tools to keep track of pages. The Internet Explorer text fields, toolbar buttons, and menu items provide you with the ability to manage pages of information that might otherwise leave you feeling overwhelmed and unfocused. Each time you open the explorer window (you can have multiple Internet Explorer windows open concurrently), you begin a new session of Internet interaction.

The author of a page supplies the content you initially see. Sometimes the content is presented as a single unit taking up the entire content area of the window. Other times the content is displayed in multiple rectangular frames that, together, form a patchwork of individual pages that fills the content area. Each frame can contain scroll bars to let you view more information. Internet Explorer allows you to resize any frame by positioning the mouse in the borders between frames (the cursor changes shape), then dragging the frame to a new size.

A frame within a page is, in essence, a smaller page within a large patchwork page. Each frame has characteristics of a page. Together, the frames form a top-level page (also called a frameset).

For example, clicking on a link within a frame can bring new information within the frame or to a different frame. Likewise, a link can bring an entirely new top-level page replacing all the frames.

When viewing a page with frames, certain menu items change to reflect that actions affect only a selected frame's page and not the set of pages in the top-level page. When you select a frame by clicking within it, other functions such as keyboard shortcuts affect only the contents of the frame.

Some pages and frames have the capability to automatically update themselves. Pages that have server-push and client-pull capabilities contain instructions that allow multiple interactions with the server computers. You can always terminate these automatic actions by going to another page or otherwise exiting the page.

To open a new Internet Explorer window, choose the File | New | Window menu item. The new window brings another copy of your home page to screen in a fully functional and independent Internet Explorer window. You can have simultaneous network connections.

5.3.4 Knowing that every page has a unique URL

To understand how a single page is kept distinct in a world of electronic pages, you should recognize its URL, short for Uniform Resource Locator. Every page has a unique URL just like every person has unique fingerprints.

A URL is text used for identifying and addressing an item in a computer network. In short, a URL provides location information and Netscape displays a URL in the location field. Most often you don't need to know a page's URL because the location information is included as part of a highlighted link; Internet Explorer already knows the URL when you click on highlighted text, press an arrow button, or select m a

menu item. But sometimes you won't have an automatic link and instead have only the text of the URL (perhaps from a friend or a newspaper article).
Internet Explorer gives you the opportunity to type a URL directly into the Address bar field (or the URL dialog box produced by the File | New | Window menu item). Using the URL, Internet Explorer will bring you the specified page just as if you had clicked on an automatic link.

On Window, Address bar field offers a pull down menu to right of the text. The menu contains URLs of pages whose locations you have most recently typed into the field and viewed. Choosing a URL Item from this menu brings the page to your screen again. If you are typing into address bar, you will need to enter the characters that exactly match the URL. For example, some pathnames contain the tilde character (~) which designates a particular home directory on the server.

5.3.5 Buttons in Internet Explorer

In addition to link in the content area, you can also access links using Internet Explorer buttons and menu items. Many of the links controlled by buttons and menu items bring pages you have viewed at least once before. Button links are particularly useful for going back and forth among recently viewed pages. Menu item links directly access a wide range of pages such as a history list of pages you have viewed or a favorites list of pages you (or others) have personally selected as noteworthy.

The toolbar offers the following button:

- Back displays the previous page in the history list. The history list is a reference to pages you have viewed.
- Forward displays the next page in the history list. (Available only after using the Back command or a history menu item.)

- Home displays the home page designated in your preferences.
- Stop terminates the page loading, if page transfer is in progress.
- Refresh makes sure you have latest version of current web page.
- Search finds occurrences of text in current web page and highlights them.
- Favorite allows quick and easy access to pages you have added in favorite list.
- History finds the web pages you have recently visited.

5.3.6 Search Engines

Search Engines play very important role in browsing as it is not possible to remember the addresses of all the web sites. Moreover they find the web pages of interest immediately. We only have to enter the worlds related to particular interest and search engines display the web site containing those worlds. Google (http://google.com), Yahoo(http://www.yahoo.com), MSN(http://www.msn.com) and Khoj (http://www.khoj.com) are the popular search engines. Here is example, suppose I want to find web pages containing my name – Parvinder Singh on Google and Yahoo. Figure 1 and Figure 2 display the result in Google and Yahoo respectively.

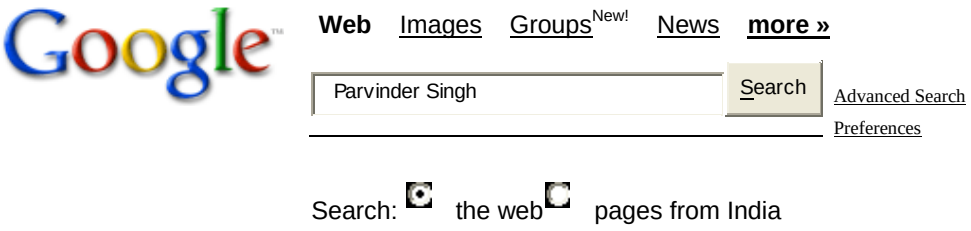

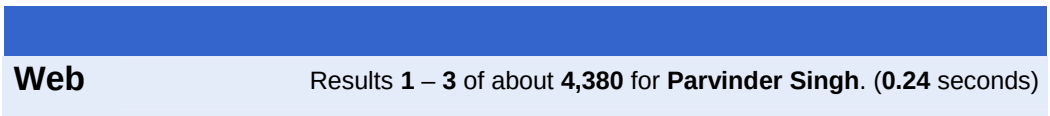

Parvinder Singh,Lecturer (Computer Science & Engineering) Guru **...**

Parvinder Singh,Lecturer (Computer Science & Engineering) Guru Jambheshwar University (GJU) Hisar Haryana(India) worked in TIT&S, Bhiwani and BRCM, Bahal, **...**

www.**parvinder**.50megs.com/ - 20k - Cached - Similar pages

Parvinder Singh offers web design india, web site india, web **...**

Parvinder Singh offers web design, web hosting, web promotion, domain registration, web designing outsourcing, software development outsourcing, **...**

www.**parvindersingh**.com/ - 11k - 22 May 2005 - Cached - Similar pages

Doon Online - Dr.**Parvinder Singh**

... In November 1998, when Dr **Parvinder Singh** was recognized as the Businessman of the **...** Dr. **Parvinder Singh** (136-J '60) passed away on July 3rd, 1999. **...**

www.doononline.net/pages/info_features/ features_spotlights/spotlights/p**singh**/ - 9k - Cached - Similar pages

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

Parvinder Singh Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Figure 1

Yahoo! My Yahoo! Mail Welcome, **Guest** [Sign In]Search Home Help

YAHOO! SEARCH

WebImagesVideoDirectoryLocalNewsProducts

Parvinder Singh Search

My Web **BETA** | Shortcuts | Advanced Search | Preferences

Search Results Results **1 - 3** of about **8,150** for **Parvinder Singh** - 0.25 sec. (About this page)

1. **Parvinder** Singh's Work

Parvinder Singh,Lecturer (Computer Science & Engineering) Guru Jambeshwar University (GJU) Hisar Haryana(India) worked in TIT&S, Bhiwani and BRCM, Bahal, offers E-Cards ... You can know more about **Parvinder Singh**, Lecturer in Guru Jambeshwar University Hisar(Haryana ... **Parvinder Singh** lecturer in India's premier institute is also working on Information ...

parvinder.rediffblogs.com - 12k - Cached - More from this site

2. **Parvinder Singh** , GJU Hisar

... **Parvinder Singh** , GJU Hisar. A blog site about **Parvinder Singh** and his works,B.E.(EC),M.Tech(CSE) presently working as Lecturer in Guru Jambheshwar ...

parvindersingh.rediffblogs.com - 9k - Cached - More from this site

3. **Parvinder Singh** offers web design india, web site india, web hosting india, software development india, software ...

Parvinder Singh offers web design, web hosting, web promotion, domain registration, web designing outsourcing, software development outsourcing, data conversion, website development, web design india, web site design india, web hosting india...

www.**parvindersingh**.com - 10k - Cached - More from this site

Results Page:

1 2 3 4 5 6 7 8 9 10 **Next**

WebImagesVideoDirectoryLocalNewsProducts

Your Search: Parvinder Singh Search

Figure 2

5.4 CGI Technology for Dynamic Web Documents

The Common Gateway Interface, or CGI, is a standard for communication between Web documents and CGI scripts you write. CGI scripting, or programming, is the act of creating a program that adheres to this standard of communication. A CGI script is simply a program that in some way communicates with your Web documents. Web documents are any kind of file used on the Web. They can be HTML documents, text files, image files, or any number of other file formats. The existence of this gateway between programs you write and your Web documents allows you to create much more dynamic and interactive Web pages than you could with HTML alone. Comparison of What HTML and CGI can do?, is given below:

Task	*CGI+HTML*	*HTML Alone*
Handle forms	Yes	No
Create almost anything non-static that needs to be on a Web page	Yes	No
Handle image maps	Yes	Yes (but only with client side image mapping)

Add searching to a Web page or set of documents	Yes	No
Create forms	Yes	Yes
Create platform-independent documents	Yes	Yes
Create applications such as chat rooms, voting booths, or anything interactive	Yes	No
Allow pages to be generated on the fly, making it easier to update a group of pages	Yes	No
Create documents tailored specifically to each user	Yes	No

5.4.1 Uses of CGI

The uses of CGI can be classified in to three categories depending upon complexity-

Simple Uses

The following is a list of simple uses in CGI:

- Hit Counters (text based)
- Programs that generate HTML for simple things like the current date, and so on

- Any Perl CGI program that is less than about 50 lines
- Any C CGI program that is less than 50 lines
- Any C++ CGI program that is less than 50 lines

Intermediate Uses

Some of the uses that fit into this category are:

- Image maps
- CGI scripts that generate entire pages of HTML
- Animation

Advanced Uses

Advance uses of CGI include:

- Database Back ending
- Search Engines
- Multiple Dynamic Pages

5.4.2 CGI and Languages

CGI applications can be written in any language that can be executed on a computer-in particular, a Web platform. In fact, you can choose any of the common languages for your CGI applications. Your choice depends on what you have to do because different languages may be specialized for different purposes. Perl, for instance, is great for string and file manipulation, while C is better for bigger, more complex programs. Perl and C are probably the most used languages for CGI programming. Feel free to choose from the following languages:

- C
- C++
- Perl
- Tcl
- Python
- Shell scripts (UNIX)
- Visual Basic
- Applescript

These languages, as well as many others, provide the programmer with the means to comply with the CGI specification and use it to its fullest potential.

5.4.3 CGI Methods

A *method* is a way of invoking a CGI program. In fact, to execute the program, you make a request to the server using a method, which defines how the program receives the data. There are three main methods, as described below:

GET Method

When you use this method, the CGI program receives the data in the `QUERY_STRING` environment variable. The program must parse (process) the string in order to interpret the data and execute the needed actions. The GET method should be used when you want to obtain data from the server and you will not change any data on the server. Exceptions may appear when the data transmitted is very long so that eventual problems in the size of the variables are prevented. In this case, the POST method is preferred.

POST Method

When you use the POST method, the Web server transmits the data to the CGI program through the stdin (standard input). The server does not mark the end of the data with an EOF character, so the program must use the `CONTENT_LENGTH` value in order to read the stdin correctly. You should use the `POST` method when the data you send will alter any data on the Web server or when you want to send large amounts of data to the CGI program (usually, more than 1024 bytes, the length limit of a URL).

HEAD Method

The HEAD method is similar to the GET method, except that with the HEAD method, only the HTTP headers (and not the data itself) are sent by the Web server to the browser.

5.5 Self Assessment Questions (SAQA)

Q1. Describe the following-

a) URL b)HTML c) VRML d) WWW e) HTTP f) CGI

Q2. Multiple Choice Questions-

i) Internet Explorer is a

a) Text based Browser b) Text and graphic based browser c) Text, Graphic, animation and sound based browser d) Animation based browser

ii) Your Home Page is

a) First page of Microsoft Site b) First Page of any site c) Google d) First Page that opens , when you open Web Browser

iii) A URL will always contain

a) http b) ftp c) Gopher d) Any of the above

iv) A Link may refer to

a) any HTML File b) any Image c) Any PDF d) All of the above

Q3. Describe the history of WWW with its origin

Q4. What is the role of browser? Explain any popular web browser.

Q5. Explain CGI technology and its usage.

Paper Code: MCA- 304	**Author: Parvinder Singh**
Paper Name: Internet Fundamentals	**Vetter: Dr. Yogesh Chaba**
Lesson Number: 06	

Network Management & Security

Objective: The objective of this lesson is to discuss the issues related with network security. Here we discuss major threats on internet and how there are dealt.

Structure

6.1 Introduction

Network Management and Security are major issues with the Internet because it is public domain. The public nature of the Internet can cause security concerns that don't exist for private intranet or dial-up applications. Because packets pass through machines over which you have no control, someone can potentially see confidential information. Any hacker with a network data scope can get credit card numbers, Social Security numbers, and other confidential information from your transmissions. We need to design for these potential security leaks.

Passing through Multiple Machines

Your transactions have the potential to pass through many computers and other devices on their way between the client and the host. On most UNIX systems, you can

issue the traceroute command to see this routing. Most of these machines are acting only as routers, but they are points where your signal can be intercepted and decoded. Anyone with a scope on any of the devices through which your information passes can trap that information. Things like Social Security numbers (999-99-9999) and credit card numbers have patterns that can be detected by automated search programs. An unscrupulous person can place one of these programs on a device routing packets along the Internet, let it work for a period of time, and then take a leisurely look at the data that it traps.

E-Mail Example

E-mail can be even more vulnerable to this type of piracy, because mail travels as plain text in a format that's easy to read, and full messages are stored and forwarded by post office machines. Although most of us don't like to look at them, and many mail readers filter them, mail headers can tell you a lot about the machines on which your mail rests. Take a look at a message header:

Received: from ns2.eds.com by mail5.netcom.com (8.6.12/Netcom)
 id NAA01582; Wed, 24 Jan 1996 13:21:17 -0800
Received: by ns2.eds.com (hello)
 id QAA07685; Wed, 24 Jan 1996 16:21:40 -0500
Received: by nnsp.eds.com (hello)
 id QAA26247; Wed, 24 Jan 1996 16:19:58 -0500
Received: from target2.sssc.slg.eds.com by dsscsun1.dssc.slg.eds.com
    (5.0/SMI-SVR4)
 id AA00143; Wed, 24 Jan 1996 15:18:57 -0600
Received: from rfbpc (rfbpc.sssc.slg.eds.com [198.132.57.4])
    by target2.sssc.slg.eds.com

The details of this heading information aren’t important for this discussion. The important thing is the fact that this piece of mail rested on four machines are not under control. At each of these points, message is simply part of a larger text file. Anyone with the proper security clearance (or anyone who can hack into that machine and obtain that clearance) can read the message. The headings are read from the bottom to the top:

- The mail originated on PC (rfbpc).
- The mail was passed to the post office machine on LAN (target2).
- The mail was forwarded by post office to division's mail handler (dsscsun1).
- The division mail post office passed the mail to corporate firewall (nnsp).
- The mail passed to the corporate post office outside the firewall (ns2).
- Finally, the mail was delivered to the post office on the Internet service provider (mail5).

Incidentally, the mail passed through several machines that aren't listed in this heading. Remember that traceroute? Mail packets have to pass through several machines on which they don't rest, making them vulnerable to snooping.

What does this mean to your application? If you're passing sensitive, private, or confidential information, consider encryption for your application.

6.2 Encryption

Many types of encryption can be used to protect your transactions. Several Web browsers and hosts are "secure" in that they encrypt information passing between them. The extent to which we want to use encryption in your application will depend on the sensitivity of the information and the cost of encryption.

Of course, if you are writing your own application in which you will provide both the client and server modules, you can provide your own custom encryption schemes. One caution about using encryption such as that used by products like Pretty Good

Privacy. These schemes are controlled by the U.S. Federal Government, which has some restrictions against exporting encryption technology overseas. Be sure to check out this issue before committing your application to specific technology or standards.

6.2.1 Secure Web Servers

If you are designing an application that will be hosted by a Web server, consider placing the application on a secure Web server. These servers establish a secure connection with the client browser and encrypt all information that passes between them. The Netscape Commerce Server, for example, uses Secure Sockets Layer (SSL) to encrypt pages during transmission.

6.2.2 Encrypting Sensitive Information

Even if you choose not to encrypt entire transmissions, never send an unencrypted password, Social Security number, credit card number, or other sensitive information over the Internet. This data can be encrypted easily by the host CGI interface program, even if you implement your program using a commercial Web hosting program. Implementing encryption at the client end of the application is more difficult if you don't rely on the encryption capabilities of the commercial server/client. Java or some other plug-in application needs to be used to encrypt the sensitive information prior to transmission.

6.2.3 Encrypting or Password-Protecting Documents

If you are going to transmit documents over the Internet, such as word processing documents, you can use the capabilities of the applications that create the documents to encrypt or password-protect the documents. For example, both Microsoft Word for Windows and Microsoft Excel can provide file-sharing passwords that must be entered before a document can be accessed.

You might also want to use the capacity of compression programs such as PKZIP to password-protect files they have compressed. With this system, even if some hacker manages to intercept a file, she will have to work hard to read it.

Following are some thoughts about using passwords:

- Use the longest password you can to protect your documents.
- Don't use common words or phrases. They are easier to remember, but also easier to crack. Random combinations of letters and numbers are best.
- Change passwords periodically. That is, if you send several documents over a period of time, make sure that you change the password at least every 30-60 days. This strategy lessens the chance that the password will be compromised.
- Don't send passwords over the Internet; transmit them via a secure means.

6.2.4 Unsecure Request, Secure Response

In the case of especially sensitive information, you can allow requests to come to your application via the public Internet. However, you might want to return the requested information via a secure medium. For example, you could allow customers to request information via the Internet and then use fax-back facilities to fax the information to their machines.

6.2.5 Verifying the Correct Client

Another difficulty in dealing with connectionless protocols is that you might need to verify that the client you are talking to is the one you think it is. Luckily, some techniques are available, as described in the following sections.

6.2.6 Trusted Addresses

Your application might accept socket connections only from "trusted" TCP/IP addresses. Web browsers send the name of the machine in the SERVER_NAME field

and the address of the remote in the REMOTE_ADDRESS field. Be aware that these fields can be faked, but they can be used in combination with user IDs and passwords to provide additional security.

6.2.7 User IDs and Passwords

Your application might ask the client for a user ID and password. For applications with custom clients, the user ID can be programmed into the client before distribution and the user can be required to enter a specific password to verify her identity. In addition, you can limit user IDs to specific TCP/IP addresses and refuse to serve ID/address pairs that don't match.

6.2.8 Cookies

If your application uses commercial browsers, you can take advantage of the capacity of some browsers-for example, Netscape or Microsoft's Internet Explorer-to store information on the client machine; that information can be returned to the server when a specific host path is requested.

CGI scripts can set data at the client's browser; this information is called a magic cookie. When a browser makes a request for a page, it sends its cookie (if it has one set) to the server along with the request. If this is the first time that this particular machine has been used to access your application, it will need to set default configurations or provide a form on which the customer can provide required information.

A magic cookie is made up of several parts:

- URLs for which the cookie will be sent
- PATH for which the cookie will be sent in the above host domain
- DATE when the browser will delete its cookie

- SECURE flag that tells the browser to send its cookie only if it has a secure connection to the server

Using this capability, you could transmit a user ID to the client and then retrieve it on subsequent visits by this client. You can match the returned cookie to security information entered by the human being on-screen as an additional security precaution.

Cookies can also be a convenient way to customize your application for a particular client; for example, when you are transmitting a page in a foreign language for international clients. Once a customer has visited your site, you can recognize the customer from his cookie, and automatically customize the page returned to him.

6.3 Firewall Concepts

The concept that stands behind the firewall approach is to allow local users to enjoy full network services within their local network and some useful services provided by the Internet while controlling outsiders' access to the local network resources. Firewall approach achieves security by isolating a specific segment of Internet topology(further Local Network) from the rest of the Internet and controlling all the traffic that comes to and leaves the Local Network.

To control the network traffic each connection of Local Network to the Internet is equipped with a firewall. Firewall's goal is to inspect and control all the traffic between the Local Network and the Internet. The traffic must be handled in such a way that all potentially "dangerous" traffic be detected and dropped and if necessary logged. What traffic is "dangerous" for the Local Network is determined by the Security Policy adopted for the site. Figure 1 shows Local Network without Firewall.

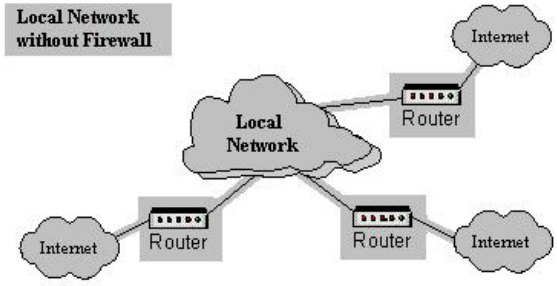

Figure 1(Local Network without Firewall)

6.3.1 Why Firewalls?

The result of firewalling the Local Network can be viewed as follows. In the case of Local Network directly connected to the Internet without any firewall, the entire network is subject to the attack. Consider a large organization with thousands of hosts. If every host is allowed to communicate directly with the outside world, the attackers will find the weakest of the hosts and penetrate it. If one of the hosts is penetrated it is not difficult to penetrate all the other hosts on the network using the resources of that compromised host. Practical experience shows that it is very difficult to ensure that every host on the network is secure. One badly chosen password and all the network security can be compromised. On the other hand if Local Network is guarded by the firewall there is direct access only to selected subset of hosts and the zone of risk is often reduced to the firewall itself or a selected subset of hosts on the network. In some sense firewall are not so much a security solution as they are a response to the engineering/administration problem: configuring a large number of hosts systems for good security.

As was mentioned above firewall must inspect all the packets that come to and leave the Local Network and filter out those packets that do not conform to the security policy adopted for the Local Network.

Remember the ISO seven layers protocol model. The packet inspection can take place on any of the layers. But packet inspection is most commonly implemented at Application layer by Application layer firewalls and at Network layer by Network layer firewalls.

Communication Layers
Application
Presentation
Session
Transport
Network
Data Link
Physical

Figure 2 (ISO Model)

When talking about TCP/IP protocol suite the Application layer firewalls are commonly called Application Gateways or Proxies(further Proxies) and Network layer firewalls Filtering Routers or Screening Routers(further Filtering Routers).

6.3.2 How packets are filtered out?

Remember the function of ordinary IP router. It receives IP datagram extracts destination IP address and consults the routing table for next hop for this datagram.

As its name indicates Filtering Router in addition to routing function performs a filtering of the packets it receives, that is before consulting the routing table it must decide whether this packet should be forwarded towards its destination.
The filtering decision is made according to the Access Control List (further ACL) associated with physical interface the packet came through.
ACL consists of the entries. Each entry specifies values for particular header fields and action to be taken if arrived packet matches these values.
Each arrived packet is matched successively against the entries in the ACL if the match occurs the action is taken. Figure 3 shows the steps of filtering decisions.

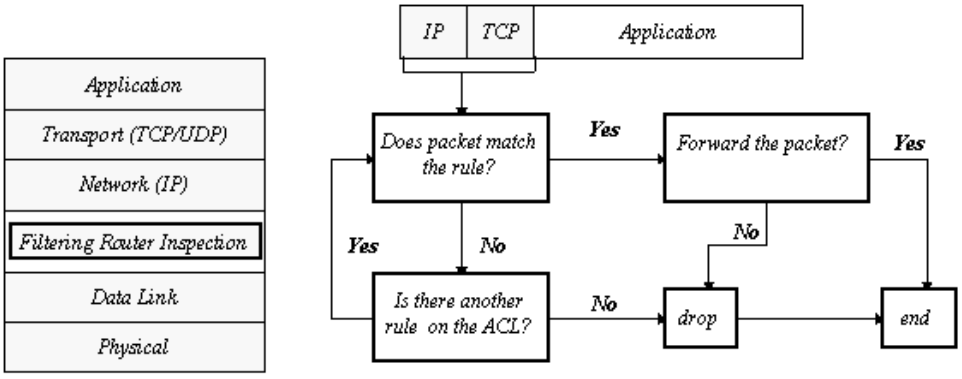

Figure 3 (Steps of Filtering Decision)

The question arises : "What is done with the packet that does not match any entry in the ACL?".
In this situation two different approaches may be adopted:

- That which is not expressively permitted is prohibited, that is these packets will be dropped by the Filtering Router;
- That which is not expressively prohibited is permitted, that is these packets will forwarded by the Filtering Router;

6.3.3 What information is used for filtering decision?

The portions that are parsed by the filtering router are IP header, and transport protocol header whether TCP or UDP. Therefore the header fields that can be used in ACL entries are:

- Source IP address (IP header),
- Destination IP address (IP header),
- Protocol Type (IP header, specifies whether the data encapsulated in the IP datagram belongs to TCP, UDP or ICMP protocol),
- Source port (TCP or UDP header),
- Destination port (TCP or UDP header),
- ACK. bit (TCP header, this bit specifies whether the packet is the acknowledgment for received TCP packet);

6.3.4 Firewall Components: Application Level Gateways

Proxy

Proxy is -The agency, function, or power of a person authorized to act as the deputy or substitute for another;

The idea that stands behind Proxy component of a firewall design is not to allow direct TCP (UDP) connection between client software on the Local Network and server software on the Internet or vice versa (the client software on the Internet and server software on the Local Network). Instead the direct connection is broken into two separate connections. The proxy program is acting as an intermediate. Operation of application level gateway is shown in figure 4.

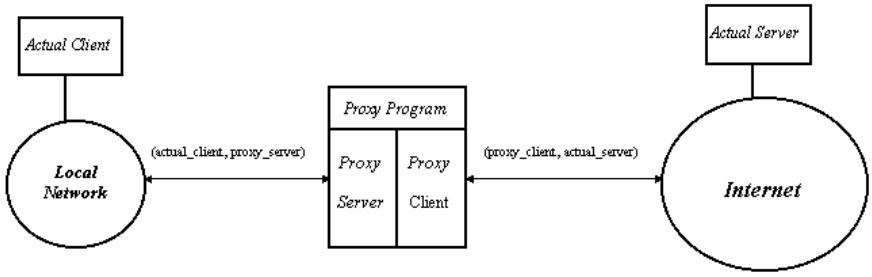

Figure 4 (Operation of Application Level Gateway)

As it relays the traffic between actual client and actual server it does checks and access controls that typical client and server software do not support.

6.3.5 How proxy program works?

Proxy program must implement enough of the client and server part of application protocol to accomplish the following:

- Accept client sessions and appear to them as a server;
- Receive from the client software the name of the actual server;
- Contact the actual server and appear to it as a client;
- Relay all the data from the client to a server;
- Perform access control function, that is according to Security Policy chosen for a site it must reject potentially dangerous connections;

6.3.6 How external client is enforced to contact proxy server?

First of all the external hosts should not know about the Local Network topology and thus should not know the IP address or name of the host machine on the Local

Network on which a specific server may be located. External hosts should only know about the Application Gateway machines.

But if external hosts does have this information it could try to contact the internal server. To prevent this the IP connectivity between Local Network and the Internet must be broken. This can be achieved in several ways please refer Firewall Architectures part of this document.

6.3.7 How internal client is enforced to contact proxy server?

Once more the IP connectivity between Local Network and the Internet must be broken.

In addition to this the client software on the Local Network must know how to contact the proxy server instead of the actual server on the Internet. To accomplish this two proxy technologies exist: classical proxy technology and transparent proxy technology.

In classical proxy technology either the client software is modified or the user is instructed to follow special setup procedures in order to make call to the actual server through the proxy server.

In transparent proxy technology the routing tables of the Local Network are configured in such a way that all the packets destined for the external servers come to the Application Gateway machine and proxy program knows how to intercept these packets and to form two connections (actual_client, proxy_server) and (proxy_client, actual_server).

6.3.8 Firewall Architectures

Screening Router Architecture

In this architecture a firewall consists of nothing more than a screening router. Figure 5 shows this architecture. Host on the Local Network and hosts on the Internet are allowed to communicate directly. The communication is restricted to the type that is

allowed by a screening router. The security of the whole Local Network depends on the correct ACL of the router and on the amount of services permitted.

Dual-Homed Host Architecture

In this architecture a firewall consists of Dual-Homed Host machine (machine having two or more IP addresses each for specific physical port). One port of the machine connects to the Local Network and the other port/ports connects to the Internet. The IP datagram forwarding is turned off on the Dual-Homed Host machine, thus there is no direct TCP/IP connection between the Local Network and the Internet. Figure 6 shows this architecture.

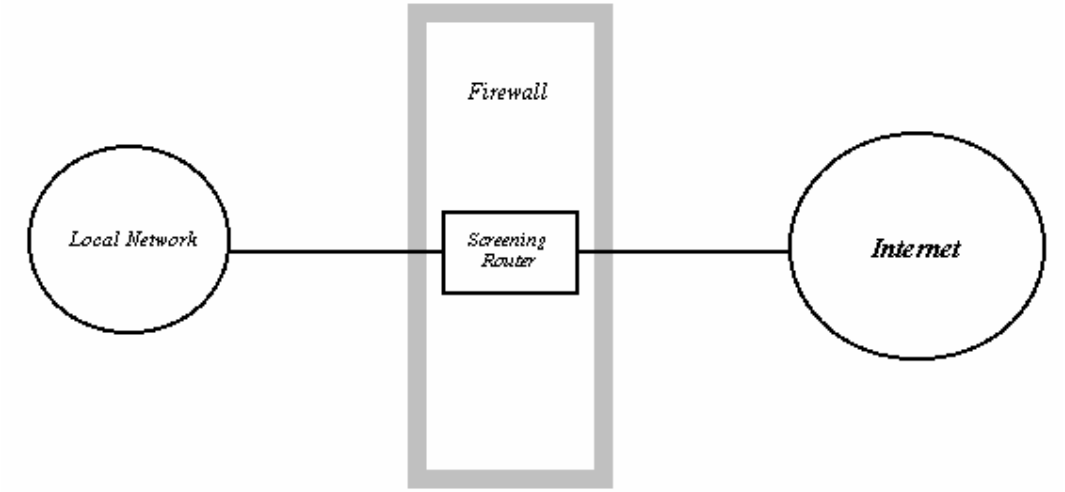

Figure 5

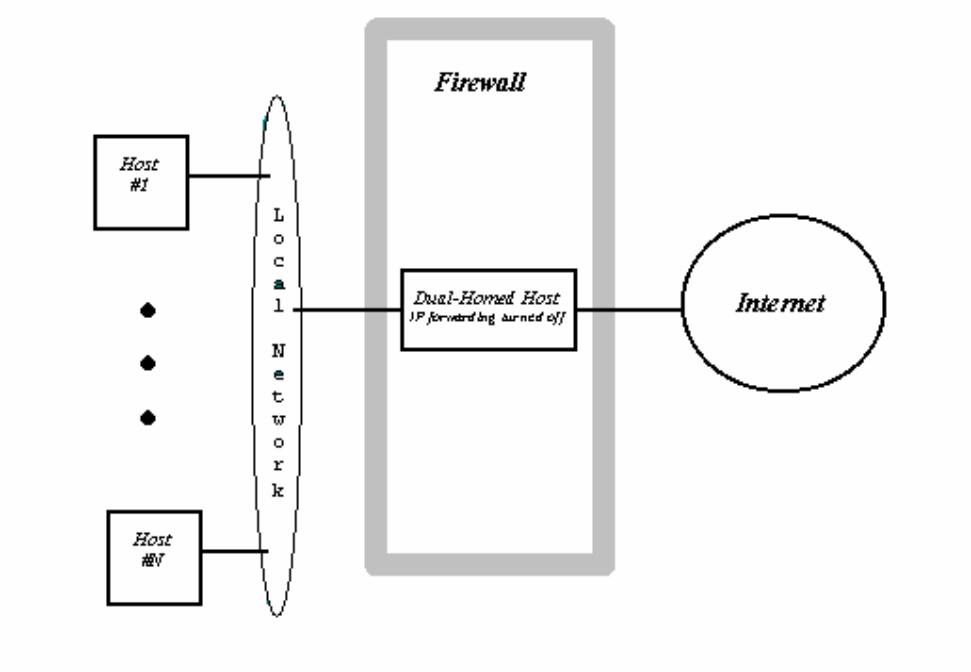

Figure 6

You permit communication between Local Network and the Internet in either of two ways:

- Users on the Local Network are given accounts on the Dual-Homed Host machine. In order to use Internet services the must rlogin on the Dual-Homed Host machine. The fact that you allow accounts on the machine weakens its security greatly (it now depends on each user and user that have access to it, more correctly it depends on the users' ability to choose "strong" passwords). Once the outsider succeeds to rlogin on the Dual-Homed Host machine he/she can access the entire Local Network.
- Dual-Homed Host runs proxy program for each service you want to permit, thus there is no more need for users to rlogin to the machine in order to access the Internet. They can communicate via proxy software.

The only host that can be accessed and thus attacked from the Internet is the Dual-Homed host machine. Thus it must have much greater level of security than the ordinary host on the Local Network. The excessive logging and auditing of system state must be performed, only secure software and necessary software installed and so on.

This architecture is much more secure than the Screening Router Architecture. But still once the Dual-Homed Host is subverted the entire Local Network is vulnerable to attack.

Screened Host Architecture

This architecture consists of the Screening Router and Screened Host. Figure 7 shows this architecture.

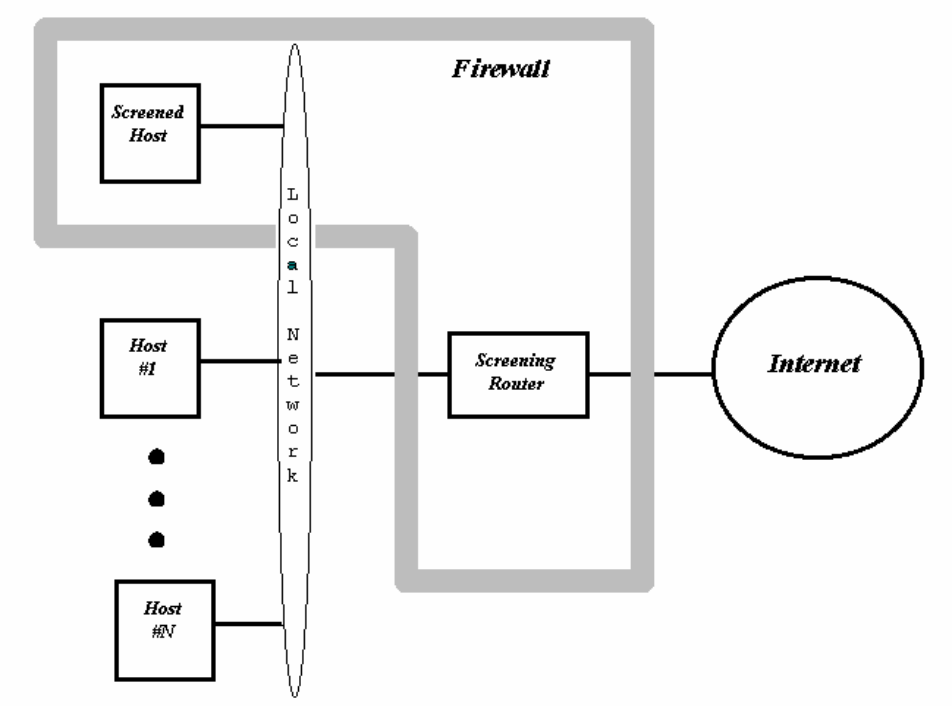

Figure 7

Screening Router is placed between the Local Network and the Internet and its role is to block all the traffic between those two networks but the one that originates on the Internet and goes to the Screened Host or the one that originates on the Screened Host and destined for the Internet. That is Screening Router stops all the attempts to setup direct communication between ordinary host on the Local Network and the host on the Internet.

Screened Host is the host on the Local Network. It is the only host on the Local Network that can be accessed from the Internet and usually will run proxy programs for the allowed services. The other hosts on the Local Network must communicate with the Internet through proxy servers located on the Screened Host.

This architecture is more flexible than that of Dual-Homed Host with proxy services, because some secure services for which proxy software does not exist can be allowed to pass through Screening Router directly to a host on the Local Network.

Screened Host is also the only host that is subject to attack on an initial attempt. Thus an extra attention is paid to its security (because of this fact it is sometimes called in the literature "Bastion Host"). Once the Screened Host is subverted the attackers have access to all the hosts on the Local Network.

Screened Subnet Architecture

This architecture consists of the Screening Routers and Screened Hosts combined in such a way that when one of Screened Hosts is subverted the Local Network is not automatically open for an attack. In the figure bellow the Screened Subnet Architecture is shown using two Screening Routers and one Screened Host. Figure 8 shows this architecture.

What are the differences from the Screened Host Architecture?

Screened Host is placed on the different physical segment than other hosts on the Local Network. Suppose that Screened Host is subverted. If it was connected to the same physical segment as other hosts in many network technologies it could to sniff all the traffic passing on the segment.

Local Network is guarded from the Screened Host by additional Screening Router. Thus in order to attack Local Network the attacker must pass through this additional router.

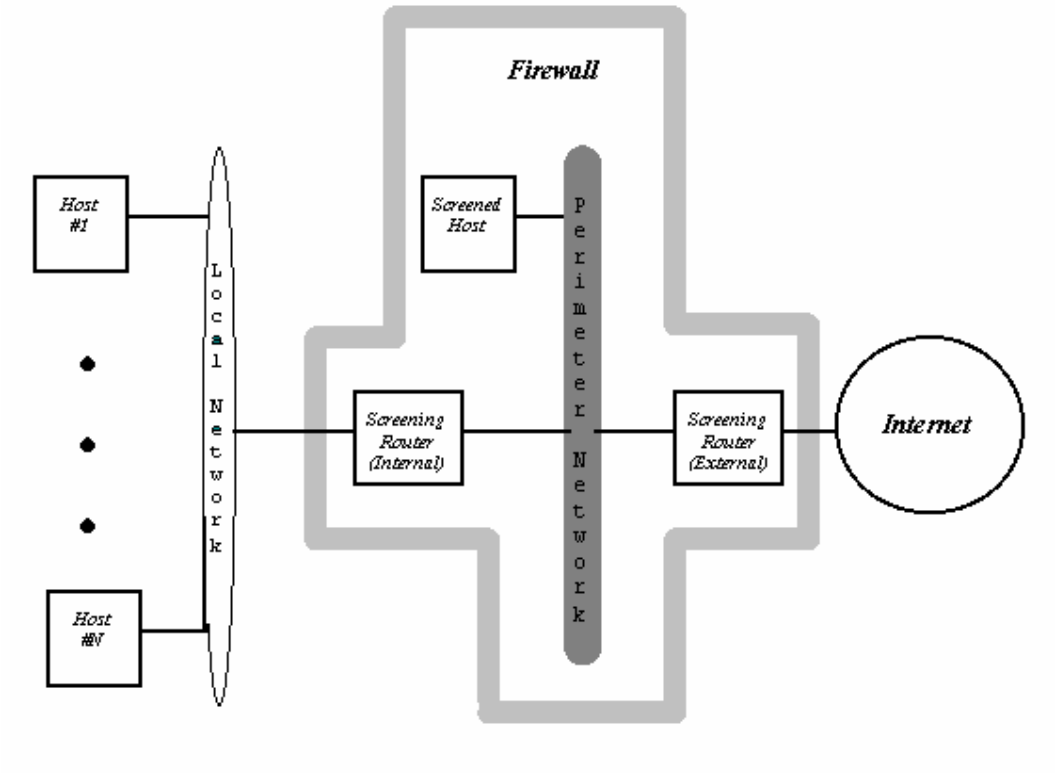

Figure 8

6.4 Dangers on the Internet

It is often the case that users are not prepared to deal with the issue of security or to protect themselves against security risks such as viruses or dialers. Many users only

react when the damage is done, to the system or their finances. Surveys show that many users only seriously deal with the issue of security when it is too late - when a virus or other damaging applications have emerged.

6.4.1 Dialers

One of the biggest threats for users who are directly connected with the Internet through a telephone connection such as ISDN or for modem users is dialers. These are dial-up programs, which use a phone line to create a connection with another system and in the past used expensive 0190 numbers.

These have become now 0900 numbers after a change in legislation. Nevertheless many PC users still incur huge financial losses due to these dialers. One of the most expensive dialers ever, called "Whirlpool", cost 900 Euros per dial-up to the Internet. Others cost approximately 80 Euros per minute. Since the change in legislation at the beginning of 2004, a maximum of 30 Euros per hour is permitted, and by law it must be disconnected after 60 minutes.

6.4.2 Viruses and worms

A virus is a program that nestles in other files and infects these. The infected file acts as a "host" through which the virus is further spread. While the computer is being infected the virus can destroy or manipulate existing data in the file and render them unusable. This can cause a loss of data.

Worms on the other hand do not require a host in order to be spread. They are independent programs which can be copied to other computers through security gaps for instance. However, they generally reach your computer by e-mail attachment. In order to activate the worm, the user simply has to run the attachment.

As a matter of fact, recently the number of worms has increased for which the user does not even need to do this in order to activate the worm. Some worms use security gaps in mail programs such as Outlook, through which the worm can be activated by

reading a mail. After successfully infecting the computer, the worms are automatically sent to all e-mail addresses saved in the address book.

Pests also nestle in many Office products. You should be particularly skeptical when working with documents that contain macros. Macro is a programming language which controls and interprets files. The macro virus exploits exactly this option. It infects other files by executing its code. It mainly affects Excel tables and Word documents.

However, you should still be wary of tempting offers on Web pages which can only be accessed by using specific software.

6.4.3 Trojan horses

These are programs which give hackers access to external computers. Such programs pose as useful software or games for the users, in order to install a dangerous pest in the system, in the background… Now the hacker has the full control of your computer. They can access, manipulate or even delete data. Other Trojans carry out attacks on other computers so your PC is involved in such activities without you knowing about it.

6.4.4 Spyware

Spyware is a relatively new type of pest program. These programs collect information about the computer and its users and re-transmit information about the online behavior of the user to the author of the program without informing the PC user about these processes.

Most programs are circulated by online companies, in order to carry out targeted advertising for their products. Frequently these programs sneak into the system through incorrectly configured browsers. Others are installed on the computer through software such as music exchange marketplaces.

However, hackers have also discovered the benefits of these programs for themselves. They are less interested in the online behavior of the user and more in the keyboard entries. Keyloggers log all keyboard entries and send these back to the hackers. Passwords, access data for bank accounts, private correspondences - everything is recorded without the PC user noticing anything. After a certain period, the recorded information is simply sent by e-mail to a mailbox set up by the hacker.

6.5 Self Assessment Questions (SAQA)

Q1. Write short notes on-

i) Encryption ii) Cookies

Q2. What is the use of Firewall? How it is implemented?

Q3. Explain different firewall architecture

Q4. What are dangers on Internet? How they can be dealed?