Internet and Intranet Administration

Course Designer and Acquisition Editor

Centre for Information Technology and Engineering

Manonmaniam Sundaranar University

Tirunelveli

Internet and Intranet Administration

CONTENTS

Lecture 1

Internet Fundamentals

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the International networking concepts.
- Describe the Linkage of the Internet.
- Describe new features of Internet tools.
- Describe about various types of Internet tools.
- Describe the Internet address.

Coverage Plan

Lecture 1

1.1 Snap Shot

In this lecture we learn about the Internet concepts is consist of a set of connected networks that act as an integrated whole. The chief advantage of an Internet is that it provides universal interconnections while allowing individual groups to use whichever network hardware is best suited to their needs. An international fabric of interconnected government, education, and business computer networks in effect, a network of networks. A person at a computer terminal or personal computer with the proper software communications across the internet by placing data in an IP packet and "addressing" the packet to a particular destination on the Internet. TCP guarantees end-to-end integrity.

1.2 Internet

Internet was a stock it would be considered a market phenomenon, with sustained doubling growth and no apparent end in sight to the upward spiral. Recent Internet numbers are stunning. Between January 1993 and January 1994 alone, the number of nodes grew from 1,313,00 to 2,217,000, an impressive 69% increase. Over 70 countries have full Internet connectivity and about 150 have at least E-mail services.

No discussion of telecommunications would be complete without discussing the Internet, referred to by its millions of users as the Net. The Net evolved from ARPANET, a research network created and subsidized in the 1960s by the Defense Department and the National Science Foundation (NSF) to link research institutions and government agencies around the world to exchange information on a wide range of topics. Although the original ARPANET no longer exists, its design and architecture laid the foundations for the Internet.

The Net, one of the oldest long distance networks in the country, is a network of networks. It links approximately 1.5 million computers, attached to more than 13,000 networks, in 100 or more countries.

Businesses use the Net for a number of tasks, such as sharing files, sending E-mail, and selling goods and services. Fewer than 1% of major companies today are not attached to the Internet. In fact, the Internet has become such an integral part of the corporate landscape that it is generating a number of entrepreneurial opportunities and highly specialized jobs: there are Internet explorers, security experts, technicians, librarians, trainers, and other service providers. In fact, the Internet is becoming so common that many employers expect their new hires to have a working knowledge of it.

Think of the Internet as a huge repository of information on almost every topic imaginable. People all over the world can search the net for information, add new information, and exchange views on different topics. The Internet is an electronic web that connects people and businesses that have access to networks and allows them to send and receive E-mail and to participate in a number of other activities, around the clock. In fact, the Internet is so huge, with such a wide variety of features, that there are few, if any, experts in the world who know everything about it.

Thus in spite of its enormous power and potential, no one really owns the Internet, although some segments of it may have their own funding and guidelines.

The primary workings of the Internet are funded by the National Science Foundation (NSF); the Internet Engineering Task Force (IETF), a committee of scientists and experts, provides technical supervision, standards, and guidelines for the Net. Any network connected to the Internet must abide by the standards established by the Internet Architecture Board. (IAB).

1.3 Linking to the Internet

Internet has undergone explosive growth, primarily because of the exponential growth in PCs, both at work and in the home. Today, 70% of PC sales are targeted at the home market. Andy Grove, CEO of Intel (which makes Pentium and there chips for the PC market), points out that every 2 years the U.S. makes as many PCs as have existed in all preceding years. Before the spread of PCs, the Internet could be accessed only if an organization was site or node on the Net. Today, one can access the Internet through a company’s connection or through commercial online services or simply from a home PC. Further, there are a number of PC-based interfaces that facilitate interaction with the Net. For example, Cyberdog, a Macintosh-based tool, allows users to directly access the Internet even without quitting an application, such as word processing. It also allows them to link and update their documents with information from the Internet. Such user-friendly interfaces are another reason for the increased usage of the Internet.

Today, there are a number of commercial online service provides, such as CompuServe, Prodigy, and America Online. It is estimated that there are approximately 6 million users of commercial online services: this number is expected to grow to more than 13 million by 1998. CompuServe claims to have 2 million subscribers, mostly businessmen: Prodigy claims 2 million, including more women and children than CompuServe; America Online claims 1 million subscribers-a younger, more consumer oriented crowd.

Famous Provider: America Online (AOL), Prodigy, CompuServe, E- World.

1.4 Internet Address

Connection to the Internet opens the door to a vast world of information and communications that resides in files on hose computers, including any computer that follows the TCP/IP protocol. Finding a file on the Internet is similar to finding the location of an individual in a city; the user needs the file's address. Addresses are the key to receiving and sending information on the Internet. All individual Internet addresses follow the same two-element pattern: the person's userID, followed by an @, followed by the name of the person's hose computer. An Internet service provider gives addresses to individuals and organizations that register with it. Let us look closely at the element of an Internet address.

The address uniquely identifies the individual or the organization that the user is trying to reach. For an individual, the name is made up of two parts separated by the character @. For example, if Joe Green is accessing the Net from his home, using his American Online (AOL) account, his internet address might be joe@aol.com. "aol" identifies the host computer or network: ".com" indicates the domain in which AOL resides.(".com" – commercial sites, ".edu" – education sites, ".mil" – military ones, and so on). If Joe Green is connecting to the internet from, say, a machine in the computer science department at New Mexico State University, his address might be jgreen@cs.NMSU.edu.jgreen is his userID; cs.NMSU.edu, which identifies the host computer, is known as the host name.

Thus, all Internet addresses have the following format or some slight variation thereof:

usrid@name-of -the-computer.name of division/department/.name of the institution / organization.type of institution or organization.
For example, let us take the Internet address

hford@hollwood.faclub.columbia.com

The userid is hford (Harrison Ford): the name of the computer which receives the Mail is Hollywood which is part of the fanclub (department) at Columbia Pictures (organization), a commercial entity.

1.5 Internet Tools

Some of the important and popular tools on the Net are discussed below. These include

- Information retrieval tools
- Communication tools
- Multimedia information tools
- Information search tools

Each of these is briefly discussed below.

1.6 Information Retrieval Tools

FTP

File transfer protocol (ftp) which was one of the first tools on the Internet, allows users to move files, such as text, graphics, sound, and so on, from one computer to another. It is a command that activates a type of client-server relationship. FTP works as follows: The user first uses the software on his or her machine called the client, to gain access to the remote machine, called the server. The user's client program communicates with a program on the remote computer, to either upload, or send to the remote computer or download certain requested files from it.

Another popular way to retrieve files is by using what is called an anonymous ftp server. In this case, the user logs on to the server using the special user ID anonymous. If the server then asks for a password, the user types in his or user ID. Freeware and shareware – programs that are available at no cost on the Net – can be obtained through anonymous ftp servers. Although it is easy to retrieve the information once the site is located, sometimes it is difficult to locate on ftp site or identify the files available on that site.

Gopher

The second type of information retrieval tool available on the Internet is Gopher, a menu-based interface that provides easy access to information residing on special servers, called Gopher sites. Although Gopher performs primarily the same tasks as the ftp command, its interface is much more user-friendly and it provides additional functions, such as links to other Internet services. By selecting an item on the Gopher menu, users can move, retrieve, or display files from remote sites. The menu also allows users to move from one Gopher site to another, where each site provides different information. The entire Gopherspace can be easily expanded by adding more servers.

1.7 Communication Tools

Communication tools facilitate written communications and in this section we discuss three types of communication tools: E-mail, Telnet and Usenet.

E-mail

E-mail, which refers to sending messages or files electronically, was one of the first Internet tools.

Telnet

Telnet is a command that connects the user to a remote machine which may be located anywhere on the Internet and the user can then type commands to the remote machine, for example to change directories in search of certain files. While the ftp only allows users to move or transfer files, the services that Telnet provides depend on the services provided by the hose machine, which may include much more than simple file transfers. For example, some servers are dedicated to the playing of board games, such as chess or go.

Usenet

The Usenet is a network that provides users with discussion groups or forums. A user posts an article to a chosen newsgroup on the Usenet, where each newsgroup is devoted to particular topic such as politics, the environment, gun control, surfing, and so on. The article is routed only to those sites that have expressed an interest in receiving information on the topic.

Many find the Usenet to be helpful for gathering information on a variety of topics. For example, vendors monitor the technical forums on the Usenet closely to answer technical questions that customers may have and to stomp out any misinformation or rumors about their company or its products. For instance, the calculation error in Intel Corp's Pentium microprocessor was first discussed on an Internet Usenet newsgroup. Another popular application on Usenet is downloading new or upgraded software from vendors and obtaining quotes for products and services. IS managers are also tapping into the Usenet to search for potential employees. One IS manager posts particularly thorny technical problems on the Usenet in the relevant discussion groups and then makes an offer to individuals who gave some of the best replies.

1.8 Multimedia Information Tools

World Wide Web. The World Wide Web (www) is one of the newest and most popular hypertext-based Internet tool. It allows users to access and display documents and graphics stores on any server on the Internet. In 1989, Tim Berners-Lee, a computer scientist at the CERN particle physics lab in Switzerland, designed WWW as a tool to help and international group of physicists exchange findings and information related to their research. It became popular only in late 1993, when WWW software was delivered for desktops that used popular operating systems, such as Windows, allowing users easy and friendly access to the Internet. Some WWW terms are shown as follows:

CGI Common gateway interface. A specification for a communication interface between external systems and a Web server.

HTML Hyper-Text Markup Language. A text formatting language used to create Web pages.

HTTP Hyper-Text Transport Protocol. The communications protocol between a Web server and a client.

Image Map A text file that defines regions using graphics files that users can click on to navigate.

MIME Multipurpose Internet Mail Extensions. An Internet standard for sending and receiving multimedia E-mail.

Server-Side A Web server feature that lets HTML pages be parsed “on the fly”, filling in parts of the page with information from external files.

S-HTTP Secure Hyper-Text Transport protocol. An extension of HTTP that provides communication and transaction security for Web clients and servers.

SSL Secure socket layer. A transport mechanism developed by Netscape Communication for transmitting secure data over a network.

URL Uniform resource locator. An address that specifies the location of a Web page.

There are many interfaces to the WWW, such as Mosaic, a GUI – based hypertext browser, and Netscape, that allows users to easily navigate the Internet and access it many services. These easy-to-use interfaces are playing an important role in popularizing user of these interfaces to create their own Web pages. According to Michael J. Walsh, president of Internet Info., in Falls Church. Virgina, the number of companies with .com addresses is expected to double within the next year, primarily because it is now easy to use the Internet and gain 24-hour access to customers all over the world.

Home page

The World Wide Web has also spawned new services, known as service providers that provide a multitude of services on the Web. For example, a Web service provider can help a company establish a home page on the Net, an electronic description of the company and its products and services, similar to a catalog or a brochure. Just as catalogs should be attractive and grab the interest of customers, the home page should also be attractive and encourage the user to further explore the company and its products.

Basic requirements to set up a home page include an Integrated Services Digital Network (ISDN) line, software to generate Hypertext Markup Language (HTML), which is the graphical interface to the internet, and a PC or Unix Server that has between 16M bytes and 32 M bytes of RAM and runs at 100 Mhs or more.

Many companies are setting up home pages to attract new customers. For example, GE Plastics found that as soon as it set up a Web page, it was getting approximately 12,000 potential inquires each month its products. Before the page was set up, it took nearly 3 days to respond to customer queries: now customers can get information within a few minutes. The company firmly believes that setting up a Web page has given it a competitive advantage. Further, the cost of setting up a Web page was much lower than that of printing pamphlets and mailing catalogs to customers. The advantages of a Web page are so great that, as one spokesperson said, "You simply can't lose. There is nothing equivalent to the Net".
Security APL is an example of another company that relies on the internet to enhance internal and external communications.

In 1995, when statewide elections were held in California, tens of thousands of voters accessed a World Wide Web Server that contained poll reports culled every five minutes from the state's Election Web Server. Clinton's 1992 presidential campaign was notable for making his E-mail address public, and the White House recently established its own World-Wide Web Home Page(try whitehouse.gov).

United Parcel Service (UPS) is another company Spokesperson says, "Our competitors aren't online yet, and that gives us an edge where we can generate some allegiance to our service".

Although there are many advantages to a home page, some experts are worried that the Internet may simply become a "de facto national post office", and may become heavily commercialized. As more and more businesses start using the Net to sell their products and services, the Net may not be capable of handling the expanded volume of traffic. Ray Hoving, chairman of the Society for Information Management's (SIM) National Data Highways Advisory Council, is concerned about the technical, administrative, and security issues that will arise as traffic on the Net continues to increase.

1.9 Information Search Tools

Archie

The tools discussed in this section help users locate information on the Net. For example, Archie, one of the first information search tools developed on the Internet, periodically searches anonymous ftp servers that participate in the Archie database and identifies all files on these servers. It then creates a central index of all files available on anonymous ftp sites and creates a central database that users can access to locate information.

Users who have the Archie software, or who can use the Telnet command to connect to an Archie server, can access this database. The only disadvantage is that the user must know at least part of the filename in order to be able to search for the file. When the desired file is located, the servicer idenfies the file's address and the user can use the ftp command to access and retrieve the file. Though Archie is a very useful tool for locating files, it must be noted that not all anonymous ftp sites participate in the Archie database, and therefore, the database is not a comprehensive one.

Veronica

Veronica is a search tool designed specifically to locate all files on Gopher sites and it is listed under Other Gopher and Information Servers on the Gopher menu. Users who have access to a Gopher server and a Veronica server can then access the database. Note however, that some Gopher serves may choose not to participate in the Veronica Service. One of the biggest advantages of Veronica over Archie is that the user does not have to know the filename; phrases descriptive of the search area will do.

WAIS

WAIS (Wide Area Information Server). Pronounced "ways", is a search system that accesses servers all over the world to locate requested files. The WAIS database has an index of keyboards that helps users to locate files in topic areas of interest to them. When given the keywords, WAIS returns the addresses where the files are located. The user can then user one of the other services discussed above to download the files. If a particular file is not found on a given WAIS server, the server will automatically direct the query to other WAIS servers on the Net.

1.10 Short Summary

The Internet represents a network of interconnected networks without a fixed boundary. The use of TCP/IP applications and protocols in a private network environment results in the creation of a corporate intranet. Both the Internet and corporate intranets represent a rapidly and dynamically evolving communications environment.

1.11 Brain Storm

1. What is Internet?

2. Explain the connection of Internet?

3. Explain the various tools of Internet?

4. Explain about the communication tools?

5. What is Internet address explain the various parts of Internet address.

Lecture 2

Intranet Fundamentals

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Intranet concepts.
- Describe the Intranets Vs. groupware.
- Describe about Intranet hardware.
- Describe about Intranet Software.
- Describe the Intranet services.
- Discuss about Extranet.

Coverage Plan

Lecture 2

2.1 Snap shot

2.2 Intranet

2.3 Intranet Vs. Groupware

2.4 Intranet hardware

2.5 Intranet software

2.6 Intranet services

2.7 Extranet

2.8 Short summary

2.9 Brain storm

2.1 Snap Shot

In this session we discuss about Intranet concepts. These networks are corporate networks, basically portions of or overlays to traditional enterprise networks, which use the same or overlays to traditional enterprise networks, which use the same lower-layer and application-level protocols as the internet, specifically WWW-related technology.

2.2 Intranet

With the explosion of intranets within the corporate MIS world, many people are asking the question: What exactly is an intranet? You could read many articles on intranets and intranet technologies without ever coming across a definition of an intranet. Here is a simple definition:

An intranet is a private system that uses hardware and software developed for the Internet to provide communication, information management, and information publishing services within an organization.

In the beginning, there were great, big computers called mainframes. The mainframe computer lived in its own specially constructed, temperature-controlled room where it was tended by specially trained computed priests and acolytes. When mortals wanted the computer to do something for them, they went to the computer room and asked very politely. Sometime – perhaps days – later, a stack of paper would be delivered to them with the computer's answer to their question. If the information they got wasn't what they really wanted, it could only be because they hadn't asked exactly the right question in exactly the right way. When that happened, they had to start all over again. This way of getting the computer to do things was called "batch processing".

By and by, a new king of computing came to be: interactive processing. Computer terminals, comprising a display screen and keyboard, were placed throughout the organization and connected by cables to the mainframe computer. With interactive processing, people could work with the computer directly, asking the questions and getting the results on their terminals. That speeded things up quite a bit, but there were still a couple of drawbacks: First, the more people who wanted to use the mainframe computer at one time, the slower it went. Second, if the mainframe quit working for some reason, everyone was out of luck.

Then came the personal computer. The advent of small, relatively inexpensive computers meant that users could have their own private computers right on their desks. No longer were they at the mercy of the mainframe.

There were a couple of new problems with personal computers, however. For one thing, having individual computers made it difficult for users to share information. Another problem was that each computer had to have its own printer and enough disk space to store all the work it was doing. Printers and disk drives were very expensive. Organizations soon found they were spending more for peripheral equipment than they were for the computers themselves.

To solve that problem, organizations connected their computers, printers, big disk drives, and other expensive peripherals on networks. With the right software, users on a network could send messages to each other, share files, and do their printing on high-speed, centrally located printers.

Network technology has progressed steadily over the years. Local-area networks (LANs) offer convenient connectivity for workgroups. Fiber-optic cables, high-speed telephone lines, and satellite links have made it possible for organizations to build private wide-area networks (WANs) that span the world.

But in spite of the advances in networking, the model for managing and sharing information has stayed pretty much the same: If we can figure out whether the information exists on the network and where it is and what format it's in, and if we have the right kind of computer and the right kind of software and the right level of access permission, we can have the information.

Simply stated, an Intranet is an application of Internet technology to private networks. This means that an intranet is based on the use of the TCP/IP protocol suite, and incorporates one or more TCP/IP applications such as e-mail FTP, Telnet, and Web client browsers and servers, An intranet can be represented by a LAN containing an FTP server or a private network consisting of several LANs linked together by WAN's with one or more LANs containing Web and/or Telnet service.

The application of Internet technology to the corporate network provides certain distinct advantages in terms of employee training and product development. Concerning training, the ability for employees to use a common set of applications for both Internet and Intranet applications can reduce training time and boost employee productivity. Concerning product development. Instead of having to master different applications for internal and external use, and employee can learn one application and use that application to satisfy internet and intranet information requests.

Because product development for the internet is attracting thousands of companies, the ability to use such applications on the corporate network enables the corporation to acquire the latest in technology rather than waiting for such technology to be ported to proprietary legacy systems, assuming they are ever ported.

Through the use of the concept of an Intranet, companies are rapidly adding applications that significantly enhance the ability of their employees to perform their daily operations. For example, many companies now operate Web server based help desks, enabling employees to use a familiar browser to access the corporate help desk, report problems, request assistance, or simply check the availability of a computer system for weekend access. As the Internet continues to gain popularity, we can expect more corporations, government agencies, and academia to convert private networks to the TCP/IP protocol suite, increasing the growth of private Intranets.

2.3 Intranets Vs. Groupware

Groupware is a broad category of software that has appeared in the last two or three years. It often includes facilities for functions such as sending and receiving e-mail, maintaining personal and group calendars, planning projects, and sharing documents.

Groupware packages also provide ways to customize their capabilities or create group applications. The insurance company in our example could use groupware to provide its employees with the capabilities we described.

But World Wide Web software has one big advantage over most groupware software: It's a whole lot less expensive.

For example, a Lotus Notes application can cost $250,000 or more to install in a network of significant size. A corporate intranet, on the other hand, can be implemented for less than $10,000 and serve the same number of users.

Lower costs are making intranets more accessible to organizations that cannot afford the groupware price tag. For example, a small distribution company with 150 users spread throughout the country justify an intranet implementation long before it could afford the equivalent groupware implementation.

Groupware will still have its market. Groupware has greater facilities for collaborative computing-applications that allow two or more people to work on a document simultaneously over the network – than intranets can currently offer.

Intranet developers are closing the gap between the two systems, however, Intranet browsers and servers are becoming more intelligent and capable, especially with the addition of programming facilities such as Microsoft's ActiveX and Netscape's java technologies.

On the groupware side, the Lotus Notes server software can now act as a Web server, making it a type of hybrid groupware / intranet system. It seems clear that the two technologies – groupware and intranets – will merge over time, each incorporating the best features of the other.

2.4 Intranet Hardware

An intranet is a client / server system. The server is a computer that has two main functions.

- It runs the intranet server software.
- It usually stores some or all of the content that is available to users.

For an intranet, the requirements for the server computer depend mainly on how many connections per hour it will be expected to handle.

A typically equipped modern personal computer – single Pentium 120 MHz processor, 32MB of random-access memory, 2GB hard disk – would be sufficient to handle a small intranet load. The load could be up to, say, a few hundred connections per hour. When the connection load on the server gets too high, users will find themselves waiting a long time to get a page. Eventually, their requests wil time out (the browser will get tired of waiting and give up) because the server can't get to their request fast enough.

When the server becomes overloaded, there are a number of things you can do to improve its performance.

- Add more memory to the server
- Add more processors to the server
- Add more servers and split the content among them

Almost any computer can be a client. All the client needs to do is run the World Wide Web browser software, which does not require a lot of processor power or memory. Clients can also be based on any technology, such as a Macintosh, Intel, or Sun, as long as the platform is capable of running TCP / IP.

- Although it is true that browser software is relatively "light" in terms of its demands on the client computer, that situation is changing. Newer browsers, such as Microsoft's

Internet Explorers 3.0 run to several megabytes in size, and they are happiest with a plentiful supply of RAM.

2.5 Intranet Software

The software that runs on the server computer and provides services such as World Wide Web (also known as HTTP) publishing, FTP file retrieval, and indexing and searching facilities is known as server software. Server software is available free on the Internet, and commercial packages such as Microsoft's Internet Information Server and Netscape Navigator are also available. Server software is available for a variety of operating systems, such as UNIX and Windows NT Server.

The software that runs on the client is called a browser. We use the browser to access and view Web pages. Because our intranet will use the same protocols as the Web, we can use any of the well-tested Web browsers already available – either free of charge or at a reasonably low cost on the commercial market. Perhaps the best-known Web browser is Netscape Navigator. Netscape provides the user with a variety of features, such as bookmarks, e-mail automation, and "hot" lists for storing the addresses of our favorite Web sites, Navigator is available commercially and, at most times, Netscape Communications Corporation allows downloading of a recent version.

The Internet is a gold mine of free software. We can find everything we need to run a full-features intranet. One word of caution: if you use software from the Internet, you may be on your own when it comes to getting it set up right and solving the inevitable problems that come up. Unless we're the adventure some soft and feel comfortable poking around in the innards of computer applications, we might be better off spending a few dollars for well-designed, tested, and supported software.

NAME	WEB ADDRESS
Midas WWW	http://www-midas.slac.stanford.edu/midasv 22 / introduction.html
Mosaic	ftp://ftp.ncsa.uiuc.edu/Mosaic/
Netscape	http://home.netscape.com
Emacs	ftp://moose.cs.indiana.edu/pub/elisp/w3/
Lynx	ftp://ftp2.cc.ukans.edu/pub/lynx/
Internet Explorer	http://www.microsoft.com

2.6 Intranet Services

All of Internet services are still available on the Internet. Chances are, however, we won't want or need all of them on our intranet.

Here's the suite of services and facilities you'll probably want:

- World Wide Web publishing
- FTP
- Some kind of indexing and searching service
- Some facility to run scripts.

Web (HTTP) Publishing

The World Wide Web is by far the most popular service on the Internet. In fact, for many people, the Web is the Internet.

The World Wide Web is based on HTTP, HTTP defines how browsers and servers communicate and move information back and forth.

- There's that word again: protocol. Now we have three layers of protocols: the HTTP uses the Internet Protocol over a network that has some kind of network communication protocol. One would think Internet was invented by the State Department.

- HTTP uses a "request and response" process. The browser on the client computer sends a request to the server for a particular page of information. The server receives the request, finds the requested page (file), and sends it to the browser. HTTP then forgets about the whole transaction. The browser and the server maintain a connection only long enough to process the request.

- Let's clear up some potential confusion here. As stated, the server and the client maintain a logical connection only long enough to process a request. However, they remain physically connected over the network. If we connect to our intranet remotely over a telephone line, the telephone connection remains open until you quit the session, even though the server and browser aren't talking to each other.

HTTP is usually used to send HTML files to the browser, but it can send any kind of file. If the browser cannot display the file, it will start another application (called a helper application) to display it, if possible. If it can't find as helper application, it will give the user option of saving the file to disk.

Hyper-Text Markup Language

The primary language used to create Web pages is HTML. HTML uses commends called tags embedded in a text file to tell the browser how to display the file.

HTML started out as fairly simple markup language but, like everything else in the computer world, it's getting more complex all the time. An official HTML standard is published by the World Wide Web Consortium, but companies that produce browser software like to add unique extensions to the set of tags their browsers understand, to give them a competitive edge.

For example, Microsoft's Internet Explorer lets we use a special tag in an HTML document to create marquee in which a line of text crawls across the page from one side to the other. A browser that does not support this tag will not display the marquee effect.

This example illustrates one of the irritating "gotchas" in the Web: All browsers do not support all the same extended HTML tags, so you can't count on them all to display the same page in the same way. For our intranet, it would be a good idea to standardize on one browser to be used by everyone, so we will always know what tags it can understand.

In the early days of the Web, creating a document in HTML format meant typing the content in a text editor and then adding HTML tags by hand. Nowadays, there are authoring tools, such as Navigator Gold, FrontPage, and bothers, that are very much like word processors. They let you create documents and add formatting and effects without ever having to touch an HTML tag. Instead, the tools add the tags behind the scenes as you create the document.

Hypertext

Hypertext – links that can take user to new page with a single click of the mouse button – is a defining feature of the World Wide Web. Hypertext is what gives we the capability to manage information on the Web in new, highly effective ways.
If we have used the World Wide Web, we have used hypertext. It shows up usually as colored text – a label, a section title, a word or two in a sentence, or even a portion of an image. When we move the mouse pointer over hypertext, the pointer changes to a pointing hand. When we click the hypertext with the mouse, something happens. Usually, we go to another page.

Communication Systems

E-Mail

Electronic mail or E-mail, as it is popularly called, is a system that allows a person or a group to electronically communicate with others through a network. A user sends an electronic message over a network; the message is stored in the electronic mailbox of the receiver. The electronic mailbox is usually a file on a server; the messages in it can be retrieved when the recipient is ready to receive them. Users can also edit, sort, save, and classify messages and forward them to other individuals on the network. If two users are logged onto the network at the same time, then they can converse through E-mail. Some E-mail systems have multimedia capabilities allowing E-mail users to send not only text, but also voice and still pictures.

There are many kinds of E-mail software; depending on the characteristics of each type used, mail can be sent between different computers or may be restricted to users on one computer. If an organization has different types of systems, say IBM, Digital, and so on, these systems must be linked together for E-mail to be successful.

E-mail has many benefits. First, it allows organizations to be responsive to customer needs. For example, when Freddie Mercury died, his fans around the world sent E-mail messages to his music company. EMI Records, to order his music. As a tribute to his memory, radio stations around the world began to play his once neglected songs. The box-office blockbuster film Wayne's World features one of Freddie's old hits and MTV began showing his music videos. The dead star was suddenly hot. EMI subsidiaries around the world were flooded with E-mail orders for the star's music. The Queen albums in EMI's catalog were particularly popular with the E-mail users. The result EMI sold 3 million Queen albums in 1 month, generating $30million to $15million of the year's profits for the $2-billion company.

The "big 3" credit reporting agencies – TRW, Equifax, and Trans Union, use E-mail to respond quickly to customers' concerns and queries. Before the system was installed, consumes who were trying to rectify errors in their credit reports often felt harassed by the credit agencies. Consumers were forced to write to all three agencies about any errors in their credit reports. Each agency had to write to retailers and lenders about the corrections. This paper-driven process was nightmarish for some consumers; there were enormous delays in rectifying errors. There were mail delays and internal processing delays in getting the paper forms to the right people, with the result that dispute resolution could take from 30 days to 8 months.

Today, all three agencies use E-mail systems to correspond with retailers and lenders. This not only saves the agencies time and money, but has greatly alleviated the problems that consumers face when there are errors in their credit records. Most disputes are resolved in under 5 days – a big improvement over the old system.

Another advantage of E-mail is that it provides instantaneous access to and dissemination of information, thereby eliminating the time lag involved in using the postal service. E-mail messages do not get lost or reach the wrong party.

In spite of its many benefits, E-mail should be used carefully. John Beamish, senior systems analyst at the Ministry of Education in Toronto, tells this E-mail story; "At my last place of employment, the unwritten rule was; If you drink the last of the coffee, make a fresh pot. Unfortunately, not many people did. In response, an angry coffee drinker sent a ragging five-page E-mail message regarding his dissatisfaction. The message was so long and went to so many people, it used all the memory of at least one dist, and no one could access their E-mail messages".

Finally, E-mail systems must be carefully selected, since a poor package may not only deter users from using E-mail but can also result in poor communications. Some factors to consider before selecting an E-mail package are: What hardware and operating system platforms must the mail system support? How many users do we currently have? How many users will we have in 2 years? In 5 years? Will remote users at multiple sites need to use the system? Must users access external services, such as Compuserve? Can messages be sent between different platforms (say, from a UNIX machine to a Macintosh or an IBM PC)? Is the package capable to meeting the growing needs of the organization? Will the vendor provide training and technical support if necessary? These and other related questions must be carefully addressed.

Fax:

Fax. or facsimile transmission, is another type of electronic publishing and processing system. Fax technology uses telephones, modems, and scanners to transmit text and graphics to individuals and organizations all over the world who have access to telephones. Note that a computer is not required to send a fax. all that is needed is a fax machine at either end of a telephone connection. The idea itself is very simple and elegant. A scanner in the fax machine scans the document at one end and a built-in modem sends it; in the fax machine at the other end, a build-in modem receives the message, a scanner scans the document, and a printer prints it.

In a short time, the fax, sometimes viewed as a "long-distance copier", has become part of the home office, the work office, and even of war zones! Journalists reporting the war in Bosnia used fax machines to transmit was stories because editors preferred to see the stories in hand-copy form. It is estimated that there are 24.5 million fax machines in the world today. The Japanese, in particular, rely heavily on fax communication because their language, with its thousands of ideographs, is easier to write in longhand than to type on cumbersome Japanese word processors.

Fax machines can send the same document to multiple users. They can also be integrated with applications such as word processors, so that faxes can be edited without being rekeyed into the computer. Fax machines can be programmed to send faxes when telephone rates are lowest. A portable fax machine allows users to retrieve documents using a touch-tone phone's keypad. An interesting development is the fax modem, which allows users to send or receive faxes using their PCs.

Voice Mail:

Voice Mail facilities oral communication. This way it works is illustrated in below figure, the sender dictates a message over the telephone. A special device, called a codec, converts the analog signal of the sender's voice into a digitized message. The message is transmitted over the network and stored in a server at the receiver's end. A blinking light on the receiver's hone indicates that he or she has a voice message. When the receiver chooses, the digitized message is retrieved from the server, reconverted into alalog form, using a codec at the receiver's end, and the receiver receives it over the phone.

Although voice mail is a simple idea, it eliminates a number of problems, such as "phone tag". It also ensures that a message reaches the right party even when that party is not available. Some voice mail systems can send the same message to several people, reroute a message to another phone, save messages for future reference, and retrieve messages from any telephone, anywhere in the world. In some systems, uses need a password to access voice messages, making voice mail a secure medium for confidential messages.

2.7 Extranet

Extranet are semi-private networks located outside the corporate firewall which provide selective internal information to business partners such as clients/customers and suppliers as well as to remote employees. Extranets support activities such as sales and marketing, supply chain management, customer/product support, employee support, and office equipment and supply procurement. In other words, an Extranet is a business-to-business intranet that allows limited, controlled, secure access between a company's intranet and designated, authenticated users from remote locations. An intranet that allows controlled access by authenticated parties.

In the last couple of years, there has been a blitz in the media with every technology company worth its name taking about intranets, extranets, and e-commerce and much more, leading to general confusion. Now, we will discuss relationship between intranets, extranets and e-commerce. First, extranets and e-commerce have in common the use of internet protocols to connect business users. Second, as mentioned above, intranets are more localized and can therefore

Move data faster than more distributed extranets. Third, the amount of control that network managers can exert over users is different for the three technologies.
On an Intranet, administrators can narrowly prescribe access and policy for a fixed group of users.

On a business-to-business Extranet (e-commerce), system architects at each of the participating companies must collaborate to ensure a common interface and consistent semantics. Since once company cannot reasonably enforce standards on its trading partners, extranet application developers must take into account a wider range of technologies than is the case for intranets.

2.8 Short Summary

An intranet is an internal, network-based infrastructure. In reality it is the same architecture as the Internet and has the same document and content standards as the World Wide Web.

Corporate networks that use the same networking/transport protocols and locally based web servers to provide access to vast amounts of corporate information in a cohesive fashion. Documents must be stored in HTML, and clients need web browser software.

2.9 Brain Storm

1. What is Intranet?
2. Explain about Intranet software and hardware.
3. What are Intranet services?
4. Explain about extranet.
5. What is groupware?
6. Differentiate the intranet and groupware services.

Lecture 3

Internet Server

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Internet server.
- Describe the web protocols.
- Describe about browser.
- Describe about the futures of Internet and Intranet application.

Coverage Plan

Lecture 3

3.1 Snap Shot

In this session we discuss about the types of server and web protocols. The Internet is a shared communication system, which means that many different communication sessions can take place over the same line at the same time. In contrast, a voice telephone call consists of a single circuit that only you and the person you called use, and the connection last only for the duration of the call.

3.2 Internet Server

The Internet is a rapidly global communication network of interconnected computers. Together, these millions of connected computers form a vast repository of hyperlinked information. The Internet is the result of standards that were originally defined by the U.S. government and that have been adopted by organizations around the world. It provides a system that lets computers anywhere communicate with one another, regardless of who manufactured them or what type of software they run.

The Internet is often called the "information superhighway", and for good reason. You can get an idea of the Internet structure by picturing the highway system in the United States. Major freeway systems cross the country, and these freeways intersect both with other freeways and with slower, secondary highways. The Internet is similar. Major carriers such as AT&T, Sprint, and MCI have constructed high-speed communication channels that cross the country. Branching from these trunks are secondary routes at major cities, and branching from those are other links to local service providers and businesses.

The foundation of the current network is the optical cables installed by major communication companies with the help of government agencies. Below figure illustrates the trunk lines of the very high-speed Backbone Network Service, which is a product of the National Science Foundation and major providers, including MCI, Ameritech, Sprint, Pacific Bell, and several other companies. vBNS uses MCI's nationwide network of advanced switching and fiber-optic transmission facilities that can transmit voice, data, and video signals at speeds of up to 600 million bits per second (Mbps) – more than 20,000 times faster than a typical home modem connection of 28,800 bits per second (bps).

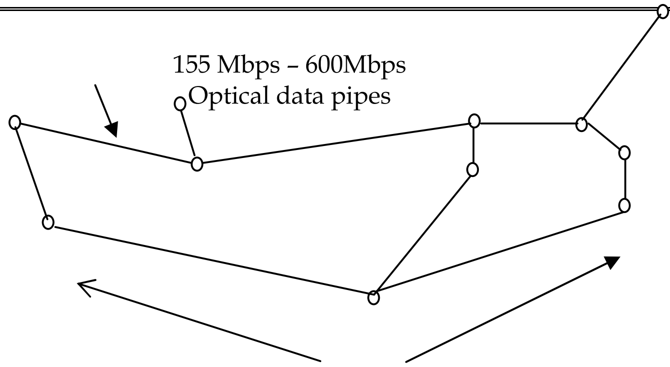

Network access points (NAPs)

The backbone networks that cross the country are like big data pipes, with the size of the pipes, with the size of the pipe measured in terms of bandwidth. The larger the bandwidth, the more information you can push through the pipe. Millions of digital messages funnel through these electronic pipes at any one time. Smaller pipes branch from the big pipes at locations called network access points (NAPs), and even smaller pipes branch from these. For example, Pacific Bell connects to the vBNS, at a NAP located in San Francisco. Internet service providers around the state of California connect directly to this NAP to gain access to the backbone as pictured in following figure. ISPs pay Pacific Bell for large chunks of bandwidth, which they resell to their customers for a monthly fee.

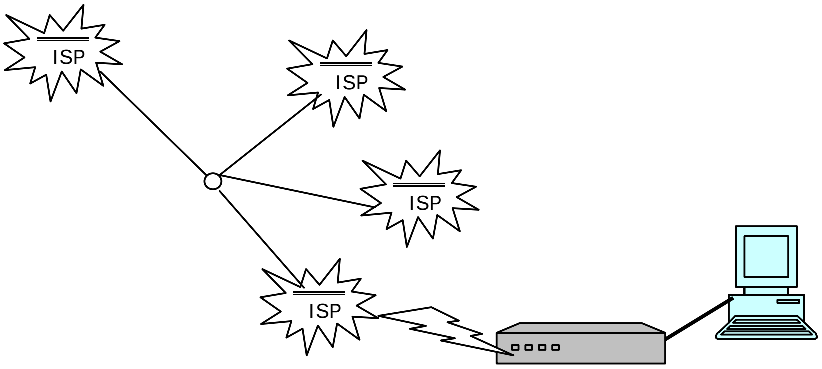

Customers are the end users of services and connect directly to the ISPs – the "on-ramp" for the information superhighway. A typical home user connects to an ISP with a dial-up modem, which doesn't make a connection to the ISP until the user initiates the telephone call. Alternatively, many larger companies have high-bandwidth dedicated leased lines that are always connected to the Internet. These companies' networks are directly attached to the leased line, and all the users on these networks can access the Internet through them.

The monthly fee paid by end users and the large fees paid by corporate users and universities pay for the Internet's infrastructure. Ultimately, your $20 to $30 monthly fee combines with everybody else's to support a communication system that lets you connect to any computer in the world almost instantaneously. In this system, there's no additional charge for long distance. Whether you talk to a computer next door or to one halfway around the globe, your monthly fee remains the same.

The Internet is a shared communication system, which means that many different communication sessions can take place over the same line at the same time. In contrast, a voice telephone call consists of a single circuit that only you and the person you called use, and the connection last only for the duration of the call. On the Internet, information is divided into chunks and placed into packets. Packets are addressed to a destination computer and sent out over the Internet. Devices called routers ensure that packets are sent along the best paths to reach their destinations.

Again, the best analogy for this is the freeway system. People get into cars and drive to their destinations. Each car is an independent delivery system. The driver can same time, each going to a separate destination. At major interchanges, cars can switch paths to get on a more direct routs to their destinations. As with a freeway system, the interchanges on the Internet are the places where traffic congestion is likely to occur.

A typical communication session might involve hundreds or thousands of individual packets of information travelling between two computers. Information is divided into packets by the sending computer and reassembled by the receiving computer. One advantage to using packets is that a small glitch on the network usually affects only one or a few packets. The receiver will know that some packets are missing and can request retransmission from the sender.

The packet nature of the Internet enables data traffic to be sent around problem areas or over multiple paths. Packets are sent over the mesh of network connections as independent objects that eventually reach their destination. Like posted letters, packets hold information and have source and destination addresses.

The Transmission Control Protocol / Internet Protocol (TCP/IP) is a set rules that define communications and packet handling. Specifically, TCP /IP defines

- How much data goes in a packet
- How to address a packet
- How to transmit packets over the network
- How to route packets around failed links
- How to detect errors or missing packets and get retransmission

TCP / IP has become the grand unifier for computers, not only in the corporate environment, but in all environments throughout the world. It didn't start out that way, though.

The U.S. Defense Advanced Research Projects Agency (DARPA) originally designed TCP/IP as a protocol for building a communication system that would ensure continuous connections during emergencies -–even during nuclear wars. Should a link fail in the mesh of inteconenctions, packets would simply be diverted around the problem and travel a different path. A testament to the reliability of this system, which used the same protocols, during the Persian Gulf war.

The Internet grew as it interconnected other government agencies as well as universities and research organizations involved in government work. Today, commercial enterprises have adopted the Internet as the means to connect with other organizations. People around the world now connect to the Internet to access a variety of services, some of which are listed here:

- Electronic mail, which provides an easy way of transferring messages and documents between systems.

- Internet discussion groups about politics, hobbies, computers, or virtually any other topic imaginable.

- Libraries of information on computers located at public and private institutions.

- The World Wide Web, perhaps the most exciting part of the internet.

Types of Internet Servers

The Internet not only provides an infrastructure for communications, but is also an information system in which servers provide information to Internet users. Web servers, the fastest growing segment of the Internet. The "traditional" information servers on the Internet – FTP, Gopher and WAIS – are not as dynamic and interesting as Web servers, Web servers introduced graphics and hypertext links, a way of navigating the Internet by pointing and clicking a mouse.

- FTP servers File Transfer Protocol servers provide users with a way to connect with an Internet server, browse its directory structure, and transfer files. Basically, an FTP server lets remote users execute text-based file system commands on a server.

- Gopher servers Gopher servers provide information in an easy-to-navigate, menu-driven format. Users select categories of information from a menu, which might open another menu. Eventually, users find what they are looking for by navigating the menus, assuming that the content provider has created a usable interface.'

- WAIS servers Wide Area Information Server servers are information servers that provide indexing services. Documents that are fed to the server are fully indexed on keywords and include unique searching tools for locating matching keywords or documents that are similar in content.

- Web servers Web servers provide hyperlinked information in a graphical format, as described in the following sections.

What is the Web?

The Web is built on top of the Internet and uses TCP / IP to transport information from place to place. Users run software browsers such as Microsoft's Internet Explorer or Netscape Navigator to connect with Web servers and FTP and Gopher sites or to send and read electronic mail. Browsers eliminate the need for users to understand arcane UNIX commands and replace the mystery of the Internet with a graphical user interface program that displays text and images and plays back sound and video as well. You can set up a Web server for your own in-house use or to provide information to Internet users outside your organization.

Below figure shows Microsoft's Web browser and the home page for Microsoft's site. Note that MS Internet address of http://www. microsoft.com / was typed in the browser's Address field. Most of the icons and graphics you see on the page are hypertext links. If you click a link, you "jump" to other locations or Web sites. You can tell whether a graphic is a hypertext link by moving the mouse pointer over it. If the pointer changes to a hand with a pointing finger, the graphic is a hypertext link.

Users get access to the Internet through local service providers or by connecting through their organization's network. Once an Internet connection is made, you can run a browser and start jumping from Web site to Web site. Some people call this "surfing the Web". A Web serer automatically displays a home page when a client connects to it. HTTP, which is discussed next, takes care of all the details in the background and handles hyperlinks to other documents or other sites.

3.3 Web Protocols

Technically, a Web server runs under the Hypertext Transfer Protocol, an Internet protocol that supports hyperlinks in documents. The server displays information on Web pages that can contain hyperlinks to other places. Those "other places" can be on the same page, another page on the same server, or an entirely different server at another site. In this respect, the Web is like a vast network of interlinked documents - it is essentially one huge interlinked information publishing system, a sort of universal repository of information that you can access with mouse clicks. It is far better than any library because you can access it from your home or office and instantly hyperlink to new information.

The best thing about hyperlinks is that users don't need to remember pathnames, document titles, and other resource properties to access information on the Web. Other Web sites and resources are usually just a click away and usually referenced in documents you read. Another great feature is that hyperlinks let you explore and do research in an intuitive way. You can follow just the path you want through a maze of documents to locate the information you need. If hyperlinks don't get you to the information you need, you can always try one of the great Internet search sites such as Yahoo. (http://www.yahoo.com/) or Alta Vista (http://www.altavista.digital.com/).

The content of Web pages is written with a special formatting language called Hypertext Markup Language. Web page authors use HTML and some associated add-ons to create appealing content. All Web browsers know how to interpret and display HTML, so it can be considered a sort of universal document formatting language. It is easy to "code" HTML documents, and once the documents are complete, they can be placed on most Web servers, not just the server or operating system on which they were created. You can even create HTML documents in word processors such as Microsoft Word.

Web servers really do most of the work on the Web. They pull information from local storage, such as text and graphics for a page, and funnel that information out over the Internet. The browser just displays things. That is the core activity of the Web's client-server relationship. However, new techniques are under development that are making pages more active and that are shifting some of the work to the client, as discussed in a moment.

The Client-Server Model

The Web is based on the client-server model. Users run Web browsers to access information on Web servers. Web servers "serve up" information and services. Below the figure illustrates how clients and servers are interconnected over a data network, which might be an in-house network, the Internet, or both.

The client-server model has been around for a while. In corporate environments, it is used to design networks and applications that distribute the processing load between a user's computer and a server. Client-server systems take advantage of the intelligence resident in both the client and the server and also take advantage of the network's communication systems.

For example, the client formulates a request for some information or service and sends it to a server. The server then cranks away to provide that information or service and returns something to the client. The client performs some processing of its own in this scheme, such as building server requests and displaying and formatting the results returned from the server.

In the traditional client-server relationship on the Web, the server does most of the work. When you connect with a Web site, the server displays an HTML page-a relatively easy task. However, if you fill out forms to request information, the server may need to do a lot of extra work, which can be a strain if a lot of other people are accessing the site. For example, the server might need to formulate your request into a query for a database server, wait for the results, then build an HTML page on the spot and send it to your browser. Meanwhile, your browser has done nothing except wait for the response.

This scheme is changing as Netscape, Sun, Microsoft, NeXT, and other companies create languages and development environments to make the Web more active by pushing more processing to the client side. When you log on to a site that uses these tools, small programs, or applets, are downloaded to run on your computer. That means your computer is no longer idle, and the server has freed itself from some of the processing so it can get on with other tasks. With the client running part of the program, content developers can create more exciting and active Web pages.

Some people fear that all of these changes and upgrades will fragment the Web, that browsers will lose their universality as Web sites come online that a particular browser can't talk to . Microsoft has countered this by offering programs and product development tools that are open to all new standards.

An interesting development is the ability to provide automatic browser updates by downloading just the component you need to keep your browser current. In other words, if you need a special program to view some new type of video format, that program is automatically downloaded to your computer. You don't even need to ask for it in most cases. The program stays on your computer waiting to be used again. If this plan works, the Web will become an automatically updated software platform!

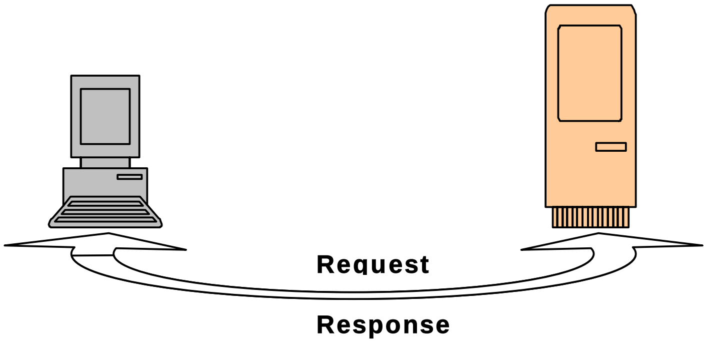

Client / Server Network Connections

However, the really interesting thing is that this same technology will probably spread to other environments. For example, there are rumors that Microsoft intends to of a button you'll be able to get online, start an automated check of your operating system, and then have

Windows automatically download and install the required components for you. This is a far cry from the days when you had to order update disks from vendors and wait for them to arrive in the mail. If this particular experiment is successful, you can count on seeing more vendors offering automated updates of their software in the future.

3.4 Browsers

Web-Browser: The Universal Client

In 1994 and 1995, much of the attention of people interested in the Web was on browsers. There was a lot of competition to make one browser better than another. The Netscape Navigator browser got the most attention because Netscape added some interesting features, in some cases without waiting for the normal standardization process. Because Netscape had somewhat of a lead in the market, it could make changes to standards without much fear of upsetting users or other vendors, who simply followed along and adopted the updates.

The first popular browser, called Mosaic, was designed by the National Center for Super Computing Applications (NCSA) at the University of Illionois. NCSA Mosaic is public-domain software, although people have paid for it (and other browsers) by buying them with add-on promotional products such as Internet. Here are a few advantages to using a browser:

- Browsers are typically free or very inexpensive. Windows 95/98 and Windows NT include a browser.

- Browsers provide an almost universal interface for accessing and displaying information. Everyone uses the same or a similar interface, even people outside your organization and in different countries.

- Browsers can connect with any Web server, no matter what operating system or platform.

- Browsers require few system resources and little if any maintenance and system configuration.

- Browsers can easily be updated by downloading the latest version from the Web page of the browser’s developer or vendor.

Today, browsers are available for all major computer platforms and operating systems. That means that anybody can log on to the Internet and view the contents of any Web site. Think about that for a minute. After years of incompatible standards, the computer industry finally has a tool that anybody can use to view information on any other computer. That means you can set up your own Web site and publish information without regard for the equipment and operating systems that people use. Well, almost.

Recently, Netscape, Microsoft, and other vendors have been waging browser feature wars, adding new features and functions in an attempt to differentiate their products. This loosens Web standards and makes it difficult for Webmasters to design pages that will support every potential client. Webmasters need to be sensitive to the different types of browsers, perhaps writing duplicate pages formatted for each browser, or just posting a tag that says, "This site requires Netscape Navigator!"

Microsoft has said that it will "embrace and extend" its Internet browser to match new browser standards. That strategy will probably make it impossible for anyone to make any money selling browsers, so eventually the interface may stabilize. Over the long run, Microsoft's strategy of supporting all standards and environments will probably force other browser developers to do the same thing.

Keep in mind that the concept of a stand-alone browser may be ending. Vendors are starting to include browser features directly in their application. After all, why should you leave your word processor just to get a document from the Web? Why not retrieve it using the same pull-down menu options that you use to search for and access files on your local computer or networks? Microsoft is already building Internet support into all of its applications. You can count on seeing an Internet interface in Windows 98 as well.

Bill Gates made an interesting statement about this in a recent speech. He said that "there is a funny inversion going on in which it's harder to find documents on the LAN than it is on the Internet. The searching tools are better on the Internet and there are links between documents. Documents created in-house do not have these features, but they can if Web servers are deployed in-house." Microsoft is putting hooks into its Office products so that individual users can publish documents and search for documents within their own organizations. This will make it important to use common, standardized document formats for in-house publications so users can search and retrieve documents just as easily as they do on the Net.

The New Enterprise Computing Model

During the 1980s and early 1990s, organizations began to install local area networks to connect computers in departments and workgroups. Unfortunately, department-level managers usually made the decision about what type of computers and networks they wanted to install. The marketing department might choose Macintosh computers, the sales department might choose PCs, and the engineering department might choose UNIX systems.

Then people started thinking about joining those systems together so people could exchange e-mail and work in collaboration with other people in the company. Companies began considering enterprise computing, client-server applications, distributed computing, middleware, open systems, and a lot of other strategies to get all of their dissimilar systems to connect and communicate together.

Industry organizations were formed to create open standards; each vendor developed its own strategy that it thought the rest of the industry might accept as a standard. The idea most people had in mind was to use existing networks as a sort of plug-and-play platform to which any computing device could be attached so that it could communicate with any other device.

This so-called enterprise computing strategy began to take root as more and more organizations adopted TCP/IP. When computer ran TCP/IP and were connected to the same network, they could exchange files and use shared resources such as printers. However, the diversity of operating systems, applications, and data formats still restricted the free flow of information among computers. People wanted to collaborate without the need for translating, reformatting, and recompiling their programs and data.

In 1994 people started to notice that the Web was providing the heterogeneous environment that people wanted all along, and that it was practically free. You could use a single interface to access information on any computer. In fact, when you access a server on the Web, you don’t need to know what operating system it runs or what type of system it is. It just serves you the information.

By the end of 1995, it was clear to many in the industry that at least until something better came along, setting up internal Web servers might be a good way to disseminate information throughout an organization. After all, almost everybody accessing the Internet had a graphic Web browser. Why not use it to access information on back-end database systems and IBM mainframe computer systems?

This new internal Web site strategy really brought a lot of things into focus. Why write several different versions of a program to access database servers and mainframe systems when you can write one application for a Web server that Macintosh, PC, and UNIX users can access? New Web development tools that are coming online benefit developers, Web site administrators, and users. They merge traditional programming languages and document processing languages, facilitating the development of applications with custom user interfaces. Getting information from database servers and legacy systems is easier because you only need to write a link between the Web server and the back-end system.

At any rate, Microsoft realized what was happening and joined the party. If organizations were going to set up internal Web sites, why not write applications that would take advantage of those sites? In late 1995, Microsoft announced that all its major applications, including the full Office suite, will come already packaged to function on the Internet, or will be Internet-enabled.

Building Public and Private Web Servers

There are basically two types of Web servers: public and private. You put up a public Web server to provide information to the world of Internet users. Anyone with a Web browser can attach to a public Web site. You build a private Web server as part of an intranet strategy to provide information only to authorized users within an organization or to authorized users outside the organization. A firewall is a server or router that puts restrictions on who can access a server or gain access to an internal network. If you have an internal LAN connected to a public Web server, you need a firewall to keep untrusted external Internet users out of your internal LAN.

Following Figure illustrates public and private Web server connections. Note that the organization has its internal LAN connected to both the public and private Web servers, and that the public Web server is available to Internet users. The firewall might be part of the Web server or a stand-alone device. The public server provides three important features:

♠ Information to external Internet users.

- A firewall to prevent external users from accessing systems on the internal LAN
- A gateway that lets internal LAN users access the Internet through single secure connection.

Note that internal users can access either the public or the private Web server. For example, a systems administrator might need to manage the public server, or a marketing person might need to update promotional material on Web pages. The private Web server is available to any authorized user on the internal network.

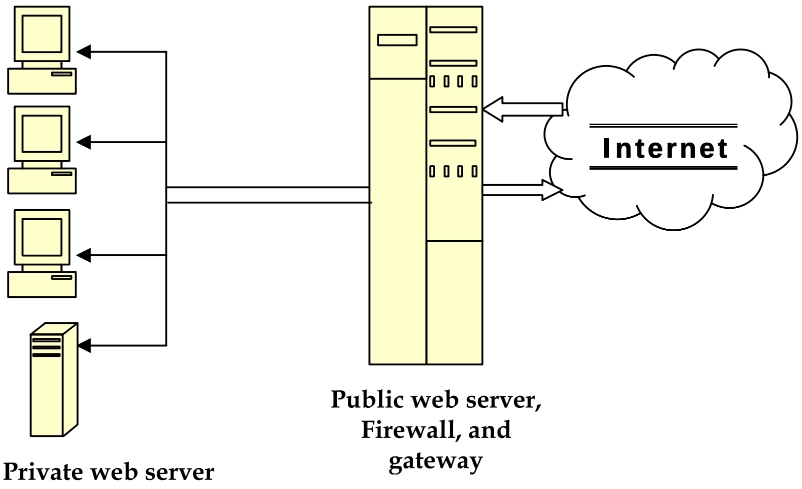

Private and public server configuration

There are currently more Web servers set up for internal use than there are public Web servers. More and more organizations are discovering the advantages of Web servers for internal communications and the distribution of information. These organizations build both local and wide area networks within the organization that mimic the Internet itself. TCP / IP provides the basis for these networks.

Public Web Servers

You can set up your own public Web site to accomplish numerous different goals. The Web is a completely new publishing system that lets anyone publish anything he or she wants. In the future, everyone may have a personal home page. Currently, there are few rules or restrictions on what you can and cannot put on your page, although recent legislation has attempted to outlaw “offensive” material on the Internet. How this will actually affect the

Internet will be decided by the courts over the next several years. Nonetheless, as more and more people access and publish on the Web, even more people will join and become a part of the phenomenon.

The next few sections describe some typical uses for public Web servers.

A Personal Home Page

If you have a full-time connection to the Internet, or if you want to serve up information to other people connected to your in-house network, you can set up your own Web server by following the instructions in this book. What do you post at the site? How about pictures of the last family reunion so relatives can display them up on their monitors. Or you could set up a site that pays homage to your favorite singer or author or sport. It's your call.

A Public Service Site

With proper funding from various community services and free connect time from a local service provider, you can set up and manage a Web server that provides useful information to the entire community – and rack up some experience toward starting your own Web server business.

A Business Site

A business site is designed to promote your products and services, so you need to be particularly concerned about the quality of content and the number of browsers your site can handle at once. In addition, many business sites now offer technical support, which means that the site has to be easy to use.

A Commerce Server

A commerce server is designed to handle-to-handle electronic business transactions for Internet users. Every transaction must be properly secured to prevent eavesdroppers and hackers from obtaining sensitive information such as customer information and credit card numbers. A commerce server is typically integrated with a database application such as Microsoft's SQL server and uses forms to collect information from users. Once users are authorized, they can access information in the database.

If you don't have the budget for a server and a full-time Internet connection, you can always lease space at another Web site or from a service provider. If you can afford a server but not a full-time Web connection, talk to your service provider about putting your at its site. Any of these methods will give you a start testing your ideas for a Web site or just experimenting. Once things start to move, you can expand your operations.

As you saw earlier figure ISP, your Web site will most likely connect through a dedicated line to an Internet service provider. The connection from your site to the ISP can vary in bandwidth, starting with a suggested minimum throughput of 28,800 bits per second and going all the way up to the megabit-per-second range. The size of the data pipe you choose depends on the number of people you expect to visit your site. For example, you could start with a simple 28.8 modem dial-up connection to the service provider. Such a line won't support a lot of users, but it's a good way to get a low-cost start. Make sure you also investigate possibilities such as cable modems, which can transmit data at least four times faster than dial-up connections and cost only marginally more. For example, a typical dial-up connection costs $19.95 per month for unlimited service, while a cable connection usually starts at $29.95 per month – a mere $10 per month more. Of course, you will have to buy a cable modem to use a cable connection to the Internet.

Private Web Servers (Intranets)

Internal Web Servers can be made available to an entire organization, to individual departments, or to specific department workgroups. You might use an internal Web server to

M Post company information such as employment policies, regulations, and events

M Host employee discussion groups

M Post forms, templates, disclaimers, and other information for creating public documents.

M Post the company phone book and employee directory

M Host a policy and orientation center for new employees

M Serve as a front-end to an organization library or information database.

M Post information about the performance of certain groups, divisions, or products in the organization

For example, an international company could post information that its employees can use to write business plans and contracts or to conduct other types of business around the world. Published material might include marketing, inventory, sales, economic, accounting, and other types of information. A real estate organization could post information that helps brokers and agents keep in touch with new and existing listings.

One of the most important uses for an internal Web server is as a place for users and departments to post their own Web documents. It is becoming easier for people to create Web pages, and an internal Web server provides space to post those pages. For example, a finance officer could post a form that helps employees calculate the value of their pension plans, the personnel department could post new job openings of interest to employees, or the marketing department could post information about the sales of new products. A Web page can even be used as a suggestion box.

Intranets and Virtual Private Networks (VPNs)

Keep in mind that not everyone who needs to access an internal Web server will be attached to the same LAN as the Web server. You can create directories on your Web server that require username and password access. Then only employees, business partners, or other designated people can access the pages by going through a normal Internet connection from anywhere in the world. This setup is like a private network on the Internet.

These private networks on the Internet are sometimes called virtual private networks(VPNs). With a VPN you basically build a network for your company over the existing Internet structure. Rather than pay for your own private leased telephone lines to connect every one of your business sites, you simply connect users and sites over the Internet.

VPNs require data encryption to ensure that your private company transmissions are protected. One method of achieving this is to use Microsoft Web browsers and servers, which together provide a way to encrypt transmissions between clients and servers, as discussed in following chapter.

Microsoft, with offices worldwide and a WAN communication bill of over $10 million per year, expects to reduce that bill significantly by running its network over the Internet. It plans to get rid of the bulk of its network infrastructure and use the Internet as the basis of its wide area network.

Collaboration and Groupware

Web servers can provide a central place where people working on the same project from different departments or locations can post and share important information. A good example of collaborative software is DEC's Workgroup Web Forum, which provides a place for teams to collaborate across the Web or across internal corporate TCP / IP LANs.

Workgroup Web Forum combines team collaboration with Internet technology. It provides file sharing and online conferencing and allows teams of users to easily upload, organize, and distribute information in a variety of formats. Users create descriptive folders and document names that other item members can recognize and access. Those folders and documents are connected through hyperlinks that span local server-based repositories, as well as the World Wide Web. Workgroup Web Forum automatically manages the hyperlinks, allowing quick and easy association of distributed information. Some Workgroup Web Forum features and benefits are listed here:

- Administrative tools provide for secure management of content and access control lists, ensuring that only authorized users have access to files and documents.

- Online conferencing allows users to conduct and participate in multiple electronic conferences.

- Real-time polling gives discussion groups a mechanism for reaching consensus on key issues. Polling results can be displayed graphically.

- Keyboard and content-based retrieval capabilities help users find information quickly. A "what's new" feature checks for recently generated material.

- Support for heterogeneous environments is accessible by any PC, Macintosh, or UNIX workstation with a Web browser – no specialized client software is needed.

3.5 Future of Internet & Intranet Application

The Internet faces a number of challenges on an ongoing basis, including the following:

- ❖ Increasing demand for backbone bandwidth
- ❖ Increasing demand for access bandwidth
- ❖ IP-related enhancements; for example,

  M IP-related issues, such as address limitations and type-of-service distinctions, and the need for newer versions of the protocol..

  M Support of multicasting over IP.

  M Support of multimedia traffic, using existing protocols such as IP and TCP. The issues are that multimedia is real time, requiring a lot of bandwidth, and the applications expect a well-defined quality of service. On the other hand, IP is a connectionless protocol, and TCP uses window-based flow control.

  M Routing issues, relating to the growth of routing tables, routing storms, and the need to use link-state protocols such as OSPF.

- ❖ The development of secure e-cash mechanisms.

New backbone connection schemes

The creation of the NSFNet backbone was one of the most important events in the recent history of the Internet. It enabled the growth of the Internet that continues to this day. Having succeeded so well, the NFS has recently gotten out of the business of providing the backbone for the U.S. information infrastructure.

NFS has devised a new scheme, under which a multiplicity of national Internet backbones would be deployed. In a sense, the NSFNet itself would retreat to its original purpose of serving the needs of scientists using supercomputer applications. The new supercomputer backbone, to be called vBNS (very high speed Backbone Network Service), would operate at a significantly higher speed than the current backbone. Parallel to this network would be a variety of commercial and government backbones, operating in an environment of competition as well as cooperation. The NSF plan calls for NSF to provide funding for regional networks, leaving it up to them to connect to national backbone services. Over a five-year period since the project award, NSF has planned to reduce those subsidies to zero. NSF called for the creation of Network Access Points (NAPs), which allow interconnection of regional networks, the vBNS, and foreign networks.

Options for connection LANs to Internet service providers will also change over the next few years. Alternatives to traditional leased lines are finally coming into their own. ISDN, talked about for many years, is finally being deployed in the United States as an access technology Asymmetric Digital Subscriber Line(ADSL) is talked about, as is Internet over cable TV. Other alternatives to leased lines, such as frame relay, SMDS, and ATM, are also being deployed, as discussed next.

The use of Asynchronous Transfer Mode.

ATM will become an increasingly important technology, not only in corporate environments, but also within the realm of the Internet. Graphics- and video-intensive applications necessitate higher speeds. By current standards, *high-speed* refers to networks that operate at 155 Mbps; ATM hardware can switch data at gigabit speeds; protocols for standardized access to these speeds are expected in the next few years. An ATM network consists of one or more high speed switches that each connect to host computers and other ATM switches. A typical ATM switch can connect between 16 and 64 ATM ready devices. ATM uses optical fibers for connections, including connections from a host computer to an ATM switch. Optical fibers provide a higher transfer rate than a copper wire. Typically, the connection between a host and an ATM switch operates at 100 or 155 Mbps.

Although a single ATM switch has finite capacity, switches can be interconnected to form a larger network. The connection between two switches differs slightly from the connection between a host computer and a switch, from a protocol point of view. Interswitch connection can operate at higher speeds and can use slightly modified protocols.

Below Figure shows the topology and shows the difference between network node interface (NNI) and a user-to-network interface (UNI). ATM is being contemplated for access as well as for backbone networks.

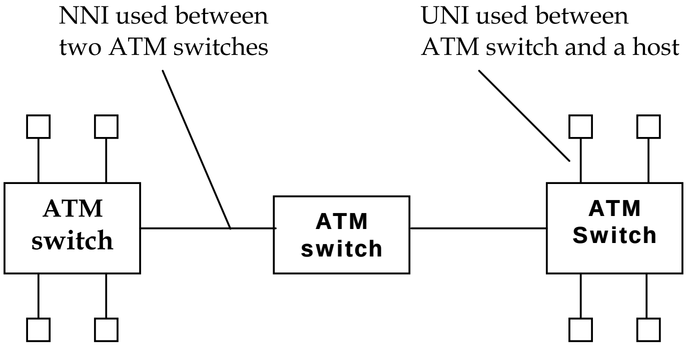
NNI used between two ATM switches
UNI used between ATM switch and a host
ATM switch
ATM switch
ATM Switch

Connection-oriented networking

ATM offers connection-oriented service. Before a host computer connected to an ATM can send cells, the host must first interact with the switch to specify a destination. The interaction is analogous to placing a telephone call. The host specifies the remote computer's ITU-T E.164 address and waits for the ATM switch to contact the remote system and establish a path. When a connection succeeds, the local ATM switch chooses a virtual identifies for the connection and passes the connection identifier to the host along with a message that informs the host of success. The hose uses along with a message that informs the host of success. The host uses the connection identifies when sending or receiving cells. When it finishes using a connection, the host again communicates with the ATM switch to request that the connection be broken. The switch disconnects the two computers. LAN emulation provides a mechanism for interworking connection-oriented ATM with connectionless LAN protocols.

ATM adaptation layers. Although ATM switches small cells at the lowest level, application programs that transfer data over an ATM do not read or write cells. Instead a computer interacts with ATM through an ATM adaptation layer, which is part of the ATM standard. The adaptation layer performs several functions, including PDU segmentation and reassembly, clock information transfer, and detection of errors such as lost or corrupted frames. Firmware that implements an ATM adaptation layer is located on a host interface along with the hardware and firmware that provide cell transmitting and receiving. Layers provide data transfer services for computer that use ATM. When a virtual circuit is created, both ends of the circuit must agree on which adaptation protocol will be used.

ATM adaptation layer 5 (AAL5) is used to send conventional data packets across an ATM network. Although ATM uses cells at the lowest level, AAL5 presents an interface that accepts and delivers large, variable-length packets. In particular, AAL5 allows each packet to contain between 1 and 65,535 bytes of data.

Datagram encapsulation and IP MTU size. A computer can use ATM adaptation AAL5 to transfer an IP Datagram over an ATM connection. On the sending host, AAL5 generates a trailer, segments the Datagram into cells, and sends each cell over the circuit. On the receiving host, AAL5 reassembles the cells to reproduce the original Datagram, strips off the trailer, and delivers the Datagram to the receiving host or the IP. ALL5 uses a bit in the header of the cell itself to mark the final cell of a given Datagram.

AAL5 uses a 16-bit length field, making it possible to send 65 KB in a single AAL PDU. Despite the capabilities of AAL5, TCP/IP restricts the size of datagrams that can be sent over ATM. The standards impose a limit of 9180 bytes per datagram. As with any network interface, when an outgoing datagram is larger than the network MTU, IP fragments the datagram and passes each fragment to AAL5. Thus, for these applications, AAL5 accepts, transfers, and delivers datagram of 9180 bytes or less.

IP address binding in an ATM network. As in other network technologies, ATM assigns to each attached computer a physical address that must be used when establishing a virtual circuit. On one hand, because an ATM physical address cannot be encoded within an IP address. Thus, IP cannot use static address binding for ATM networks. On the other hand, ATM hardware does not yet generally support broadcast. Thus, IP cannot use conventional ARP to bind addresses on ATM networks. Software on the host may not know the IP address or the ATM hardware address on the remote endpoint. Thus, an IP address-binding mechanism must provide for the identification of a remote computer connected over a virtual circuit.

TCP/IP allows a subset of computers attached to an ATM network to operate as an independent LAN using the Classical IP Over ATM model. Such a group is called a logical IP subnet (LIS) in RFC 1577. Computers in an LIS belong to a single IP subnetwork. A computer in an LIS can communicate directly with any other computer in the same LIS, but is required to use a router when communicating with a computer in another LIS. When a host creates a virtual circuit to a computer in its LIS, the host must specify an ATM hardware address for the destination. How can a host map a next-hop address in an appropriate ATM hardware address? The host cannot broadcast a request to all computers in the LIS because ATM does not offer hardware broadcast. Instead, it contacts a server to obtain the mapping.

Communication between the host and the server uses ATM Address Resolution Protocol (ATMARP).

Improved access approaches

On the dial-up side, ISDN promises a two-to threefold improvement in speed, but that will soon be swamped by requirements. A technology called asymmetrical digital subscriber line, which can be viewed as a mix between a private line and primary rate ISDN and supports up to 6 Mbps

For corporate-headquarters applications, dedicated access using ISP-provided routers connected over dedicated DSI (1.544 Mbps) or DS3 (45Mbps) lines can be utilized. However, the issue of how to handle telecommuting employees remains.

A replacement for IP

The number of networks on the Internet doubles approximately every year. With this explosion of growth, there is an increased concern about the availability of IP addresses. The Internet Engineering Steering Group has been working on a new version of IP, called IPng (also called IPv6). The next-generation Internet Protocol will increase the address size from 32 bits to 128 bits, support authentication and security capabilities, and feature quality-of-service labeling of packets to distinguish real-time data from e-mail.

The “IP” part of TCP/IP was designed under the assumption that perhaps a few thousand hosts would be attached to the network. With over several million hosts now IP-accessible, IP is approaching its design limit. One problem IP faces is the consumption of address space. Due to the way IP addresses are segmented into classes of networks, much of the address space is “waster.” In particular, addresses for Class B networks are largely used up. The Internet community has adopted a scheme called Classless Internet Domain Routing (CIRD) that will preserve addresses by abandoning the old class rules. The specifications for CIDR can be found in RFC 1518-1520. CIDR works on the underlying fact that contiguous blocks of IP addresses share the same most significant bits. CIDR uses these summarized contiguous blocks, called supernets, to enable route table aggregation. CIRD is expected to provide relief until the late 1990s, when a new scheme is required.

The IETF is working on the next-generation IP suite of protocols. In its current draft state, the next generation of IP is known as IP version 6 (IPv6). The version of IP now in use is known as IPv4. At press time, the IPv6 specifications were in draft form only. While the overall architecture is unlikely to change significantly, this can theoretically occur until the specification is released.

The IETF chose IPv6 from several competing proposals that were discussed and debated during the early 1990s. The overall goal of IPv6 is to create an architectural framework that enables the Internet to grow into a system with millions of interconnected networks. Another equally important goal is to allow a gradual migration to IPv6 from IPv4 with minimal disruption to existing systems. Subgoals within this framework are portable computing and dynamic IP address discovery and assignment. The major competitor was called TCP and UDP with Bigger Addresses (TUBA), which embraces an OSI standard called CLNP. Some engineers believe that a successor to IP needs to be designed to handle the ultimate load imaginable – let's say, 10 devices for every person on the planet, with room for growth in population as well as devices.

Support of electronic commerce

Many now claim that the "future belongs to electronic commerce." Effective methods of payment over the Internet have to be developed for this activity to take off. The issues of concern here have been security and privacy. A number of technologies are becoming available to improve both. Two examples are covered here.

One example of an approach that can be used is based on CyberCash Incorporated's technology for cost-effective, convenient, and rapid payments over the Internet. Its system is designed to facilitate Internet commerce by enabling financial transactions between individuals, businesses, and financial institutions. The solution is based on establishing a trusted link between the Internet and the traditional banking world. Three services are available: credit card, electronic check, and electronic coin services, which are the Internet counterparts to credit cards, checks, and low-denomination cash payments. The CyberCash Wallet software is distributed to individuals free of charge through private-label arrangements with on-line service providers and financial institutions. It may be directly downloaded from Web sites of the company (www.cybercash.com) and participating merchants. CyberCash provides merchants with free server software and leverages a financial institution's existing infrastructure through its gateway server software. Recently an Internet-based entertainment service based on pay-as-play has come on-line. The DigiCash electronic coin service is employed to enable secure payments of as little as $0.95.

Another example is provided by DigiCash (www.digicash.com). The company addresses the issue of privacy. Most e-cash systems have an audit trail in their system for reasons of security, to allow a transaction, or e-cash value of money, to be tracked back to the source, no matter how many times the "until" of currency is traded between e-cash users. DigiCash is different from other e-cash systems, since it offers privacy in only one direction. As a DigiCash user requests a supply of e-cash tokens from the DigiCash server, those tokens are encrypted using a proprietary encryption system and e-mailed back to the holder. From there, when the holder buys something from a company or another e-cash account holder, he or she mails those tokens to the other person, who then banks them with DigiCash. While the recipient knows where the tokens have come from, and the sender knows where the e-cash was sent, DigiCash itself knows only that the e-cash has been banked by the person or organization who deposited the e-cash. So DigiCash knows where the e-cash has come from directly, but it does not know where the e-cash originated. This system is called one-way privacy, because no one, not even the police or legal authority can backtrack the chain to find out who has the e-cash and where it originated. Currently, DigiCash's e-cash system is provided by Mark Twain Bank in the United States. Recently, Eunet, the largest European ISP, also started issuing e-cash in cooperation with Merita Bank, Finland's largest commercial bank.

3.6 Short Summary

Server An application program that accepts connections in order to service requests by sending back responses.

There are basically two types of Web servers: public and private. You put up a public Web server to provide information to the world of Internet users. Anyone with a Web browser can attach to a public Web site.

3.7 Brain Storm

1. Explain about information highway?
2. Explain the types of server.
3. Distinguish between the public and private web server.
4. Explain the futures of Internet and intranet application.
5. Differentiate the web-server and client/server model.

Lecture 4

TCP / IP

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the evolution of TCP/IP
- Describe the rise of the Internet.
- Describe about TCP/IP protocol architecture.
- Describe about the TCP/IP Core protocols.
- Discuss about TCP/IP application interfaces.

Coverage Plan

Lecture 4

4.1 Snap shot

4.2 The evolution of TCP/IP

4.3 The rise of the internet

4.4 TCP/IP protocol architecture

4.5 TCP/IP core protocols

4.6 TCP/IP application interfaces

4.7 Short Summary

4.8 Brain Storm

4.1 Snap Shot

In this lecture we focuses on the evolution of TCP/IP. It notes the relationship of the two protocols to the ISO Open system interconnection reference model. Since a variety of applications were developed to use TCP/IP. The TCP / IP protocol suite, which is the entire set of layers in a protocol to include applications such as FTP, Telnet, and www program residing at the top layer of a protocol.

4.2 The Evolution of TCP/IP

The Transmission Control Protocol (TCP) and the Internet Protocol (IP) represent two specific protocols within a protocol stack that are commonly and collectively referred to as TCP/IP. Due to the evolution of those protocols, they are also commonly referred to as DOD protocols and the Internet protocol.

TCP/IP notes the relationship of the two protocols to the International Standards Organization (ISO) Open System Interconnection (OSI) Reference Model. Since a variety of applications were developed to use TCP/IP, below this also briefly discusses the operation and utilization of a number of those applications. You will gain an appreciation for the TCP/IP protocol suite, which is the entire set of layers in a protocol to include applications such as FTP, Telnet, and World Wide Web programs residing at the top layer of a protocol.

The work involved in establishing ARPANET resulted in the development of three specific protocols for the transmission of information: the Transmission Control Protocol (TCP), the Internet Protocol (IP), and the User Datagram Protocol (UDP), Both TCP and UDP represent transport-layer protocols. TCP is a transport-layer protocol that provides end-to-end reliable transmission, whereas UDP represents a connectionless-mode, layer-4 transport protocol. TCP includes such functions as flow control, error control and the exchange of status information, and is based on a connection between source and destination being established prior to the exchange of information. Thus, TCP provides an orderly and error-free mechanism for the exchange of information. In comparison, UDP is well suited for transaction-based applications that can work on a best-effort delivery scheme, such as the transmission of network management information where transmission efficiency is more important than reliability.

At the network layer, the IP protocol was developed as a mechanism to route messages between networks. To do so, IP was developed as a connectionless-mode, network layer protocol and includes the capability to segment or fragment and reassemble messages that must be routed between networks that support different packet sizes than the size supported by the source and / or destination networks. The set of networking standards developed by DARPA was officially named the TCP/IP protocol suite after its two main standards.

4.3 The Rise of The Internet

Although ARPANET was established as a research network, by the early 1980's DARPA began to convert other computers to use TCP/IP protocols. Gradually other networks were interconnected to ARPANET, which by default became the backbone of a network of interconnected networks that was referred to as the Internet. In January 1983 the U.S. Office of the Secretary of Defense mandated that all computers connected to long-distance networks should use the TCP/IP protocol suite for communications. At the same time, the U.S. government's Defense Communications Agency (DCA), which was responsible for ARPANET, split that network into two entities. One of the two networks retained the name ARPANET and continued in its role as a research network. The second network, which was for military use, was appropriately named MILNET.

Because a considerable amount of communications research was being performed by university computer science departments, DARPA made the TCP/IP protocol stack available at a low cost, which encouraged its adoption. Because most university computer science departments were using a version of the UNIX operating system, a large number of UNIX applications were initially developed for use with TCP/IP. Because such applications as electronic mail, remote login, and file transfer were developed, other groups began to adopt TCP/IP. In 1986 the National Science Foundation established a network that interconnected its six supercomputer centers. Known as NSFNET, this network was connected to the ARPANET. At approximately the same time, computer networks at universities, government agencies, and corporate research laboratories located throughout the world were being connected to the evolving Internet. To provide a mechanism for the adoption of standards necessary to maintain communications interoperability, DARPA formed the Internet Activities Board (IAB) in 1983. Recognizing that the expansion of the Internet included communications activities that now included basic file transfer and electronic mail necessary for daily business activities, the IAB was reorganized in 1989. That reorganization resulted in the formation of two major groups: the Internet Research Task Force (IRTF) and the Internet Engineering Task Force (IETF).

4.4 TCP/IP Protocol Architecture

TCP/ IP protocols map to a four-layer conceptual model known as the DARPA model, named after the U.S. government agency that initially developed TCP / IP. The four layers of the DARPA model are: Application, Transport, Internet, and Network Interface. Each layer in the DARPA model corresponds to one or more layers of the seven-layer Open Systems Interconnection (OSI) model.

Network Interface layer

The network Interface Layer (also called the Network Access Layer) is responsible for placing TCP/IP packets on the network medium and receiving TCP/IP packets off the network medium. TCP/IP was designed to be independent of the network access method, frame format, and medium. In this way, TCP/IP can be used to connect differing network types. This includes LAN technologies such as Ethernet or Token Ring and WAN technologies such as X.25 or Frame Relay. Independence from any specific network technology gives TCP/IP the ability to be adapted to new technologies such as Asynchronous Transfer Mode (ATM).

The Network Interface Layer encompasses the Data Link and Physical layers of the OSI Model. Note that the Internet Layer does not take advantage of sequencing and acknowledgment services that may be present in the Data Link Layer. An unreliable Network Interface Layer is assumed, and reliable communications through session establishment and the sequencing and acknowledgment of packets is the responsibility of the Transport Layer.

Internet Layer

The Internet Layer is responsible for addressing, packaging, and routing functions. The core protocols of the Internet Layer are IP, ARP, ICMP, and IGMP.

- The Internet Protocol (IP) is a routable protocol responsible for IP addressing and the fragmentation and reassembly of packets.
- The Address Resolution Protocol (ARP) is responsible for the resolution of the Internet Layer address to the Network Interface Layer address, such as a hardware address.

- The Internet Control Message Protocol (ICMP) is responsible for providing diagnostic functions and reporting errors or conditions regarding the delivery of IP packets.
- The Internet Group Management Protocol (IGMP) is responsible for the management of IP multicast groups.

The Internet Layer is analogous to the Network layer of the OSI model.

Transport Layer

The Transport Layer (also known as the Host-to-Host Transport Layer) is responsible for providing the Application Layer with session and datagram communication services. The core protocols of the Transport Layer are TCP and the User Datagram Protocol (UDP).

- TCP provides a one-to-one, connection-oriented, reliable communications service. TCP is responsible for the establishment of a TCP connection, the sequencing and acknowledgment of packets sent, and the recovery of packets lost during transmission.

- UDP provides a one-to-one or one-to-many, connectionless, unreliable communications service. UDP is used when the amount of data to be transferred is small, when the overhead of establishing a TCP connection is not desired, or when the applications or upper layer protocols provide reliable delivery.

The Transport Layer encompasses the responsibilities of the OSI Transport Layer and some of the responsibilities of the OSI Session Layer.

Application Layer

The Application Layer provides applications the ability to access the services of the other layers and defines the protocols that application use to exchange data. There are many Application Layer protocols and new protocols are always being developed.

The most widely known Application Layer protocols are those used for the exchange of user information.

- ✓ The Hyper Text Transfer Protocol (HTTP) is used to transfer files that make up the Web pages of the World Wide Web.

- ✓ The File Transfer Protocol (FTP) is used for interactive file transfer.

- ✓ The Simple Mail Transfer Protocol (SMTP) is used for the transfer of mail messages and attachments.
- ✓ Telnet, a terminal emulation protocol, is used for remote login to network hosts.

Additionally, the following Application Layer protocols help facilitate the use and management of TCP / IP networks:

- ∗ The Domain Name System (DNS) is used to resolve a host name to an IP address.

- ∗ The Routing Information Protocol (RIP) is a routing protocol that routers use to exchange routing information on an IP internetwork.

- ∗ The Simple Network Management Protocol (SNMP) is used between network management console and network devices to collect and exchange network management information.

Examples of Application Layer interfaces for TCP / IP applications are Windows Sockets and NetBIOS. Windows Sockets provides a standard application programming interface (API) under the Microsoft Windows operating system. NetBIOS is an industry-standard interface for accessing protocol services such as sessions, datagrams, and name resolution.

4.5 TCP / IP Core Protocols

IP

IP is a connectionless, unreliable datagram protocol primarily responsible for addressing and routing packets between hosts. Connectionless means that a session is not established before exchanging data. Unreliable means that delivery is not guaranteed. IP will always make a best effort attempt to deliver a packet. An IP packet might be lost, delivered out of sequence, duplicated, or delayed. IP does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets is the responsibility of a higher-layer protocol, such as TCP. IP is defined in RFC 791.
An IP packet consists of an IP header and an IP payload.

Key fields in the IP header

IP Header Field	Function
Source IP Address	The IP address of the original source of the IP datagram
Destination IP Address	The IP address of the final destination of the IP datagram.
Identification	Used to identify a specific IP Datagram and to identify all fragments of a specific IP datagram if fragmentation occurs.
Protocol	Informs IP at the destination host whether to pass the packet up to TCP,UDP, ICMP, or other protocols.
Checksum	A simple mathematical computation used to verify the integrity of the IP header.
Time to Liver (TTL)	Designates the number of networks on which the datagram is allowed to travel before being discarded by a router. The TTL is set by the sending host and is used to prevent packets from endlessly circulating on an IP internet work. When forwarding an IP packet, routers are required to decrease the TTL by at least one.

Fragmentation and Reassembly

If a router receives an IP packet that is to large for the network onto which the packet is being forwarded, IP will fragment the original packet into smaller packets that will fit on the downstream network. When the packets arrive at their final destination, IP at the destination host reassembles the fragments into the original payload. This process is referred to as fragmentation and reassembly. Fragmentation can occur in environments that have a mix of networking technologies, such as Ethernet and Token Ring.

The fragmentation and reassembly works as follows:

1. When an IP packet is sent by the source, it places a unique value in the identification field.

2. The IP packet is received at the router. The IP router notes that the maximum transmission unit (MTU) of the network onto which the packet is to be forwarded is smaller than the size of the IP packet.

3. IP fragments the original IP payload into fragments that will fit on the next network. Each fragment is sent with its own IP header which contains:

   - The original Identification field identifies all fragments that belong together.
   - The More Fragments Flag indicates that other fragments follow. The More Fragments Flag is not set on the last fragment, because no other fragments follow it.
   - The Fragment Offset field indicates the position of the fragment relative to the original IP payload.

4. When the fragments are received by IP at the remote host, they are identified by the identification field as belonging together. The Fragment Offset is then used to reassemble the fragments into the original IP payload.

ARP

When IP packets are sent on shared access, broadcast-based networking technologies such as Ethernet or Token Ring, the Media Access Control (MAC) address corresponding to a forwarding IP address must be resolved. ARP uses MAC-level broadcasts to resolve a known forwarding IP address to its MAC address. ARP is defined in RFC 826.

For more information on ARP see the “physical address resolution” section later in this paper.

ICMP

Internet control message protocol provides troubleshooting facilities and error reporting for packet that are undeliverable. For example if IP is unable to deliver a packet to the destination host, ECMP will send a Destination unreachable message to the source host. Below table shows the most common ICMP messages.

Common ICMP messages

ICMP Message	Function
Echo Request	Simple troubleshooting message used to check IP connectivity to a desired host.
Echo Reply	Response to an ICMP Request
Redirect	Sent by a router to inform a sending host of a better route to a destination IP address.
Source Quench	Sent by a router to inform a sending host that its IP datagrams are being dropped due to congestion at the router. The sending host then lowers its transmission rate. Source Quench is an elective ICMP message and is not commonly implemented.
Destination Unreachable	Sent by a router or the destination host to inform the sending host that the datagram cannot be delivered.

To send ICMP Echo Request messages and view statistics on the responses on a Windows NT-based computer, use the ping utility at a Windows NT command prompt.

Common ICMP Destination Unreachable messages

Destination Unreachable Message	Description
Network Unreachable	Sent by an IP router when a route to the description network can not be found.
Host Unreachable	Sent by an IP router when a destination host on the destination network can not be found. This message is only used on connection-oriented network technologies. IP routers on connectionless network technologies do not send Host Unreachable messages.
Protocol Unreachable	Sent by the destination IP node when the Protocol field in the IP header cannot be matched with an IP client protocol currently loaded.
Port Unreachable	Sent by the destination IP node when the Destination Port in the UDP header cannot be matched with a process using that port.
Fragmentation Needed and DF Set	Sent by an IP router when fragmentation must occur but is not allowed due to the source node setting the Don’t Fragment (DF) flag in the IP header.

There are a series of defined Destination Unreachable ICMP messages. Above Table describes the most common ICMP Destination Unreachable messages.

ICMP does not make IP a reliable protocol. ICMP attempts to report errors and provide feedback on specific conditions. ICMP messages are carried as unacknowledged IP datagrams and are themselves unreliable. ICMP is defined in RFC 792.

IGMP

Internet Group Management Protocol (IGMP) is a protocol that manages host membership in IP multicast groups. An IP multicast group, also known as a host group, is a set of hosts that listen for IP traffic destined for a specific multicast IP address. Multicast IP traffic is sent to a single MAC address but processed by multiple IP hosts. A given host listens on a specific IP multicast address and receives all packets to that IP address. Some additional aspects of IP multicasting:

- Host group membership is dynamic; hosts can join and leave the group at any time.
- A host group can be of any size.
- Members of a host group can span IP routers across multiple networks. This situation requires IP multicast support on the IP routers and the ability for hosts to register their group membership with local routers. Host registration is accomplished using IGMP.
- A host can sent traffic to an IP multicast address without belonging to the corresponding host group.

For a host to receive IP multicasts, an application must inform IP that it will be receiving multicasts at a specified destination IP multicast address. If the network technology supports hardware-based multicasting, then the network interface is told to pass up packets for a specific multicast address. In the case of Ethernet, the network interface card is programmed to respond to a multicast MAC address corresponding to the desired IP multicast address.

A host supports IP multicast at one of the following levels:

- Level 0
  No support to send or receive IP multicast traffic.
- Level 1
  Support exists to send but not receive IP multicast traffic.
- Level 2

Support exists to both send and receive IP multicast traffic. Windows NT TCP / IP supports level 2 IP multicasting

The protocol to register host group information is IGMP. IGMP is required on all hosts that support level 2 IP multicastings. IGMP packets are sent using an IP header.

IGMP messages take two forms:

1. When a host joins a host group, it sends an IGMP Host Membership Report message to the all-hosts IP multicast address or to the desired multicast address declaring its membership in a specific host group by referencing the IP multicast address.
2. When a router polls a network to ensure there are members of a specific host group, it sends an IGMP Host Membership Query message to the all-hosts IP multicast address. If no responses to the poll are received after several polls, the router assumes no membership in that group for that network and stops advertising that group-network information to other routers.

For IP multicasting to span routers across in internetwork, multicast routing protocols are used by routers to communicate host group information so that each router supporting multicast forwarding is aware of which networks contain members of which host groups.

TCP

TCP is a reliable, connection-oriented delivery service. The data is transmitted in segments. Connection-oriented means that a connection must be established before hosts can exchange data. Reliability is achieved by assigning a sequence number to each segment transmitted. An acknowledgment is used to verify that the other host received the data. For each segment sent, the receiving host must return an acknowledgment (ACK) within a specified period for bytes received. If an ACK is not received, the data is retransmitted. TCP is defined in RFC 793.

TCP uses byte-stream communications, wherein data within the TCP segment is treated as a sequence of bytes with no record or field boundaries. Below table describes the key fields in the TCP header.

Key fields in the TCP header

Field	Function
Source Port	TCP port of sending host
Destination Port	TCP port of destination host
Sequence Number	The sequence number of the first byte of data in the TCP segment.
Acknowledgment Number	The sequence number of the byte the sender expects to receive next from the other side of the connection.
Window	The current size of a TCP buffer on the host sending this TCP segment to store incoming segments.
TCP Checksum	Verifies the integrity of the TCP header and the TCP data.

TCP Parts

A TCP port provides a specific location for delivery of TCP segments. Port number below 1024 is well-known ports and is assigned by the Internet Assigned Numbers Authority (IANA). Below the Table lists a few well-known TCP ports.

Well-known TCP ports

TCP Port Number	Description
20	FTP (Data Channel)
21	FTP (Control Channel)
23	Telnet
80	Hyper Text Transfer Protocol (HTTP) used for the World Wide Web
139	NetBIOS session service

For a complete list of assigned TCP ports, see RFC 1700.

The TCP Three-Way Handshake

A TCP connection is initialized through a three-way handshake. The purpose of the three-way handshake is to synchronize the sequence number and acknowledgment numbers of both sides of the connection, exchange TCP Window sizes, and exchange other TCP options such as the maximum segment size. The following steps outline the process:

1. The client sends a TCP segment to the server with an initial Sequence Number for the connection and a Window size indicating the size of a buffer on the client to store incoming segments from the server.

2. The server sends back a TCP segment containing its chosen initial Sequence Number, an acknowledgment of the client's Sequence Number, and a Window size indicating the size of a buffer on the server to store incoming segments from the client.

3. The client sends a TCP segment to the server containing an acknowledgement of the server's Sequence Number.

TCP uses a similar handshake process to end a connection. This guarantees that both hosts have finished transmitting and that all data was received.

UDP

UDP provides a connectionless datagram service that offers unreliable, best-effort delivery of data transmitted in messages. This means that the arrival of datagram is not guaranteed; nor is the correct sequencing of delivered packets. UDP does not recover from lost data through retransmission. UDP is defined in RFC 768.

UDP is used by applications that do not require an acknowledgment of receipt of data and that typically transmit small amounts of data at one time. The NetBIOS name service, NetBIOS datagram service, and the Simple Network Management Protocol (SNMP) is examples of services and applications that use UDP. Below Table describes the key fields in the UDP header.

Key fields in the UDP header

Field	Function
Source Port	UDP port of sending host
Destination Port	UDP port of destination host.

UDP Checksum	Verifies the integrity of the UDP header and the UDP data.
Acknowledgment Number	The sequence number of the byte the sender expects to receive next from the other side of the connection.

UDP Ports

To use UDP an application must supply the IP address and UDP port number of the destination application. A port provides a location for sending messages. A port functions as a multiplexed message queue, meaning that it can receive multiple messages at a time. A unique number identifies each port. It's important to note that UDP ports are distinct and separate from TCP ports even though some of them use the same number.

Well-known UDP ports

UDP Port Number	Description
53	Domain Name System (DNS) Name Queries
69	Trivial File Transfer Protocol(TFTP)
137	NetBIOS name service
138	NetBIOS datagram service
161	Simple Network Management Protocol (SNMP)

For a complete list of assigned UDP ports, see RFC 1700.

Windows Sockets Interface

The windows Sockets API are a standard interface under Microsoft Windows for applications that use TCP and UDP. Applications written to the Windows Sockets API will run on many versions of TCP/IP. TCP/IP utilities and the Microsoft SNMP service are examples of applications written to the Windows Sockets interface.

Windows Sockets provides services that allow applications to bind to a particular port and IP address on a host, initiate and accept a connection, send and receive data, and close a connection. There are two types of sockets:

1. A stream socket provides a two-way, reliable, sequenced, and unduplicated flow of data using TCP.

2. A datagram socket provides the bi-directional flow of data using UDP.

A socket is defined by a protocol and an address on the host. The format of the address is specific to each protocol. In TCP / IP, the address is the combination of the IP address and port. Two sockets, one for each end of the connection, form a bi-directional communications path.

To communicate, an application specifies the protocol, the IP address of the destination host, and the port of the destination application. Once the application is connected, information can be sent and received.

NetBIOS Interface

NetBIOS (Network Basic Input/Output System) was developed for IBM in 1983 by Sytek Corporation to allow applications to communicate over a network. NetBIOS defines two entities, a session level interface and a session management / data transport protocol.

The NetBIOS interface is a standard API for user applications to submit network I/O and control directives to underlying network protocol software. An application program that uses the NetBIOS interface API for network communication can be run on any protocol software that supports the NetBIOS interface API for network communication can be run on any protocol software that supports the NetBIOS interface.

NetBIOS also defines a protocol that functions at the session/transport level. This is implemented by the underlying protocol software, such as the NetBIOS Frames Protocol (NBFP, a component of NetBEUI) or NetBIOS over TCP/IP (NetBT), to perform the network I/O required to accommodate the NetBIOS interface command set. NetBIOS over TCP/IP is defined in RFCs 1001 and 1002.

NetBIOS provides commands and support for NetBIOS Name Management, NetBIOS Datagrams, and NetBIOS Sessions.

NetBIOS Name Management

NetBIOS Name Management services provide the following function:

Name Registration and Release

When a TCP/IP host initializes, it registers its NetBIOS names by broadcasting or directing a NetBIOS name registration request to a NetBIOS Name Server such as Windows Internet Name Service (WINS) server. If another host has registered the same NetBIOS name, either the host or a NetBIOS Name Server responds with a negative name registration response. The initiating host receives an initialization error as a result.

When the workstation service on a host is stopped, the host discontinues broadcasting a negative name registration response when someone else tries to use the name and sends a name release to a NetBIOS Name Server. The NetBIOS name is said to be released and available for use by another host.

Name Resolution

When a NetBIOS application wants to communicate with another NetBIOS application, the IP address of the NetBIOS application must be resolved. NetBIOS over TCP/IP performs this function by either broadcasting a NetBIOS name query on the local network or sending a NetBIOS name query to a NetBIOS Name Server.

NetBIOS Datagram

The NetBIOS datagram service provides delivery of datagrams that are connectionless, non-sequenced, and unreliable. Datagrams can be directed to a specific NetBIOS name or broadcast to a group of names. Delivery is unreliable in that only the users who are logged on to the network will receive the message. The datagram service can initiate and receive both broadcast and directed messages. The datagram service uses UDP port 1`38.

NetBIOS Sessions

The NetBIOS session service provides delivery of NetBIOS messages that are connection-oriented, sequenced, and reliable. NetBIOS sessions use TCP connections and provide session establishment, keepalive, and termination. The session service allows concurrent data transfers in both directions using TCP port 139.

4. 6 The TCP/ IP Application Interfaces

7. Applications 6. Presentation DNS 5. Session	File Transfer Protocol	Telnet	HTTP	SMTP	SNMP	
4. Transport Layer	TCP				UDP	
3. Network Layer	IP					
2. Data Link Layer	Data Link Layer					
1. Physical Layer	Physical Layer					

A portion of the TCP/IP protocol suite

As a layered communications protocol, TCP/IP groups functions into defined network layers. Above figure illustrates a portion of the TCP / IP protocol suite, and indicates the relationship between TCP / IP protocols and the seven-layer OSI Reference Model.

The Left portion of figure indicates the seven layers of the ISO's OSI Reference Model. At the lowest position in the model, the physical layer provides a set of rules that governs the electrical and physical connection between devices. At the next layer in the reference model, the data link layer denotes how a device obtains access to the medium specified in the physical layer. In addition, the data link layer defines the framing of information within messages as well as error control procedures to control the flow of data between two nodes on a network. On a LAN, examples of the data link layer include Ethernet and token ring, while wide are network (WAN) examples of data link protocols include binary synchronous communications (BSCs) and high-level data link control (HDLC).

At the third layer in the ISO OSI Reference Model, the network layer has the responsibility for arranging the establishment of a logical connection between source and destination devices. Services provided at the network layer support the movement of data through a network and include addressing, routing, switching, sequencing, and flow control procedures. Because large, complex networks might result in several paths requiring traversal for a source to communicate with the desired destination, the routing of data through the network is an important feature provided by the network layer.

After a route has been established through a network, it is important to ensure that the transfer of information occurs correctly. This function is the responsibility of the transport layer of the OSI Reference Model, which uses error control, sequence clicking, and other end-to-end reliability-enhancing functions to guarantee the reliable transfer of information.

While the first four layers in the OSI Reference Model are fairly well defined and incorporated into most modern protocols, the so-called "upper" layers are not fully followed by most protocols. Since TCP / IP preceded the development of the OSI Reference Model, it comes as no surprise that it deviates from that model. That deviation is in the upper layers, where TCP/IP application functions correspond to the upper three layers of the reference model.

The OSI Reference Model's session layer provides a set of rules for establishing and terminating data transfers between nodes in a network. In comparison, the presentation layer is responsible for the conversion of transmitted data into a display format suitable for a receiving device. To provide this function, the presentation layer will transform and format data. Finally, at the top layer of the OSI Reference Model, the applications layer acts as a window through which the application gains access to the services provided by the model. Here applications can include file transfer, electronic messaging, and the capability to access a distant computer as if the computer executing the application were directly cabled to the distant device.

In examining figure note that only six of hundreds services are shown. Those applications essentially correspond to the upper three layers of the OSI Reference Model. Also note that the subdivision of the transport layer in above Figure is used as a mechanism to indicate which applications are transported by TCP and those carried by UDP. Concerning those applications, the File Transfer Protocol (FTP) provides a mechanism for the orderly transfer of files between computer systems while Telnet represents an application that enables a terminal to be connected to a remote host as if it were directly connected to the remote computer. Although Telnet was originally developed as a protocol to enable dumb terminals to connect to remote hosts, Telnet applications now enable PCs access mainframes, minicomputers, and other PCs as if they were directly cabled to those distant computers.

The Hypertext Transfer Protocol (HTTP) enables the transportation of World Wide Web (WWW) pages from Web servers to browsers operating on client computers, whereas the

Simple Mail Transfer Protocol (SMTP) provides a standard method for exchanging electronic mail. Each of the previously discussed applications runs on top of TCP. In comparison, the Simple Network Management Protocol (SNMP), which supports the management of network devices, and the domain name system (DNS), which provides a mechanism for locating the address of computers referenced by English mnemonics, are transported by UDP.

One of the key reasons for the dramatic growth in the use of the TCP/IP protocol suite can be traced to the development and structure of TCP/IP. Because TCP/IP was developed using taxpayer funds, its specifications were placed in the public domain, and they are available for use royalty-free, for vendors to develop "protocol stacks" in software or firmware to implement TCP/IP's operational features. Another key reason for the growth in the use of TCP/IP for both private and corporate networks, as well as for the interconnection of academic, commercial, and private networks to the Internet, is its structure, which makes it suitable for both LAN and WAN operations.

As the data link layer shown in above figure, TCP/IP can be transported within Ethernet, token ring, FDDI, or another type of local area network frame. Because a considerable amount of effort was expended in developing LAN adapter cards to support the bus structure used in Apple Macintosh, IBM PCs, and compatible computers, Sun Microsystems's workstations, and even IBM mainframes, the development of software-based protocol stacks to facilitate the transmission of TCP/IP on LANs provides the capability to interconnect LAN-based computers to one another whether they are on the same network and merely require the transmission of frames on a common cable, or whether they are located on networks thousands of miles from one another. Concerning the latter, through the use of transport layer and network layer services, TCP/IP can transport data via simple to complex wide area network structures. This provides a mechanism to satisfy academic, commercial, and governmental communications requirements from LAN to LAN on both an intra-and inter-LAN basis. Therefore, the use of TCP/IP has significantly increased and Microsoft Corporation has for some time recognized its importance as a networking protocol by the inclusion of a TCP/IP networking in Windows NT.

4.7 Short Summary

The TCP/IP suite both as a protocol and a set of services for connectivity and management of IP internetworks.

To allow applications to access the services offered by the core TCP/IP protocols in a standard way, network operating systems the Windows NT make industry standard application programming interfaces (APIs) available. Application programming interfaces are sets of functions and commands that are programmatically called by application code to perform network functions. For example, a Web browser application connecting to a Web site needs access to TCP's connection establishment service.

4.8 Brain Storm

1. Explain the architecture of TCP/IP protocol.
2. Explain the evolution of TCP/IP.
3. What is OSI model explain the architecture of OSI.
4. Explain about TCP/IP application interfaces.

Lecture 5

IP Addressing

Objectives

After completing this lesson, you should be able to do the following

- Discuss about IP addressing
- Describe the Basic address scheme.
- Describe about address classes.
- Describe about dotted-decimal notation.
- Discuss about networking basics.
- Discuss about host restrictions.
- Describe about subnet and subnet mask.
- Discuss the domain name system.

Coverage Plan

Lecture 5

5.1 Snap Shot

In this session we discuss about IP addressing. TCP allows two computer programs to exchange information over a network or a group of networks in a reliable manner. When a computer wants to establish a dialog with another system over the Internet, it opens a TCP connection to the other computer.

5.2 IP Addressing

Each TCP/IP host is identified by a logical IP address. The IP address is a network layer address and has no dependence on the data link layer address. A unique IP address is required for each host and network component that communicates using TCP/IP.

The IP address identifies a system's location on the network in the same way a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IP address must be globally unique and have a uniform format.

Each IP address includes a network ID and a host ID.

- The network ID identifies the systems that are located on the same physical network bounded by IP routers. All systems on the same physical network must have the same network ID. The network ID must be unique to the internetwork.

- The host ID identifies a workstation, server, router, or other TCP/IP host within a network. The address for each host must be unique to the network ID.

An IP address is 32 bits long. Rather than working with 32 bits at a time, it is a common practice to segment the 32 bits of the IP address into four 8-bit fields called octets. Each octet is converted to a decimal number in the range 0-255 and separated by a period.(a dot). This format is called dotted decimal notation. Below table provides an example of an IP address in binary and dotted decimal formats.

Example of an IP addresses in binary and dotted decimal format

Binary format	Dotted Decimal Notation
11000000 10101000 00000011 00011000	192.168.3.24

The notation w.x.y.z is used when referring to a generalized IP address.

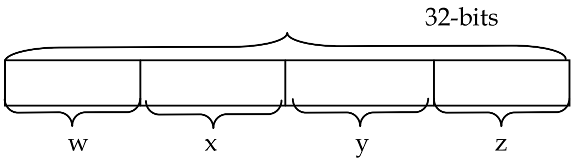

5.3 Basic Address Scheme

In most situations, an IP address is used to identify a network as well as a host connected to the network. Thus, an IP address normally represents a two-level addressing hierarchy consisting of a network prefix and a host address. The hierarchy is two level IP addressing structure.

Network prefix	Host address

All hosts on the same network must be assigned the same network prefix but must have a unique host address to differentiate one host from another. Similarly, two hosts on different networks must each be assigned a different network prefix; however, those hosts can both have same host address.

5.4 Address Classes

The Internet community originally defined five address classes to accommodate networks of varying sizes. Microsoft TCP/IP supports class A,B, and C addresses assigned to hosts. The class of address defines which bits are used for the host ID. It also defines the possible number of networks and the number of hosts per network.

Class A

Class A addresses are assigned to networks with a very large number of hosts. The high-order bit in a class A address is always set to zero. The next seven bits complete the network ID. The remaining 24 bits represent the host ID. This allows for 126 networks and 16,777,214 hosts per network. Below figure illustrates the structure of class A addresses.

Class A

Network ID Host ID

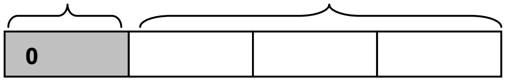

Class B

Class B addresses are assigned medium-sized to large-sized networks. The two high-order bits in a class B address are always set to binary 1 0. The next 14 bits complete the network ID. The remaining 16 bits represent the host ID. This allows for 16,384 networks and 65,534 networks and 65,534 hosts per network per network. Below figure illustrates the structure of class B addresses.

Class A

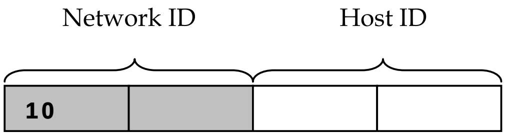

Class C

Class C addresses are used for small networks. The three high-order bits in a class C address are always set to binary 1 1 0. The next 21 bits complete the network ID. The remaining 8 bits represent the host ID. This allows for 2,097,152 networks and 254 hosts per network. Below figure illustrates the structure of class C addresses.

Class C

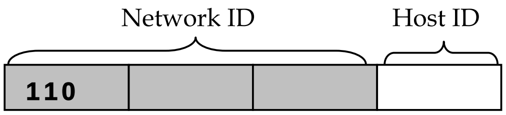

Class D

Class D addresses are reserved for IP multicast addresses. The four high-order bits in a class D addresses are always set to binary 1 1 1 0. The remaining bits are for the address that interested hosts will recognize. Microsoft supports class D addresses for applications to multicast data to multicast-capable hosts on an internetwork.

Class E

Class E addresses are experimental addresses reserved for future use. The high-order bits in a class E address are set to 1 1 1 1.

Below table is a summary of address classes A,B and C that can be used for host IP addresses.

IP address class summary

Class	Value for w	Network ID Portion	Host ID Portion	Available Networks	Hosts per Network
A	1-126	W	x.y.z	126	16,777,214
B	128-191	w.x	y.z	16,384	65,534
C	192-223	w.x.y	Z	2,097,152	254

Network ID Guidelines

The network ID identifies the TCP/IP hosts that are located on the same physical network. All hosts on the same physical network must be assigned the same network ID to communicate with each other.

Follow these guidelines when assigning a network ID

- ❑ The network address must be unique to the IP internetwork. If you plan on having a direct routed connection to the public Internet, the network ID must be unique to the Internet. If you do not plan on connecting to the public Internet, the local network ID must be unique to your private internetwork.

- ❑ The network ID cannot begin with the number 127. The number 127 in a class A address is reserved for internal loopback functions.

- ❑ All bits within the network ID cannot be set to 1. All 1's in the network ID are reserved for use as an IP broadcast address.

- ❑ All bits within the network ID cannot be set to 0. All 0's in the network ID are used to denote a specific host on the local network and will not be routed.

Below the table lists the valid ranges of network IDs based on the IP address classes. To denote IP network IDs, the host bits are all set to 0. Note that even though expressed in dotted decimal notation, the network ID is not an IP address.

Address Class	First Network ID	Last Network ID
Class A	1.0.0.0	126.0.0.0

Class B	128.0.0.0	191.255.0.0
Class C	192.0.0.0	223.255.255.0

Host ID guidelines

The host ID identifies a TCP/IP host within a network. The combination of IP network ID and IP host ID is an IP address.

Follow these guidelines when assigning a host Id.

The host ID must be unique to the network ID.

All bits within the host ID cannot be set to 1, because this host ID is reserved as a broadcast address to send a packet to all hosts on a network.

All bits in the host ID cannot be set to 0, because this host ID is reserved to denote the IP network ID.

Below table lists the valid ranges of host IDs based on the IP address classes.

Address Class	First Host ID	Last host ID
Class A	W.0.0.1	w.255.255.254
Class B	w.x.0.1	w.x.255.254
Class C	w.x.y.1	w.x.y.254

5.5 Dotted Decimal Notation

Recognizing that the direct use of 32-bit binary addresses is both cumbersome and unwieldy to deal with, a technique more acceptable for human use was developed. That technique is referred to as dotted-decimal notation, in recognition of the fact that technique developed to express IP addresses occurs via the use of four decimal numbers separated from one another by decimal points.

Dotted-decimal notation divides the 32-bit Internet protocol address into four 8-bit fields, with the value of each field specified as a decimal number. That number can range from 0 to

255 in bytes 2,3 and 4. In the first byte of an IP address, the setting of the first 4 bits in the byte used to denote the address class limits the range of decimal values that can be assigned to that byte. Class A address is defined by setting the first bit position in the first byte to 0. Thus, the maximum value of the first byte in a Class A address is 128.

To illustrate the formation of a dotted-decimal number, Let's first focus on the decimal relationship of the bit positions in a byte. Below fig indicates the decimal values of the bit positions within an 8-bit byte. Note that the decimal value of each bit position corresponds to 2^n, where n is the bit position in the byte. Using the decimal values of the bit positions. Let's assume that the first byte in an IP address has its positions set as 11000000. Then, the value of that byte expressed as a decimal number becomes 128+64, or 192. Now let's assume that the second byte in the IP address has the bit values 01001000. The decimal value of that binary byte is 64+8, or 72. Let's further assume that the last 2 bytes in the IP address have the bit values 00101110 and 10000010. then, the third byte would have the decimal value 32+8+4+2, or 46, while the last byte would have the decimal value 128+2 or 130.

Based on the preceding, the dotted-decimal number 192.72.46.130 is equivalent to the binary number 11000000100100000101110100000010. Obviously, it is easier to work with, as well as remember, four decimal numbers separated by dots than a string of 32bits.

5.6 Networking Basics

Each network has a distinct network prefix, and each host on a network has a distinct host address. When two networks are interconnected by the use of a router, each router port is assigned an IP address that reflects the network it is connected. The connection of two networks via a router, Indicating possible addresses assignments. Note that the first decimal number of the 4-byte, dotted-decimal numbers associated with two hosts on the network on the left portion denotes a Class C address. This is because 192 decimal is equivalent to 11000000 binary. Because the first 2 bits are set to the bit value 11. Also note that the first 3 bytes of a Class C address indicate the network while the fourth byte indicates the host address. Thus, the network shown in the left portion of below figure is denoted as 192.78.46, with device addresses that can range from 192.78.46.0 through 192.78.46.255.

In the lower-right portion of below figure two hosts are shown connected to another network. Note that the first byte for the 4-byte, dotted-decimal number assigned to each host and the router port is decimal 226, which is equivalent to binary 11100010. Because the first 2 bits in the first byte are again set to 11, the second network also represents the use of a Class C address. Thus, the network address is 226.42.78, with device addresses on the network ranging from 226.42.78.00 to 226.42.78.255.

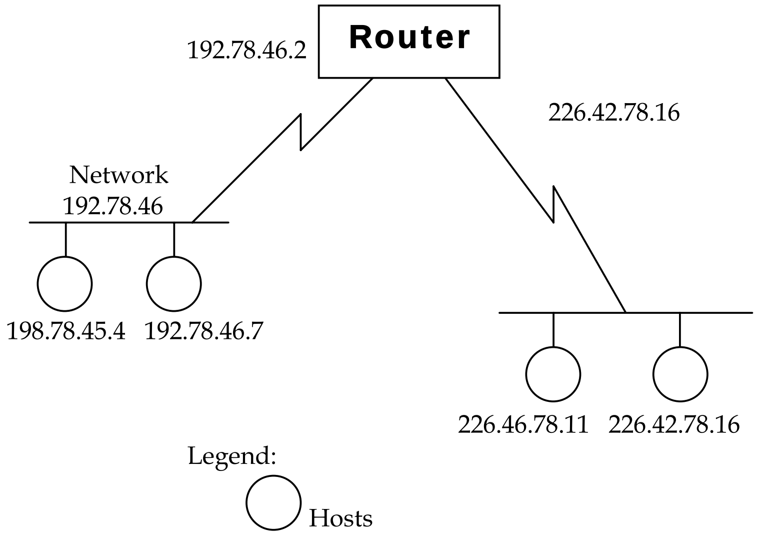

5.7 Host Restrictions

Although it would appear that 256 devices could be supported on a Class C network, in actuality the host portion field of an IP address has two restrictions. First, the host portion field cannot be set to all 0 bits. This is because an all-zeros host number is used to identify a base network or subnetwork number. Second, an all-ones host number represents the broadcast address for a network or subnetwork. Due to these restrictions, a maximum of 254 devices can be defined for use on a Class C network. Similarly, other network classes have the previously discussed addressing restrictions, which reduces the number of distinct addressable devices that can be connected to each type IP network by two. Since, as previously explained, an all-zeros host number identifies a base network, the two networks more commonly are shown numbered as 192.78.46.0 and 226.42.78.0.

5.8 Subnets

One of the problems associated with the use of IP addresses is the necessity to assign a distinct network address to each network. This can result in the waste of many addresses as well as require a considerable expansion in the use of router tables. To appreciate these problems let’s return to which illustrates that connection of two class C networks via a router.

Assume that each class C network supports 24 workstations and servers. Adding an address for the router port, each class C network would use 25 out of 254 available addresses. Therefore, the assignment of two class C addresses to an organization with a requirement to support two networks with a total of 50 devices would result in 458 available IP addresses, in effect, being wasted, in additions, routers would have to recognize two network addresses instead of one. When this situation is multiplied by numerous organizations requiring multiple networks, the effect on routing tables becomes more pronounced, resulting in extended search times as routers sort through their routing tables to determine an appropriate route to a network. Recognizing the preceding problems, RFC 950 became a standard in 1985. that standard defines a procedure to subnet or divide a single class A B, or C network into subnetworks.

Through the process of subnetting, the tow level hierarchy of Class A, b and C networks shown in turned into a three level hierarchy. In doing so, the host portion of an IP address is divided into a subnet portion and a host portion. Provides a comparison of the two level hierarchy initially defined for class A B, and C network and the three level subnet hierarchy.

Through the process of subnetting, a Class A B, or C network address can be divided into different subnet numbers, with each subnet used to identify a different network internal to an organization. Since the network portion of the address remains the same, the route from the Internet to any subnet of a given IP network address is the same. This means that routers within the organizations must be able to differentiate between different subnets but routers outside the organization consider all subnets as one network.

Let’s examine the process to illustrate the subnet process as well as to obtain an appreciation for how it facilitates the use of IP addresses in a less wasteful manner and reduces routing table entries. In doing so, we will discuss the concept of masking and the user of the subnet makes, both of which are essential to the extension of the network portion of an IP address beyond its network portion of the address.

To illustrate the concept of subnetting let’s return to the two networks illustrated in 192. 78.46.0 and 226.42.78.0. let’s assume that instead of two networks geographically separated from one another at two distinct locations, you require the establishment of five networks at one location. Let us further assume that each of the five networks will supports a maximum of 15 stations. Although your organization could apply for four additional class C addresses, doing so would waste precious IP address space since each class C address supports a maximum of 254 devices. In addition your internal network were connected to the Internet, entries for four additional networks would be required in a number of routers in the Internet in addition to your internal routers. Instead of requesting four additional class C addresses, lets use subnetting, dividing the host portion of the IP address into a subnet number and a host number. Because you need to support five networks at one location, you must use a minimum of 3 bits from the host portion of the IP address as the subnet number. Since a class C address uses one 8 bit byte for the host identification, this means that a maximum of five bit positions can be used for the host number. Assuming that you intend to use the 192.78.46.0 network address for the subnetting effort, you would construct an extended network prefix based on combining the network portion of the IP address with its subnet number.

Illustrates the creation of five subnets from the 192.78.46.0 network address. The top entry in which is labeled base network represents the class C network address with a host address but field set to all zeros. Since you previously decided to use 3 bits from the host portion of the class C IP address to develop an extended network prefix, the five entries in below the base network entry indicate the use of 3 bits from the host position in the address to create extended prefixes that identify five distinct subnets created from one IP Class C address. To the internet, all five networks appear as the network address 198.78.460, with the router at an organization responsible for directing traffic to the appropriate subnet. It is important to note that externally (that is to the Internet) there is not knowledge that the dotted decimal numbers shown in the right column represent distinct subnets. This is because the Internet views the first byte of each dotted decimal number and notes that the first 2 bits are set. Doing so tells routers on the internet that the address is a class c address for which the first 3 bytes represent the network portion of the IP address and the fourth byte represents the host address. Therefore, to the outside world address 198.78.46.32 would not be recognized as subnet 1. Instead, a router would interpret the address as network 198.78.46.0. with host address 32. similarly, subnet 4 would appear as network address 198.78.46.0 with host address 128. However internally within an organization, each of the addresses listed in the right column in would be recognized as a subnet. To visualize this dual interpretation of network addresses, consider which illustrates the internet versus the private network view of subnets.

As you might logically assume from the previous discussion of class C addresses, any address with the network prefix 198.78.46.0 will be routed to the corporate router. However although you noted how subnet addresses are formed, you have yet to learn how to assign host addresses to devices connected to different subnets and how the router can break down a subnet address so that it can correctly route traffic to an appropriate subnet.

Subnets and subnet masks

The Internet address classes were designed to accommodate three different scales of IP internetworks, where the 32 bits of the IP address are appropriate between network IDs and host IDs depending on how many networks and hosts per network are needed. However, consider the class A network ID, which has the possibility of over 16 million hosts on the same network. All the hosts on the same physical network bounded by IP routers share the same broadcast traffic; they are in the same broadcast domain.It is not practical to have 16 million nodes in the same broadcast domain. The result is that most of the 16 million host addresses are not assignable and are wasted. Even a class B network with 65 thousand hosts is impractical.

In an effort to create smaller broadcast domains and to better utilize the bits in the host ID, an IP network can be subdivided into smaller networks, each bounded by an IP router and assigned a new subnetted network ID, which is a subset of the original class-based network ID.

This creates subnets, subdivisions of an IP network each with their own unique subnetted network ID. Subnetted network IDs are created by using bits from the host ID portion of the original class-based network ID.

Consider the below figure The class B network of 139.12.0.0 can have up to 65,534 nodes. This is far too many nodes, and in fact, the current network is becoming saturated with broadcast traffic. This subnetting of network 139.12.00 should be done in such a way so that it does not impact, nor require, the reconfiguration of the rest of the IP internetwork.

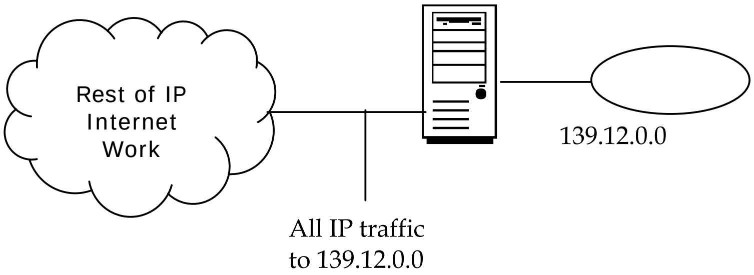

Net work 139.12.0.0 is subnetted by utilizing the first 8 host bits for the new subnetted network ID. When 139.12.0.0 is subnetted, as shown in next fig. Separate networks with their own subnetted network IDs are created. The router is aware of the separate subnetted network IDs and will route IP packets to the appropriate subnet.

Note that the rest of the IP internetwork still regards all the nodes on the three subnets as being on network 139.12.0.0. The other routers in the IP internetwork are unaware of the subnetting being done on network 139.12.0.0, and therefore require no reconfiguration.

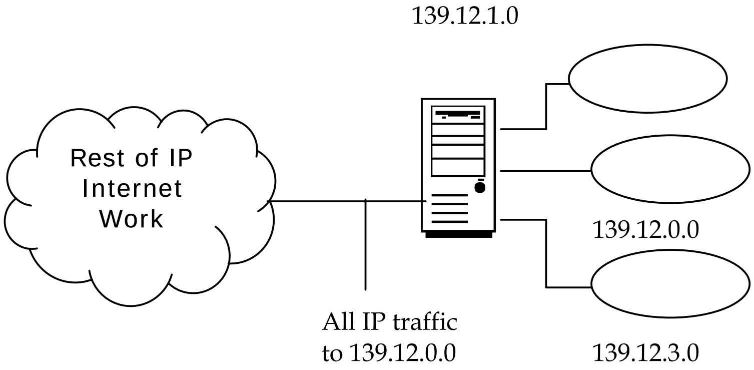

A key element of subnetting is still missing. How does the router who is subdividing network 139.112.0.0 know how the network is being subdivided and which subnets are available on which router interfaces? To give the IP nodes this new level of awareness, it must be told exactly how to discern the new subnetted network ID regardless of internet address classes. To tell an IP node exactly how to extract a network ID, either class-based or subnetted, a subnet mask is used.

Host addressees on subnets

You previously subdivided the host portion of a class c address into a 3 bit subnet filed and a 5 bit host field. Since the host field of an IP address cannot contain all) bits or all 1 bits the use of 5 bits in the host potion of each subnet address means that each subnet can support a maximum of 2.2 or 30, addresses. Therefore you could use host address 1 through 30 on each subnet. Illustrates the assignment of host addresses for subnet 3 whose creation is indicated in. In examine note that you start with the subnet address 198.78.46.96, for which the first 3 bits in the 4 byte dotted decimal number are used to indicate the subnet. They you use the remaining 5 bits to define the host address on each subnet. Therefore, the address 198.78.46.96 represents the third subnet, while addresses 198.78.46.97 through 198.78.46.126 represent hosts 1 through 30 that can reside on subnet3.

Although you now have an appreciation for creating subnets and host addresses on subnets, an unanswered question is How do devices on a private network recognize subnet addressing? For example, if a packet arrives at an organization router with the destination address 198.78.46.97, how does the router know to route that packet onto subnet 3?. The answer to this question involves what is knows as the subnet mask.

The subnet Mask

The subnet mask represents a mechanism that enables devices on a network to determine the separation of an IP address into its network, subnet, and host portions. To accomplish this, the subnet mask consists of a sequence of 1 bit that denotes the length of the network and subnet portions of the IP network address associated with a network. For example let’s assume that your network address is 198.78.46.96 and you want to develop a subnet mask that can be used to identify the extended network. Since you previously used 3 bits from the host portion of the IP address, the subnet mask would become this:

11111111.11111111.11111111.11100000

similar to the manner by which IP addresses can be expressed using dotted decimal notation, you can also express subnet masks using that notation. Doing so you can express the subnet mask as this:

255.255.255.224

The subnet mask tells the device examining an IP address which bits in the address should be treated as the extended network address consisting of network and subnet addresses. Then the remaining bits that are not set in the mask indicate the host on the extended network address. However how does a device determine the subnet of the destination address? Since the subnet mask indicates the length of the extended network to include the network and subnet fields knowing the length of the network portions of the address provides a device with the capability to determine the number of bits in the subnet field. After this is accomplished the device can determine the value of those bits which indicates the subnet. To illustrate this concept let's use the IP address 198.78.46.97 and the subnet mask 255.255.255.224, with the latter used to define a 27 bit extended network. The relationship between the IP address and the subnet mask is shown.

The first 2 bits in the IP address are set, which indicates a Class C address. Because a Class C address consists of 3 bytes used for the network address and 1 byte for the host address, the subnet must be 3 bits in length . Thus bits 25 through 27, which are set to 011 in the IP address, identify the subnet as subnet 3. because the last 5 bits in the subnet mask are set to 0, those bit positions in the IP address identify the host on subnet 3. Because those bits have the value 00001, the IP address references host 1 on subnet 3 on network 1908.78.46.0. Now that you have an appreciation for the formation and use of IP address, subnets, and subnet masks, let's turn to the second key addressing related topic to be covered in this chapter the domain name system.

5.9 Domain Name System

Although the use of dotted decimal notation is certainly easier to employ that 32 bit binary addresses, it is still difficult to remember a sequence of four numerics separated by decimal points. This is certainly true when a user requires connectivity with a large number of hosts. In addition, the use of dotted decimal notation does not provide an indication of the user of the desired host nor its organization. Recognizing these limitations, ARPANET incorporated a mechanism that allowed English type mnemonics and names to be assigned to hosts.

During the 1970 when ARPANET provided connectivity for a small number of computers, a single file host.txt was used to provide a hostname to network address translation for each host on the network. That file was maintained at the ARPANET network information center and as the network grew in size and complexity, the user of a centrally located file for name to address translations resulted in a series of problems that were eventually alleviated by the implementation of the domain name system. Two of the major problems associated with the use of a single host.txt file were the traffic directed to the host maintaining the file and available name assignment. As more and more hosts were added to the network, traffic routed to the host containing the file host.txt literally exploded, resulting in both transmission and host processing delays. Second, the use of a single file for name to address translations precluded two hosts from having the same name, even if they were requested by different organizations. Recognizing these and other problems associated with using a single name to address translation point resulted in the development of a distributed database to perform the same process. That database is the key to the operation of DNS.

DNS represents a client/ server approach that was developed to translate English type names assigned to hosts to their IP addresses. Servers, more formally referred to as name server programs, contain information about portions of the DNS database, information captained in name servers is requested via clients, which are more formally referred to asresolvers. Resolvers are commonly implemented as library routines within an application program, such as FTP. Such routines create queries that are transmitted to a name server. Therefore, the entry of an English type name such as ftp.ipretzel.com as the destination address in an FTP connection, would result in the client application sending a request to a name server for the IP address of the destination.

Database Structure

The DNS database is structured similarly to an inverted tree with its root at the top of the structure. Directly under the root are primary domain, with each domain further divided into partitions referred to as subdomains. Indicates some of the major domain name suffixes assigned by interNIC to include a description of the type of organization that is assigned to each suffix.

Readers should not that the entries in indicate initially defined generic top-level domains. In addition to domains based on organizational category, top-level domains are also defined based on the use of two letter country codes from the international standards organization (ISO) 3166 standard. For example identifies the top level Australian domain whereas .for identifies the top-level French domain.

Provides a general schematic of the DNS database indicating the relationship of six top level generic domains in the DNS to each other as well as to possible subdomains. In addition, the lower left portion of the database indicates a possible registered subdomain managed by the commercial company internationals pretzel, the address of each application is formed by adding the mnemonics assigned to each node upward tot eh top level domain they reside in, separating each name from the other by a dot. Thus, what is knows as the fully qualified domain name for an FTP server would become ftp.ipretzel.com.

In examining note that each node represents a portion of a database, and the number of domains that can be nested under one another is essentially unlimited. Similar to file system directory structure, each domain can be identified both relative and absolute to its position in the domain. When a domain is identified with respect to its parent domain, relative addressing is used. In comparison, when a sequence of labels is used to identify a domain with respect to the root of the database, absolute addressing is employed.

To facilitate the use of domain, names each organization registers its name with interNIC. The registration process results in a domain name consisting of the organization's name and its domain assignment being registered. For example, returning to the international pretzel example, assume that the company is a commercial organization and registered under the name IPRETZEL. It would be assigned the registered domain name ipretzel.com. once an organization has a registered domain name, it can prefix that name to indicate specific hosts or applications residing on a host. For example www.ipretzel.com and ftp.ipretzel.com could represent a world wide web server and file transfer protocol server. They could also represent two application residing on a common server.

When you enter the destination IP addresses in an application as a name, a translation process is required to convert that name into a 32 bit IP address. To accomplish this translation process, each TCP/IP network has a server that keeps track of the hostnames on the network. When a network user enters a name instead of an IP address, the application's name revolver, functioning as a client, transmits a request to the local name server. If the name resides on the network or was recently obtained from another network and is in each memory, the server returns the IP address associated with the name. If the name is not in the name server nor in its cache memory, the server will send a request to a higher DNS server on another network, in effect, one DNS server has a pointer to another server up the inverted tree structure illustrated in . This forwarding effect can be replicated several times and traverse around the globe until the IP address associated with the English type name is found and returned.

To facilitate the location of names, each server has pointers to servers in other domain. This alleviates, for example the necessity for searching through several servers in the .edu domain to obtain the IP address for a name registered in the.gov domain.

One of the key advantages of thee hierarchical database structure of DNS is its capability to support near duplicate names, enabling, for example the international pretzel university to set up a www server using the address www.ipretzel.ude . Similarly, a women's clothing store named white house locate in the mall at shelter cove on Hilton head island could set up a WWW server using the address www.whitehous.com, and the NDS process of name to address translation would allow appropriate requests to obtain the IP address of that commercial entity separate from the government entity. Note that the near English mnemonics and names used to represent distinct host addresses play no role in the actual routing of IP packets. Instead, the DNS process translates those names into IP addresses, which are then used by routers to establish a connection to the desired destination initially defined using near English mnemonics and names.

5.10 Short Summary

Transmission control protocol and Internet protocol is commonly used to communicate across any set of interconnected networks. It was used both in the Internet context and in any (corporate) Internet context. The TCP/IP forms the base technology for a global Internet that connects homes, university campuses and other schools, corporations, and government labs in dozens of countries.

5.11 Brain Storm

1. Explain about IP addressing.
2. Explain the basic address scheme.
3. What is subnet?
4. Explain about domain name system.
5. Explain about subnet mask?

Lecture 6

Subnets

Objectives

After completing this lesson, you should be able to do the following

- Discuss about subnet mask
- Describe the subnetting concepts
- Describe about variable length subnetting.
- Describe about supernetting and classless interdomain routing.
- Discuss about Public and private addresses.

Coverage Plan

Lecture 6

6.1 Snap Shot

In this session we learn about subnet mask. Subnet masks are frequently expressed in dotted decimal notation. Once the bits are set for the network ID and host ID portion, the resulting 32-bit number is converted to dotted decimal notation.

6.2 Subnet Mask

With the advent of subnetting, one can no longer rely on the definition of the IP address classes to determine the network ID in the IP address. A new value is needed to define which part of the IP address is the network ID and which part is the host ID, regardless of whether class-based or subnetted network IDs are being used.

RFC 950 defines the use of a subnet mask (also referred to as an address mask) as a 32 bit value which used to distinguish the network ID from the host ID in an arbitrary IP address. The bits of the subnet mask are defined as:

- All bits that correspond to the network ID are set to 1.

- All bits that correspond to the host ID are set to 0.

Each host on a TCP/IP network requires a subnet mask even on a single-segment network. Either *a default subnet mask*, which is used when using class-based network IDs, or a *custom subnet mask*, which is used when subnetting or supernetting, is configured on each TCP/IP node.

Dotted Decimal Representation of Subnet Masks

Subnet masks are frequently expressed in dotted decimal notation. Once the bits are set for the network ID and host ID portion, the resulting 32-bit number is converted to dotted decimal notation. Note that even though expressed in dotted decimal notation, a subnet mask is not an IP address.

A default subnet mask is based on the IP address classes and is used on TCP/IP networks that are not divided into subnets. Below the Table lists the default subnet masks using the dotted decimal notation for the subnet mask.

Default subnet masks in dotted decimal notation

Address	Bits for Subnet Mask	Subnet Mask
Class A	11111111 00000000 00000000 00000000	255.0.0.0
Class B	11111111 11111111 00000000 00000000	255.255.0.0
Class C	11111111 11111111 11111111 00000000	255.255.255.0

Custom subnet masks are those that differ from the above default subnet masks when doing subnetting or supernetting. For edample, 138.96.58.0 is an 8-bit subnetted class B network ID. Eight bits of the class-based host ID are being used to express subnetted network ID. The subnet mask uses a total of 24 bits (255.255.255.0) to define the subnetted network ID. The subnetted network ID and its corresponding subnet mask is then expressed in dotted decimal notation as:

138.96.58.0, 255.255.255.0

Network Prefix Length Representation of Subnet Masks

Since the network ID bits must be always chosen in a contiguous fashion from the high order bits, a shorthand way of expressing a subnet mask is to denote the number of bits that define the network ID as a network prefix using the network prefix notation:/<#of bits>. Below the table lists the default subnet masks using the network prefix notation for the subnet mask.

Default subnet masks in network prefix notation

Address	Bits for Subnet Mask	Network Prefix
Class A	11111111 00000000 00000000 00000000	/18
Class B	11111111 11111111 00000000 00000000	/16
Class C	11111111 11111111 11111111 00000000	/24

For example, the class B network ID 138.96.0.0 with the subnet mask of 255.255.0.0 would be expressed in network prefix notation as 138.96.0.0/16.

As an example of a custom subnet mask, 138.96.58.0 is an 8-bit subnetted class B network ID. The subnet mask uses a total of 24 bits to define the subnetted network ID. The subnetted network ID and its corresponding subnet mask is then expressed in network prefix notation as:

138.96.58.0/24

Determining the Network ID

To extract the network ID from an arbitrary IP address using an arbitrary subnet mask, IP uses a mathematical operation called a logical AND comparison. In an AND comparison, the result of two items being compared is true only when both items being compared are true, otherwise, the result is false. Applying this principle to bits, the result is 1 when both bits being compared are 1; otherwise, the result is 0.

IP takes the 32-bit IP address and logically ANDs it with the 32-bit subnet mask. This operation is known as a bit-wise logical AND. The result of the bit-wise logical AND of the IP address and the subnet mask is the network ID.

For example, what is the network ID of the IP node 129.56.189.41 with a subnet mask of 255.255.240.0?

To obtain the result, turn both numbers into their binary equivalents and line them up. Then perform the AND operation on each bit and write down the result.

10000001 00111000 10111101 0101001 IP Address

11111111 11111111 11110000 00000000 Subnet Mask

10000001 00111000 10110000 00000000 Network ID

The result of the bit-wise logical AND of the32 bits of the IP address and the subnet mask is the network ID 129.56.176.0

6.3 Subnetting

While the conceptual notion of subnetting by utilizing host bits is straightforward, the actual mechanics of subnetting are a bit more complicated. Subnetting is a three-step procedure:

1. Determine the number of host bits to be used for the subnetting.
2. Enumerate the new subnetted network IDs.
3. Enumerate the IP addresses for each new subnetted network ID.

Step 1: Determining the Number of Host Bits

The number of host bits being used for subnetting determines the possible number of subnets and hosts per subnet. Before you choose how many host bits, you should have a good idea of the number of subnets and hosts you will have in the future. Using more bits for the subnet mask than required will save you the time of reassigning IP addresses in the future.

The more host bits that are used, the more subnets you can have – but with fewer hosts. If you use too many host bits, it will allow for growth in the number of subnets, but will limit the growth in the number of hosts. If you use too few hosts, it will allow for growth in the number of hosts, but will limit the growth in the number of subnets.

For example, below figure illustrates the subnetting of up to the first 8 host bits of a class B network ID. If we choose one host bit for subnetting, we obtain 2 subnetted network IDs with 16,382 hosts per subnetted network ID. If we choose 8 host bits for subnetting, we obtain 256 subnetted network IDs with 254 hosts per subnetted network ID.

Class B

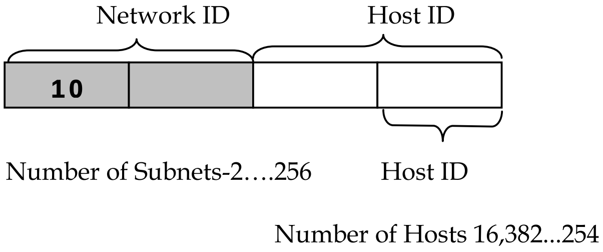

In practices, network administrators define a maximum number of nodes they want on a single network. Recall that all nodes on a single network share all the same broadcast traffic; they reside in the same broadcast domain. Therefore, growth in the amount of subnets is favored over growth in the amount of hosts per subnet.

Follow these guidelines to determine the number of host bits to use for subnetting.

1. Determine how many subnets you need and will need in the future. Each physical network is a subnet. WAN connections may also count as subnets depending on whether your routers support unnumbered connections.

2. Use additional bits for the subnet mask if;

❖ You will never require as many hosts per subnet as allowed by the remaining bits.

❖ The number of subnets will increase in the future, requiring additional host bits.

To determine the desired subnetting scheme, you will start with an existing network ID to be subnetted. The network ID to be subnetted can be a class-based network ID, a subnetted network ID, or a supernet. The existing network ID will contain a series of network ID bits which are fixed, and a series of host ID bits which are variable. Based on your requirements for the number of subnets and the number of hosts per subnet, you will choose a specific number of host bits to be used for the subnetting.

Following Table shows the subnetting of a class A network ID. Based on a required number of subnets, and a maximum number of hosts per subnet, a subnetting scheme can be chosen.

Subnetting a class A network ID

Required no of subnets	Subnet mask	No of hosts per subnet
1-2	255.128.0.0 or /9	8,388,606
3-4	255.192.0.0 or /10	4,194,302
5-8	255.224.0.0 or /11	2,097,150
9-16	255.240.0.0 or /12	1,048,574
17-32	255.248.0.0 or /13	524,286
33-64	255.252.0.0 or /14	262,142
65-128	255.254.0.0 or /15	131,070
129-256	255.255.0.0 or /16	65,534
257-512	255.255.128.0 or /17	32,766
513-1,024	255.255.192.0 or /18	16,382
1025-2,048	255.255.224.0 or /19	8,190
2,049-4,096	255.255.240.0 or /20	4,094
4,097-8,192	255.255.248.0 or /21	2,046
8,193-16,384	255.255.252.0 or /22	1,022
16,385-32,768	255.255.254.0 or /23	510
32,769-65,536	255.255.255.0 or /24	254
65,537-131,072	255.255.255.128 or /25	126
131,073-262,144	255.255.255.192 or /26	62
262,145-524,288	255.255.255.224 or /27	30
524,289-1,048,576	255.255.255.240 or /28	14
1,048,577-2,097152	255.255.255.248 or /29	6
2,097,153-4,194,304	255.255.255.252 or /30	2

Subnetting a class B network ID

Required no of subnets	Subnet mask	No of hosts per subnet
1-2	255.255.128.0 or /17	32,766
3-4	255.255.192.0 or /18	16,382
5-8	255.255.224.0 or /19	8,190
9-16	255.255.240.0 or /20	4,094
17-32	255.255.248.0 or /21	2,046
33-64	255.255.252.0 or /22	1,022
65-128	255.255.254.0 or /23	510
129-256	255.255.255.0 or /24	254
257-512	255.255.255.128 or /25	126
513-1,024	255.255.255.192 or /26	62
1,025-2,048	255.255.255.224 or /27	30
2,049-4,096	255.255.255.240 or /28	14
4,097-8,192	255.255.255.248 or /29	6
8,193-16,384	255.255.255.252 or /30	2

Subnetting a class C network ID

Required number of subnets	Subnet mask	No of hosts per subnet
1-2	255.255.255.128 or /25	126
3-4	255.255.255.192 or /26	62
5-8	255.255.255.224 or /27	30
9-16	255.255.255.240 or /28	14
17-32	255.255.255.248 or /29	6
33-64	255.255.255.252 or /30	2

Step 2: Enumerating Subnetted Network IDs

Based on the number of host bits you use for your subnetting, you must list the new subnetted network IDs. There are two main approaches:

- ❖ Binary – List all possible combinations of the host bits chosen for subnetting and covert each combination to dotted decimal notation.

- ❖ Decimal – Add a calculated increment value to each successive subnetted network ID and covert to dotted decimal notation.

Either method produces the same result – the enumerated list of subnetted network IDs.

1. Based on n, the number of host bits chosen for subnetting, create a 3-column table with 2^n entries. The first column is the subnet number, the second column is the binary representation of the subnetted network ID, and the third column is the dotted decimal representation of the subnetted network ID. For each binary representation, the bits of the network ID being subnetted are fixed to their appropriate values and the remaining host bits are set to all 0's. The host bits chosen for subnetting will vary.

2. In the first table entry, set the subnet bits to all 0's and convert to dotted decimal notation. The original network ID is subnetted with its new subnet mask.

3. In the next table entry, increase the value within the subnet bits.

4. Convert the binary result to dotted decimal notation.

5. Repeat steps 3 and 4 until the table is complete.

As an example, a 3-bit subnet of the private network ID 192.168.0.0 is needed. The subnet mask for the new subnetted network IDs is 255.255.224.0 or /19. Based on n=3, construct a table with 8 (=2^3) entries. The entry for subnet 1 is the all 0's subnet. Additional entries in the table are successive increments of the subnet bits as shown in the following Table. The host bits used for subnetting are underlined.

Subnet	Binary Representation	Subnetted network ID
1	11000000.10101000.00000000.00000000	192.168.0.0/19
2	11000000.10101000.00100000.00000000	192.168.32.0/19
3	11000000.10101000.01000000.00000000	192.168.64.0/19
4	11000000.10101000.01100000.00000000	192.168.96.0/19
5	11000000.10101000.10000000.00000000	192.168.128.0/19
6	11000000.10101000.10100000.00000000	192.168.160.0/19
7	11000000.10101000.11000000.00000000	192.168.192.0/19
8	11000000.10101000.11100000.00000000	192.168.224.0/19

Decimal Subnetting Procedure

1. based on n, the number of host bits chosen for subnetting, create a 3-column table with 2^n entries. The first column is the subnet number, the second column is the decimal representation of the 32-bit subnetted network ID, and the third column is the dotted decimal representation of the subnetted network ID.

2. Convert the network ID (w.x.y.z) being subnetted from dotted decimal notation to N, a decimal representation of the 32-bit network ID.
   N=W* 16777246 + X* 65536 + Y* 256 + Z

3. Compute the increment value I based on h, the number of host bits remaining.
   $I = 2^h$

4. In the first table entry, the decimal representation of the subnetted network ID is N and the subnetted network ID will be w.x.y.z. with its new subnet mask.

5. In the next table entry, add I to the previous table entry's decimal representation.

6. Convert the decimal representation of the subnetted network ID to dotted decimal notation (W.X.Y.Z) through the following formula (where s is the decimal representation of the subnetted network ID):

   W=INT (s/16777216)
   X=INT ((s mod (16777216))/65536)
   Y=INT ((s mod (65536))/256)
   Z=s mod(256)

INT( ) denoted integer division, mod( ) denotes the modulus, the remainder upon division.

7. Repeat steps 5 and 6 until the table is complete.

As an example, a 3-bit subnet of the private network ID 192.168.0.0 is needed. Based on n = 3 we construct a table with 8 entries. The entry for subnet 1 is the all 0's subnet. N, the decimal representation of 192.168.0.0, is 3232235520, the result of 192*16777216 + 168*65536. Since there are 13 host bits remaining, the increment I is 2^{13} =8192. Additional entries in the table are successive increments of 8192 as shown in following Table.

Subnet	Decimal Representation	Subnetted network ID
1	3232235520	192.168.0.0/19
2	3232243712	192.168.32.0/19
3	3232251904	192.168.64.0/19
4	3232260096	192.168.96.0/19
5	3232268288	192.168.128.0/19
6	3232276480	192.168.160.0/19
7	3232284672	192.168.192.0/19
8	3232292864	192.168.224.0/19

The All-Zeros and All-Ones Subnets

RFC 950 originally forbade the use of the subnetted network IDs where the bits being used for subnetting are set to all 0's and all 1's (the all-ones subnet). The all-zeros subnet caused problems for early routing protocols and the all-ones subnet conflicts with a special broadcast address called the all-subnets directed broadcast address.

However, RFC 1812 now permits the use of the all-zeros and all-ones subnets in a Classless Interdomain Routing (CIDR) – compliant environment. CIDR – complaint environments use modern routing protocols which do not have a problem with the all-zeros subnet and the use of the all-subnets directed broadcast has been deprecated.

Before you use the all-zeros and all-ones subnets, verify that they are supported by your hosts and routers. Windows NT supports the use of the all-zeros and all-ones subnets.

Step 3: Enumerating IP Addresses for Each Subnetted Network ID

Based on the enumeration of the subnetted network IDs, you must now list the valid IP addresses for new subnetted network IDs. To list each IP address individually would be too tedious. Instead, we will enumerate the IP addresses for each subnetted network ID by defining the range of IP addresses for each subnetted network ID. There are two main approaches:

- Binary – Write down the first and last IP address for each subnetted network ID and convert to dotted decimal notation.
- Decimal – Add values incrementally, corresponding to the first and last IP addresses for each subnetted network ID and convert to dotted decimal notation.

Either method produces the same result – the range of IP addresses for each subnetted network ID.

Binary Procedure

1. Based on n, the number of host bits chosen for subnetting, create a 3-column table with 2^n entries. Alternately, add two columns to the previous table used for enumerating the subnetted network IDs. The first column is the subnet number, the second column is the binary representation of the first and last IP address for the subnetted network ID, and the third column is the dotted decimal representation of the first and last IP address of the subnetted network ID.

2. For each binary representation, the first IP address is the address where all the host bits are set to 0 except for the last host bit. The last IP address is the address where all the host bits are set to 1 except for the last host bit.

3. Covert the binary representation to dotted decimal notation.

4. Repeat steps 2 and 3 until the table is complete.

As an example, the range of IP addresses for the 3 bit subnetting of 192.168.0.0 is shown in the following table. The bits used for subnetting are underlined.

Subnet	Binary representation	Range of IP addresses
1	11000000.10101000.00000000.00000001 – 11000000.10101000.00011111.11111110	192.168.0.1- 192.168.31.254
2	11000000.10101000.00100000.00000001- 11000000.10101000.00111111.11111110	192.168.32.1- 192.168.63.254
3	11000000.10101000.01000000.00000001-	192.168.64.1-

	11000000.10101000.01011111.11111110	192.168.95.254
4	11000000.10101000.01100000.00000001- 11000000.10101000.01111111.11111110	192.168.96.1- 192.168.127.254
5	11000000.10101000.10000000.00000001- 11000000.10101000.10011111.11111110	192.168.128.1- 192.168.159.254
6	11000000.10101000.10100000.00000001- 11000000.10101000.10111111.11111110	192.168.160.1- 192.168.191.254
7	11000000.10101000.11000000.00000001- 11000000.10101000.11011111.11111110	192.168.192.1- 192.168.223.254
8	11000000.10101000.11100000.00000001- 11000000.10101000.11111111.11111110	192.168.224.1- 192.168.2255.254

Decimal Procedure

1. Based on n_1 the number of host bits chosen for subnetting, create a 3-column table with 2^n entries. Alternately, add two columns to the previous table used for enumerating the subnetted network IDs. The first column is the subnet number, the second column is the decimal representation of the first and last IP address for the subnetted network ID, and the third column is the dotted decimal representation of the first and last IP address of the subnetted network ID.

2. Compute the increment value J base on h, the number of host bits remaining.

   $J= 2^h – 2$

3. For each decimal representation, the first IP address is N + 1 where N is the decimal representation of the subnetted network ID. The last IP address is N + J.

4. Convert the decimal representation of the first and last IP addresses to dotted decimal notation (W.X.Y.Z) through the following formula (where is the decimal representation of the first or last IP address):

   W=INT (s/16777216)
   X=INT ((s mod (16777216))/65536)
   Y=INT ((s mod (65536))/256)
   Z=s mod(256)

INT ( ) denotes integer division, mod( ) denotes the modulus, the remainder upon division.

5. Repeat steps 3 and 4 until the table is complete.

As an example, the range of IP addresses for the 3 bit subnetting of 192.168.0.0. The increment j is 2^13-2 = 8190.

Decimal enumeration of IP addresses

Subnet	Decimal representation	Range of IP addresses
1	3232235521-3232243710	192.168.0.1-192.168.31.254
2	3232243713-3232251902	192.168.32.1-192.168.63.254
3	3232251905-3232260094	192.168.64.1-192.168.95.254
4	3232260097-3232268286	192.168.96.1-192.168.127.254
5	3232268289-3232276478	192.168.128.1-192.168.159.254
6	3232276481-3232284670	192.168.160.1-192.168.191.254
7	3232284673-3232292862	192.168.192.1-192.168.223.254
8	3232292865-3232301054	192.168.224.1-192.168.255.254

6.4 Variable Length Subnetting

One of the original uses for subnetting was to subdivides a class-based network ID into a series of equal-sized subnets. For example , a 4-bit subnetting of a class B network ID produced 16 equal-sized subnets. However, subnetting is a general method of utilizing host bits to express subnets and does not require equal sized subnets.

Subnets of different size can exist within a class based network ID. This is well suited to real world environments, where networks of an organization contain different amounts of hosts, and different sized subnets are needed to minimize the wasting of IP addresses. The creation and deployment of various sized subnets of a network ID is knows as variable length subnetting and uses variable length subnet masks .

Variable length subnetting is a technique of allocating subnetted network IDs that use subnet masks of different sizes. However, all subnetted network IDs are unique and can be distinguished from each other by their corresponding subnet mask.

The mechanics of variable length subnetting are essentially that of performing subnetting on a previously subnetted network ID. When subnetting the network bits are fixed and a certain amount of hosts bits are chosen to express subnets. With variable length subnetting, the network Id being subnetted has already been subnetted.

Variable length subnetting example

For example, given the class based network ID of 135.41.0.0/16, a required configuration is to reserve half of the addresses for future use create 15 subnet with up to 2,000 hosts, and 8 subnets with up to 250 hosts.

Reserve half of the addresses for future use

To reserve half of the addresses for future use, a 1 bit subnetting of the class based network ID of 135.41.0.0 is done producing 2 subnets, 135.41.0.0/17 and 135.41.128.0/17. the subnet 135.41.0.0/17 is chosen as the portion of the addresses which are reserved for future use.

To achieve a requirement of 15 subnets with approximately 2,00 hosts, a 4 bit subnetting of the subnetted network IF of 135.41.128.0/17 us done. This produces 16 subnets (135.41.128.0/21) allowing up to 2,046 hosts per subnet. The first 15 subnetted network ids(135.41.128.0/121 to 135.41.240.0/21) are chosen as the network IDs, which fulfills the requirement.

6.5 Supernetting and Classes Interdomain Routing

With the recent growth of the internet, it became clear to the internet authorities that the class B network IDs would soon be depleted. For most organizations, a class C network ID does not contain enough host IDs and a class B network ID has enough bits to provide a flexible subnetting scheme within the organization.

The internet authorities devised a new methods of assigning network IDs to prevent the depletion of class B network, IDs. Rather than assigning a class B network ID, the internet network information center assigns a range of class C network IDs that contain enough network and host IDs for the organization's needs. This is knows as supernetting. For example rather than allocating a class B network ID to an organization that has up to 2,000 hosts, the interNIC allocates a range of 8 class C network IDs. Each class C network ID accommodates 254 hosts, for a total of 2,032 host IDs.

While this technique helps conserve class B network IDs it crates a new problem. Using conventional routing techniques, the routers on the internet now must have 8 class C network ID entries in their routing tables to route IP packets to the organization. To prevent internet routers from becoming overwhelmed with routes, a technique called *classless interdomain routing* is used to collapse multiple network ID entries into a single entry corresponding to all of the class C network IDs allocated to that organization.

Conceptually CIDR creates the routing table entry (starting Network ID, count,) where starting network ID is the first class C network ID and the count is the number of class C network IDs allocated. In practice, a supernetted subnet mask is used to convey the same information. To express the situation where 8 class C network IDs are allocated staring with networking ID 220.78.168.0:

In network prefix notation, the CIDR entry is 220.78.168.0/21.

A block of addresses using CIDR is knows as a CIDR block.

In order to support CIDR routers must be able to exchange routing information in the form of(network ID subnet mask) pairs. RIP for IP version 2, OSPF, and BGPV4 are routing protocols that support CIDR. RIP for IP version 1 does not support CIDR.

The address space perspective

The use of CIDR to allocate addresses promotes a new perspective on IP network IDs. In the above example the CIDR block (220.78.168.0,255.255.248.0) can be thought of in two ways.

- A block of 8 class C network IDs.
- An address space in which 21 bits are fixed and 11 bits are assignable.

In the latter perspective, IP network IDs loss their class based heritage and become separate IP address spaces, subnets of the original IP address space defined by the 32 bit address. Each IP network ID (class based, subnetted, CIDR block) is an address space in which certain bits are fixed (the network ID bits) and certain bits are variable (the host bits). The host bits are assignable as host IDs or using subnetting techniques, can be used in whatever manner best suits the needs of the organization.

6.6 Public and Private Addresses

If you intranet is not connected to the internet, any IP addressing can be deployed. If direct or indirect connectivity to the internet is desired, then there are two types of addresses employed on the internet, public addresses and private addresses.

Public addresses

Public addresses are assigned by interNIC and consist of class based network IDs or blocks of CIDR based address that are guaranteed to be globally unique to the internet.

When t he public addresses are assigned, routes are programmed into the routers of the internet so that traffic to the assigned public addresses can reach their locations. Traffic to destination public addresses are reachable on the internet.

For example when an organization is assigned a CIDR block in the form of a network ID and subnet mask, that (network ID, subnet mask) pair also exists as a route in the routers of the internet. IP packets destined to an address within the CIDR block are routed to the proper destination.

Illegal addresses

Private intranets that have no intent on connecting to the internet can choose ay addresses they want even public addresses that have been assigned by the interNIC. If an organization later decides to connect to the internet, its current address scheme may include addresses already assigned by the InterNIC to other organizations. These addresses would be duplicate or conflicting addresses and are known as illegal addresses. Connectivity from illegal addresses to internet locations is not possible.

For example a private organization chooses to use 207.46.130.0/24 as its intranet address space. The public address space 207.46.130.0/24 has been assigned to the Microsoft corporation and routes exists on the internet routers to route all packets destined to IP addresses on 207.46.130.0/24 to Microsoft routers. As long as the private organization does not connect to the internet there is no problem, since the two address spaces are on separate IP internetworks. If the private organization then connected directly to the internet and continued to use 207.46.130.0/24 as its address space, then any internet response traffic to locations on the 207.46.130.0/24 network would be routed to Microsoft routers, not to the routers of the private organization.

Private addresses

Each IP node requires an IP address that is globally unique to the IP internetwork. In the case of the Internet, each IP node on a network connected to the Internet requires an IP address that is globally unique to the Internet. As the Internet grew, organizations connecting to the Internet required a public address for each node on their intranets. This requirement placed a huge demand on the pool of available public addresses.

When analyzing the addressing needs of organizations, the designers of the Internet noted that for many organizations most of the hosts on the organization's intranet did not require direct connectivity to internet hosts. Those hosts that did require a specific set of internet services, such as the World wide web access and e-mail typically access the internet services through application layer gateways such as proxy servers and e-mail servers. The result is that most organizations only required a small amount of public addresses for those nodes (such as proxies , routers, firewalls, and translators) that were directly connected to the internet.

For the hosts within the organization that do not require direct access to the internet, IP address that do not duplicate already assigned public addresses are required. To solve this addressing problem, the internet designers reserved a portion of the IP address space and named this space the private address space. An IP address in the private address space is never assigned as a public address. IP addresses within the private address space are known as private addresses. Because the public and private address spaces do not overlap, private addresses never duplicate public addresses.

The private address space specified in RFC 1597 is defined by the following three address blocks:

- 10.0.0.0/8

The 10.0.0/8 private network is a class A network ID that allows the following range of valid IP addresses. 10.0.0.1to 10.255.255.254. the 10.0.0./8 private network has 24 hours host bits which can be used for any subnetting scheme within the private organization.

- 172.16.0.0/12

The 172.16.0.0/12 private network can be interpreted either as a block of 16 class B network IDs or as a 20 bits assignable address space which can be used for any subnetting scheme withizn the private organization. The 172.16.0.0/12 private network allows the following range of valued IP addresses 172.16.0.1 to 172.31.255.254.

- 192.168.0.0/16

The 192.168.0.0/16 private network can be interpreted either as a block of 256 class C network IDs or as a 16 bit assignable address space which can be used for any subnetting scheme within the private organization. The 192.168.0.1 to 192.168.255.254.

The result of many organizations using private addresses is that the private address space is re used helping to prevent the depletion of public addresses.

Since the IP addresses in the private address space will never be assigned by the InterNIC as public addresses, there will never exist routes in the Internet routers for private addresses. Traffic to destination private addresses are not reachable on the internet. Therefore, internet traffic from a host that has a private must either send its requests to an application layer gateway (such as a proxy server), which has a valued public addresses, or have its private address translated into a valued public address by a network address translator before it is sent on the internet.

6.7 Short Summary

One of the original uses for subnetting was to subdivides a class-based network ID into a series of equal-sized subnets. For example , a 4-bit subnetting of a class B network ID produced 16 equal-sized subnets. However, subnetting is a general method of utilizing host bits to express subnets and does not require equal sized subnets.

6.8 Brain Storm

1. Explain about subnet.
2. Explain about subnet mask.
3. What is variable length subnetting
4. What are private and public addresses?

Lecture 7

TCP / IP Operation & Applications

Objectives

After completing this lesson, you should be able to do the following

- Discuss about TCP/IP operation & applications
- Describe the Internet protocols
- Describe about the IP header.
- Describe about address resolution.
- Discuss about TCP source and destination port
- Describe about file transfer protocol
- Discuss about TELNET and Hyper Text Transfer Protocol.

Coverage Plan

Lecture 7

7.1 Snap Shot

In this session we discuss about the operation and application of TCP. And also we discuss about other various protocols ftp, telnet and http. Internet protocol is one of the important protocols in network layer. IP provides for the transfer of a basic unit of information referred to as a datagram.

7.2 Internet Protocol

IP is a network layer protocol. IP provides for the transfer of a basic unit of information referred to as a datagram. In doing so, IP operates as an unreliable, connectionless protocol. Although your first reaction when reading those terms is to have heartburn when considering the use of IP, they are not as bad as the words imply. First, although IP provides an unreliable transmission method, the term should be viewed in the context that delivery is not guaranteed. This means that queuing delays or other problems can result in the loss of data; however, higher layers in the protocol suite, such as TCP, can provide error detection and correction, which results in the retransmission of IP datagrams. Here the term datagram is used to represent a unit of data, or a portion of a single message. Second the term connectionless references the fact that each datagram is treated independently from preceding and succeeding datagrams. This means that IP does not require a connection to be established between source and destination prior to transferring the first datagram or succeeding datagrams.

The routing of datagrams through a network can occur over different paths, with some datagrams arriving out of sequence from the order transmitted. In addition, as datagrams flow between networks, they encounter physical limitations imposed on the amount of data that can be transported based on the transport mechanism used to move data on the network. For example, the information field in an Ethernet frame is limited to 1,500 bytes. Thus, as a datagram flows between networks, it may have to be fragmented into two or more datagrams to be transported through different networks to their ultimate destination. For example, consider the transfer of a 12,000-byte file from a file server, connected to a token ring network, to a workstation connected to an Ethernet LAN via a pair of routers providing a connection between the two local area networks. A 4MBPS token ring LAN can support a maximum size of the Information field of 4,500 bytes in a frame, whereas the maximum size of the Information field in an Ethernet frame is 1,500 bytes. In addition, depending on the protocol used on the wide area network connection between routers, the information field might be limited to between 512 and 1,024 byte, Thus, the IP protocol must break up the file transfer into a series of datagrams whose size is acceptable for transmission between networks. Upon receipt at the destination, each datagram must be put back into its correct sequence so that the file can be correctly reformed.

The routing of two datagrams from workstation A on a token ring network to server B connected to an Ethernet LAN. Since the routing of datagrams is a connectionless service, no call setup is required, which enhances transmission efficiency. In comparison, when TCP is used, it provides a connection-oriented service regardless of 2the lower-layer delivery system.

TCP requires the establishment of a virtual circuit in which a temporary path is developed between source and destination, this path is fixed, and the flow of datagrams is restricted to the path established. When the user datagram protocol (UDP), a different layer 4 protocol in the TCP /IP protocol suite, is used in place of TCP, the flow of data at the transport layer continues to be connectionless and results in the transport of datagrams over available paths, rather then a fixed path resulting from the establishment of a virtual circuit.

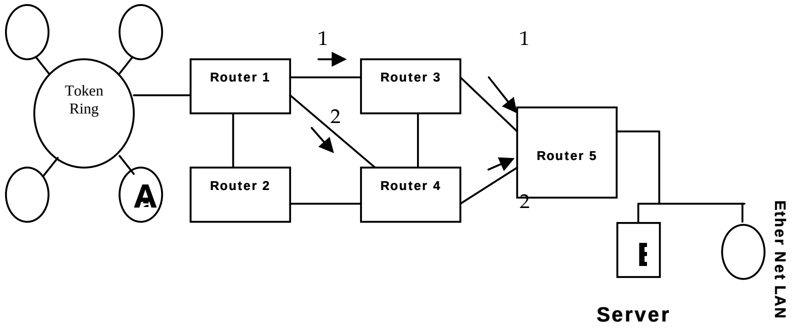

Routing of datagrams can occur over different paths

The actual division of a message into datagrams is the responsibility of the layer 4 protocol either TCP or UDP. In addition when the TCP protocol is used, that protocol is responsible for reassembling the datagrams at their destination as well as for requesting the retransmission of lost datagrams. In comparison, IP is responsible for the routing of individual datagrams from source to destination. When UDP is used as the layer 4 protocol, there is no provision for the retransmission of lost or garbled datagrams. Note that this is not necessarily a bad situation because the applications that use UDP then become responsible for managing communications.

The relationship of an IP datagram, a UDP datagram, and a TCP segment to a LAN frame. The headers represent a group of bytes added to the beginning of a datagram to allow a degree of control over the datagram. For example, the TCP header will contain information that allows this layer 4 protocol to track the sequence of the delivery of datagrams so that they can be placed into their correct order if they arrive out of sequence. We can obtain an appreciation for how datagram fragmentation is accomplished, as well as additional information about the flow of data, by examining the composition of the fields in the IP header.

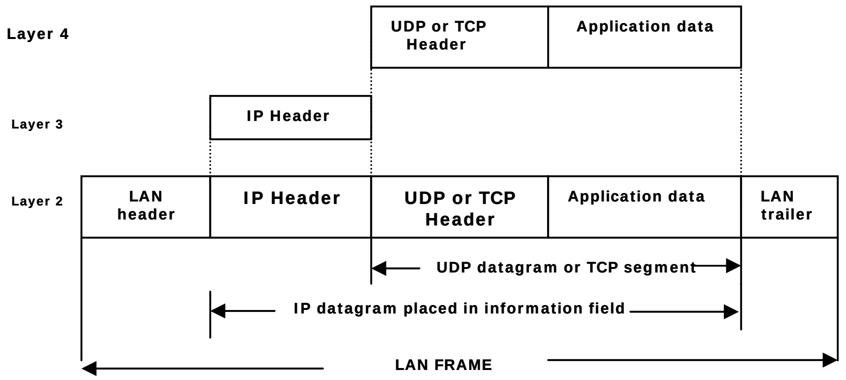

Forming a LAN frame

7.4 The IP Header

The fields contained within the IP header. Note that the header contains a minimum of 20 bytes of control information, with the width of each field with respect to a 32-bit word. Also note that since it is the job of IP to find a route for each datagram, it does not care about the contents of the datagram or the contents of the TCP header. Thus, the IP header primarily consists of source and destination 32 bit addresses, a field that protects the header, and other information primary relevant to the routing of datagrams that we will shortly cover. So lets examine the function of the fields in the IP header.

0 4 8 16 31

<table>
<tr><td>Vers</td><td colspan="2">Hlen</td><td>Service type</td><td colspan="3">Total Length</td></tr>
<tr><td colspan="4">Identification</td><td>Flags</td><td colspan="2">Fragment offset</td></tr>
<tr><td colspan="2">Time to live</td><td colspan="2">Protocol</td><td colspan="2">Header</td><td>Checksum</td></tr>
<tr><td colspan="7">Source IP address</td></tr>
<tr><td colspan="7">Destination IP address</td></tr>
<tr><td colspan="7">Options + Padding</td></tr>
</table>

The format of the IP header

The Vers Field

The Vers field consists of 4 bit that identify the version of the IP protocol used to create the datagram. The current version of the IP protocol is 4.

The Hlen and total length fields

The Hlen field, which follows the Vers fields, also contains 4 bits. This field indicates the length of the header in 32 bit words. In comparison, the total length field indicates the total length of the datagram to include in its header and higher layer information. Because 16 bits are used for this filed, an IP datagram can be up to 2^16 or 65,535 octets in length.

The service type field

The service type field defines how the datagram is handled. Three of the 8 bits in this field are used to denote the precedence or level of importance assigned by the sender. Therefore, this field provides a priority mechanism for routing IP datagrams.

The identification and fragment off set fields

The identification field enables each datagram or fragmented datagram to be identified. If a datagram was fragmented, the fragment offset field specifies the offset in the original datagram of the data being carried. In effect, this field indicates where the fragment belongs in the complete message. The actual value in this field is an integer, which corresponds to a unit of 8 octets, providing offset in 64 bit units.

The time to live field

The time to live (TTL) fields specifies the maximum time that a datagram can live. Since an exact time is difficult to measure, many routers decrement this field by 1 as a datagram flow between networks, with the datagram being discarded when the field value reaches 0. You can consider this field to represent a fail-safe mechanism because it prevents misaddressed datagrams from continuously wandering the Internet.

The flags field

The flags field contains 2 bits that indicate how fragmentation occurs while a their bit is currently unassigned. The setting of 1 bit can be viewed as a direct fragment control mechanism, as a value of 0 indicates the datagram can be fragmented, while a value of 1 denotes not to fragment. The second bit is set to 0 to indicate that a fragment in a datagram is the last fragment and set to a value of 1 to indicate that more fragments follow the current protocol.

The protocol field

The protocol field specifies the higher-level protocol used to create the message carried in the datagram. For example, a value of decimal 6 would indicate TCP, while a value of decimal 17 would indicate UDP.

The source and destination address fields

The source and destination address fields are both 32 bits in length. As previously discussed each address represents both a network and a host computer on the network.

7.5 Address Resolution

The physical address associated with a LAN workstation is referred to as its hardware address. For an Ethernet network, that address is 48 bits, or 6 bytes, in length. At the data link layer, IP uses a 32 bit logical address. One common problem associated with the routing of an IP datagram to a particular workstation of a network is the delivery of the datagram to its correct destination; to correctly deliver the datagram requires knowledge of the relationship between physical and logical addresses. This relationship is obtained by two protocols that map the physical and logical address to each other. One protocol, known as the address resolution protocol translates an IP address into a hardware address. The reverse address resolution protocol as its name implies, performs a reverse mapping converting a hardware address into an IP address.

To illustrate the use of ARP, assume that one computer user wants to send a datagram to another compute and both computers are located on the same Ethernet frame to all devices on the LAN. That packet would contain the destination IP address, since it is knows, setting the hardware address field to zeros because its value is unknown. Although each device on the LAN will read the ARP packet as it is transmitted as a broadcast packet, only the device that recognizes its own logical physical address is inserted, in the ARP address filed previously set to 0. To reduce the necessity to constantly transmit ARP packets, the originator will record received information in a table known as an ARP cache, allowing subsequent datagrams to be quickly directed to the appropriate address on the LAN. Thus, ARP provides a well-though-out methodology for equating physical hardware addresses to IPs logical addresses and allows IP addressing at layer 3 to occur independently from LAN addressing at layer 2. Now that you have an appreciation for IP and the method by which IP addresses are equated to media Access Control layer 2 hardware addresses, let's focus our attention to layer 4 of TCP /IP Protocol suite.

7.6 TCP Source and Destination Port

TCP

The Transmission Control Protocol is a layer 4 connection-oriented protocol. This protocol is responsible for providing reliable communications between hosts and processes on different hosts. TCP is structured to permit multiple application programs on a host communicate concurrently with processes on other hosts, as well as for a host to de-multiplex and service incoming TCP traffic among different applications or processes running on the host. To accomplish, this TCP header includes a destination port number to identify the ultimate destination in a computer. To obtain an appreciation for the functionality and capability of TCP, let's turn out attention to its header.

<table>
<tr><td colspan="3">Source Port</td><td>Destination Port</td></tr>
<tr><td colspan="4">Sequence port</td></tr>
<tr><td colspan="4">Acknowledgement number</td></tr>
<tr><td>Hlen</td><td>Reserved</td><td>Code bits</td><td>Window</td></tr>
<tr><td colspan="3">Checksum</td><td>Urgent pointer</td></tr>
<tr><td colspan="4">Option + padding</td></tr>
</table>

TCP header

The Source port and Destination port fields

The source port and destination port fields are each 16 length and are used to identify a user process or application. The source port field is optional and is padded with zeros when not used. The term well knows port, which is commonly used to denote an application layer protocol or process, actually references the port address used within a TCP/IP header. The well knows ports associated with weight popular TCP/IP application layer protocols, such as FTP use two port addresses or logical connections. One address is for commands and replies and functions as a control path. The second port address is used for the actual file transfer.

Name	Abbreviation	Description	Well-known port
Domain Name Protocol	DOMAIN	Defines the DNS	53
File transfer Protocol	FTP	Provides file transfer between computers	20,21
Finger Protocol	FINGER	Provides information about specified user	79
Hypertext Transfer protocol	HTTP	Conveys information between a web browser and a web server	80
Post Office Protocol	POP	Enables PC users to access mail from a mail server	110
Simple Mail Transfer Protocol	SMTP	Provides electronic mail transfer	25
Simple Network management protocol	SNMP	Provides the exchange of management information	161,162
Telnet Protocol	TELNET	Provides remote terminal access to a host	23

Using Addresses

If you re-examine in conjunction with you will not that the user of a TCP/IP application requires three addresses at both source and destination. A port address is required to identify the process or application and is contained within the TCP header. Within the IP header, an IP address is required to identify the network and host computer where the process or applications resides. Finally the delivery of information on a LAN require the user of a hardware address, which is used within the LAN header to deliver the IP datagram. Of these addresses, routers commonly use two and firewalls to enable or bar access to predefined applications on a network. Those addresses are the IP address, either source or destination, and sometimes both, and the destination well know port.

When a private network is connected to the Internet many organizations prefer to limit access from the Internet to their internal computational facilities. To obtain a degree of security, some organizations user the filtering capability of routers to control access, while other organizations install a firewall to obtain an additional level of protection.

Most routers and firewalls operate on a set of rules that follow the basic cardinal that anything not explicitly permitted is denied. This means that unless you configure your organization router or firewall to allow access to hosts and applications on hosts, users on one side of the router or firewall. For example, assume that your organization decided to establish an FTP server on an internal corporate network and wanted members if the internet community to be able to access the server. Assuming that the FTP server's IP address is you 123.45.67.8, you would program the router or firewall to enable in bound traffic from any IP address to IP address 123.45.67.8, on port 20. Note that port 21 is used for commands and responses while port 20 is used for the actual file transfer process. This explains why your organization's router or firewall must be configured to allows traffic using ports 20 and 21.

Returning to the examination of the TCP/IP header, the sequence number field identifies the position in the sender's byte stream of the data transported in the TCP segment. Thus, this filed provides a mechanism to maintain the sequential nature of the data stream.

The Acknowledgement number filed

The acknowledgement number field identifies the number of the octet that the source expects to receive next. Thus, the acknowledgement number verifies the receipt of the prior segment when the sequence number is n.

The Hlen and code bits fields

The Hlen field is 4 bits in length and denotes the length of the segment header in 32 bit multiples. The code bits field contains 6 flag bits. Some of those bits, when set, indicate that a specific field in the header is significant and the field value should be interpreted, while other bits are used to control the connection and data transfer operation.

Code Bit	Code bit field Use
URG	Urgent pointer field significant
ACK	Acknowledgement field significant
PSH	Push function
RST	Reset the connection
SYN	Synchronize sequence numbers
FIN	No more data from sender

Code bit values

The window files

The 16-bit window filed indicates the number of octets, beginning with the one in the acknowledgement filed that the originator of the segment can control. Since TCP represents a full duplex communications path, each end of the path can use the window fields to control the quantity of data being sent to it. This provides the recipient with the ability to, in effect, have a say over its destiny. That is, if it becomes overloaded with processing or another reason results in its inability to receive large chunks of data, it can user the window field to reduce the size of the chunks of data being sent to it.

The checksum field

The checksum field provides reliability for the TCP header, the IP header, and data carried in the segment. Thus, this field provides the mechanism for the detection of errors in the segment.

The urgent pointer field

Other than some options beyond the scope of this book, the urgent pointer field completes the header. This field enables the position of urgent data within a TCP segment to be identified, and a value in the field is interpreted only when the previously mentioned URG bit is set. When that bit position is set, the value in the urgent pointer filed indicates the beginning of routine data.

To understand the interrelated role of the sequence, acknowledgement, and windows fields lets examine the transmission of data between two hosts via the user of a time chart that indicates some values for each field. We will assume that window size of 8 segments is in use. Although TCP supports full duplex transmission, for simplicity of illustration we will user a half-duplex model in the time chart.

Assuming that host A is downloading a program or performing a lengthy file transfer, the first series of segments will have sequence numbers 64 through 71. Assuming no errors occurred, host B returns an ack value of 72 to indicate the next segment it expects to receive. Lets also assume that host B is running out of buffer space and reduces the window size to 4. Thus, host A uses the window field value in the TCP header sent to it and reduces the number of units of data it will transmit to four, using an initial SEQ field value of 72 and increasing that value by 1 to 75 as it transmits four units of data to host B. Assuming all data is received error free, host B then returns an ACK value of 76 to acknowledge receipt of sequence field number though 75.

Once again, host A transmits to host B, this time using Sequence field values of 75 to 79. However, as this transmission occurs, lets assume that some type of transmission impairment occurs that sends the data into the proverbial bit bucket so that it is never received at host B. Since host B does not receive anything, it does not transmit anything's back to host A. Although host A could wait forever, this would not be a good idea when data becomes lost. Instead, an internal timer clicks down to zero while host A waits for a response. When one does not appear and the timer expires, host retransmits the segment, which is then acknowledged at the bottom.

The altering of the windows filed values provides a " sliding window" that can be used to control the flow of information by adjusting the value of the window field. In doing so, there are two special field values 0 and 1 that further control the flow of information. A windows field value of 0 means that a host has shut down communications, whereas a window value of 1 requires an acknowledgement for each unit of data transmitted. Now that we have a basic understanding of IP and TCP, let’s turn our attention to a few of the applications that use the TCP/IP suite.

7.7 FTP

The file transfer protocol was developed as a mechanism to facilitate the transfer of files between computers. FTP uses two well-known ports. Port 21 for passing control information and port 20 for the actual data transfer.

FTP supports approximately 20 commands. Those commands enable a user to change directories, obtain a directory list (dir), initiate a file transfer (get)., or tranfer a file(put). FTP permits multiple file transfers with the mget and mput commands when used with a filename containing one or more wildcard characters. For example, the command mget *.gif would result in the transfer of all files in the current directory that have the extension .gif.

One of the key advantages associated with the use of FTP is that various implementations exist that operate on a range of computers, from DOS PCs to Pentium-based windows NT servers to IBM mainframes. This enables FTP to provide a mechanism to exchange files between computers as long as both computers support FTP and can be reached via a TCP/IP connection. Connecting the latter, FTP relies on TCP at the transport layer to provide a reliable transmission path, ensuring the error-free arrival of data at its destination.

Commonly used FTP commands. Because FTP was originally developed as a command driven application,, users had to remember command names and any associated parameter values to effectively use the application.

Windows NT server includes an FTP server software module that enables your server to support file transfers using the Internet file transfer protocol. You can configure you NT server to support anonymous users, which in effect means that any user with communications access to your NT server will be able to view you server’s directory structure and may be able, if permission is granted during the FTP setup process, to read and write toe ach physical and logical disk. As an alternative, you can configure the FTP server module to support only users with valid NT server accounts. However when performing the latter, it should be noted that the FTP protocol does not provide a mechanism for encrypting passwords sent between client and server. Thus, the user of FTP in a non-anonymous mode could compromise user account password especially if they flow over the Internet.

Windows NT workstation contains a command driven FTP client. The execution of that client and the list of commands it supports are illustrated in. In that illustration, the help command was first used to display commands supported by the FTP client. Next, two additional help commands were issued, followed by specific FTP commands to obtain a description of those commands.

Although you can easily use the command based FTP client to effect file transfers, many prefer to use graphical user interface based programs developed by a number of third party vendors. Those programs facilitate the user of FTP because they hide the necessity to know FTP commands through the use of a graphical user interface, illustrates the net manage chameleon FTP client program screen after the program's connect menu was used to initiate a connection to a computer named webserver operating windows NT server with its FTP module activated. Note that the FTP server prompted the client software program for the entry of an FTP password, and the chameleon program automatically generated a dialog box labeled FTP password, into which the author was entering his password.

Although many servers with an anonymous or "guest" account do not require a user to enter a password other servers may request a user to enter his or her e-mail address. To provide a level of security and prevent the potential distribution of viruses, many organizations that operate anonymous FTP guest accounts restrict data transfer an prohibit public users from downloading files into the server.

When examining the background of note that the screen is subdivided into three sections labeled local, transfer, and remote, the latter references the right portion of the screen whose label is hidden by the display of the FTP password dialog box. The local section provides a mechanism to display the directory structure to include files in each directory on the local or client computer. The transfers section includes a series of button that enable you to perform predefined operations by simply pointing and clicking a button. At the top of the transfer section are two buttons associated with the method of file transfer, with the button associated with binary shown selected. Thus clicking either button is the equivalent of entering an ASCII or binary FTP command. The two columns of buttons in the transfer section are used to perform predefined operations on the local host(button arrow pointing to left) or remote host(button arrow pointing to right) with the capability to perform an operation indicated by the highlighting of the button. At the time the screen was printed, this author was in the process of establishing a connection to a remote computer. Thus, no operation in the transfer section could be performed with respect to the remote computer, and each button in the right column of the transfer section was not highlighted. However, once you establish a connection to the remote host, all of the buttons in the transfer section will be highlighted. However, once you establish a connection to the remote host, all of the buttons in the transfer section will be highlighted. Note that by simply entering a filename in the local directory section of the screen and clicking on the left arrow on button associated with "copy" one can initiate a file transfer from the remote to the local computer. No muss, no fuss, and non-need to remember FTP command, with the later a key attribute associated with the user of different GUI TCP/IP based application programs.

7.8 Telnet

Telnet is an interactive remove access terminal protocol that was developed to allow users to access a remote computer as if they were directly connected to the computer. Similar to FTP Telnet is based on a client/server model. Although Telnet was developed to provide terminal connectivity to hosts, the client does not have to be a physical terminal, such as one of the popular digital equipment corporation's VT products. Instead, Telnet client programs have been developed that turn PCs, Macintoshes, and even a variety of IBM, Sun, and HP workstations into interactive terminals.

As previously indicated in Telnet uses a common TCP connection to transmit both data and control information with the data flow occurring on well knows port number 23. Similar to other TCP/IP applications software developers over the past few years concentrated their development efforts on GUI based applications to include Telnet.

Currently there are wide range of Telnet client products available for user, including one built into Windows NT. Each client requires you to specify a host address, type of terminal to emulate, and part number. Concerning the port number, although 23 is the default port used for Telnet communications; some organizations select another port number for user by Telnet. Often, the selection of a much higher port number is used to "hide" the presence of a Telnet server from curious wanderers of the Internet.

Although Microsoft does not presently include a Telnet server support capability in windows NT, several third party vendors provide that capability. A few examples of third party windows NT Telnet server software include pragma systems, Inc inter access Telnet server (http: / /www.ccsi.com:80/pragma/) and SLnet Telnet Server (http:// www.seattlelab.com/oridsInet.htm).

7.9 HTTP

Perhaps the most popular Internet TCP/IP applications are Hypertext Transfer Protocol (HTTP), which is used to convey information between a web browser and a web server. Here, the term browser represents a software product that uses HTTP to transport information and supports the display of information encoded using Hypertext Markup Language.

A key reason for the growth in the use of the Internet is the growth in the World Wide Web the unstructured collection of web servers containing text, graphic, and audio files whose contents can be viewed and heard through the use of a browser.

One of the primary applications for Windows NT server is its use as a World Wide Web server platform. To facilitate its use as a web server, Microsoft now bundles its Internet information server with its Windows NT operating system. Although many people choose to user the internet information server as their web server program, other people prefer to user one of the numerous third party web server programs currently being marketed. Similarly, although Microsoft's Internet explorer and Netscape communication Corporation's Netscape navigator programs dominate the browser market, other vendor product can be use to access a windows NT server running Microsoft's internet information server or another world wide web server program. Illustrates the user of the NetManage websurfer browser program to access the O'Reilly & Associates Web Site Professional World Wide Web server program operating on a Windows NT server based computer. Thus, although window NT workstation and Windows NT server represent important development in operating system technology, they also represent platforms from which you can operate numerous third party products.

7.10 Short Summary

The Internet protocol supports an unreliable transmission method, in which datagrams form the basic unit of transmitted information. In comparison, TCP provides an error-detection and –correction capability that results in the retransmission of IP datagrams on as-required basis. The actual routing of data requires three addresses at each source and destination – a port address that identifies the process or application contained in the TCP header, an IP address that identifies the network and host computer where the process or application resides, and a hardware address that defines a station on a LAN.

7.11 Brain Storm

1. Explain about the Internet protocol.
2. What is IP header explain
3. Explain about hypertext transfer protocol.
4. Explain about file transfer protocols.
5. Explain the fields of TCP header fields.

Lecture 8

Network Infrastructure

Objectives

After completing this lesson, you should be able to do the following

- Discuss about Network infrastructure
- Describe the component of network infrastructure
- Describe about Ethernet and Token ring.
- Describe about fiber distributed data interface.
- Discuss about fast Ethernet and asymmetric digital subscriber line.
- Describe about virtual local area network.

Coverage Plan

Lecture 8

8.1 Snap Shot

Infrastructure is somewhat of an overbearing term. According to the dictionary, it means the fundamental facilities serving a country, city, or areas, as transportation and communication systems power plants and roads. Applied to the internet, you can see that definition does fit., yet internet is such a simple term for the incredibly complex and expansive matrix of networks that circle the globe.

In this chapter you will find some basic information on the Internet as a whole; is say some because to go into any detail at all on the current structure of the Internet would require many volumes. Also between the time that these words are written and published, the Internet will have undergone significant change. For instance, in the first six months of 1996, the Internet grew by about 3,500,000 host computers and about 240 domains. In addition, this chapter discusses the major components that make up the networking infrastructure Ethernet, token ring FDDI, ATM, ISDN and more.

8.2 Network Infrastructure

The Internet was intended for a wide range of applications, many of which were only concepts when the original designs were laid out in the late 1960. The Internet is more like a computer than a network, in that it is general purpose in design and capable of handling a variety of tasks simultaneously. While being capable of handling many different tasks, the Internet is also designed to operate over a wide range of technologies. Some are local area networks Ethernet, token ring, and dial up telephone, for example. Over the past 20 years, the infrastructure has progressed significantly, operating over cables, light beams, radio signals and satellite transmission. Each new technology that becomes part of the Internet infrastructure call for the development of new computer code to handle the data security, and protocols.

Today office personnel use PCs attached to a corporate network, which in turn connects to the public switched network. The public stitched network (the AT&Ts, MCIs sprints, regional Bell operating companies, and so on) is rapidly replacing their main trunks (backbones). Copper conductors are being replaced with ultra high speed, high capacity glass fiber. Glass fiber bit rates have soared from about 400 Mbps to almost 3Gbps. Yes that’s 3 billion bits per second. That means a single pair of fibers can carry more than 600,000 voice circuits at the same time. These fibers are part of the Internet backbone infrastructure. The backbone is a broadband system that also incorporates microwave radio and satellite technology. That means that many channels of information can travel simultaneously through a single glass fiber or on a single microwave or satellite link.

Many progressive companies that have an internal network(Intranet) have established a dial in function for employees. They install modem banks(tacks of modems) and a phone line with rotation. That means there is only one phone number to call from outside the company to reach a modem that is currently not in use. By dialing the remote number, the employee can log in to the corporate network, using passwords and secure identification, of course. Once connected to the corporate network, the employee can open a web browser on his computer, select his favorite bookmark, and off he goes. Since he is connected via the corporate network he can also use the same functions he has while sitting at his office desk, if the corporate information system staff has enabled this type of access.

A major part of the infrastructure for the individual is the online and internet service provider. It has only been in the past two to three years that online service providers have linked into the backbone internet infrastructure and provided their customers with internet access. These providers (CompuServe, America Online Prodigy, and so on) essentially represent the intermediate level of infrastructure.

Small office and home computers generally do not have the volume of data traffic to justify the expense of a direct internet connection to the public switched network. Current technology provides for up to 34 Kbps, using an analog modem connected to a normal dial up phone line. So that phone line represents the lower level part of the internet infrastructure. It also provides POTS.(POTS is the acronym of plain old telephone service that we have been using since the invention of the telephone.

8. 3 Component of Network Infrastructure

Now you can achieve the connectivity offered by today's computing and networking environments through any number of techniques. The important thing is that you can achieve connectivity. The windows NT system covered in this set of volumes is one of the most powerful connectivity enabling products to be released by Microsoft into the world of networked computing. But more of that is available else where in these volumes; here we want to stick to the basics of networking data through mediums such as the famous and most widely implemented infrastructure, Ethernet.

8. 4 Ethernet

Ethernet was developed at Xerox corporation in the 1970s. Ethernet is based on packet radio technology. Ethernet like the myriad other terms seem to magically make the internet work, is a standard. The standards are written first as requests for comments, which are circulated through the community of researchers and engineers that read such things for daily enjoyment. The process results in a robust mixture of the best thinking in the world. Once the specification is fairly solid, it is released to the world, where it is usually catapulted into being by commercialization or the military or some other entity with worldly powers.

The first version of Ethernet was released in 1980. Ethernet is a local area network technology. It is rated at two standard speeds 10mbps. Most people think of a cable when they hear the term Ethernet. But Ethernet is not just the cable or any other medium of transmission. Ethernet is a protocol and a cable system. Protocols are definitions of how two devices will connect to each other and what language they will speak so they can communicate. Cable systems are conglomerations of cables, naturally and the matching connections, repeater hubs, swtiching hubs, and other miscellaneous small gadgets.

Ethernet is not vendor specific. You can mix devices from most vendors without problems, but if possible, you should stick to one brand of equipment. That way, you can avoid the finger pointing when something doesn't work between boxes from two different vendors. Today, Ethernet is based on the institute of electrical and electronics engineers standard IEEE 802.3. This standard is called carrier sense multiple access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications. The International Standards Organization (ISO) adopts it, and therefore it is the global networking standard. This standard also contains the specifications for the 100Mbps versions of Ethernet.

Most computing machinery built today uses the popular 10Mbps Ethernet connectors. Even at what seems to be a very high speed of data transmission, Ethernet works over sloppy wiring, even cheap telephone wire. What is not seen, however, is that even though the collision rate of the data, transmission errors, and so on can go up dramatically, the connection will still work. The effect on a large LAN could be disastrous, bogging down the efficient flow of data for all users. You will know there is a problem when the users complain about how slow the network is running.

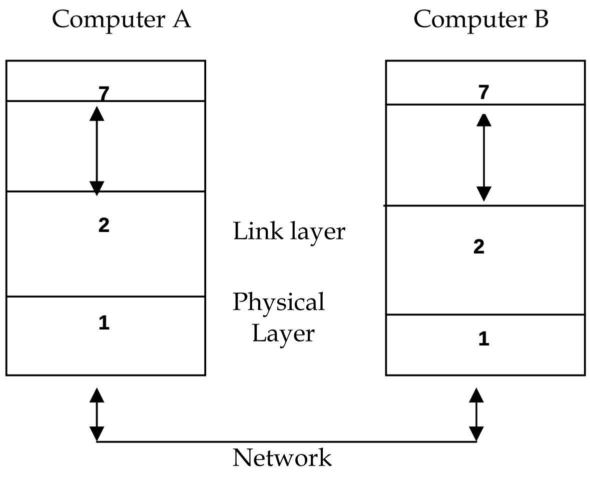

OSI Model

The Ethernet system is actually composed of several layers that relate to the layered design of the Open Systems Interconnect (OSI). Below the figure shows that at the bottom of the stack is the physical layer. The physical layer specifies the connection type, electrical levels, function, and procedure for making the Ethernet connection happen. It contains the specifications for voltage levels, timing, data rates, cable distance, and other "physical" aspects of the link.

The Ethernet protocol, like all packet-switched protocols, has a frame specification. That is a description of what that string of ones and zeros means as it squirts through the Ethernet cable. Below the figure represents the division of the data frame for Ethernet.

Problem	Start of Frame Delimiter	Source address	Length	Header and data	Frame check sequence
7 bytes	1 bytes	6 bytes	6 bytes	46 to 1500 bytes	4 bytes

Ethernet frame format

The Preamble tells the receiving computers that a frame of information is on the way. This is represented by 7 bytes of alternate 1s and 0s. the receiving computer knows when to start treating the frame as information when the start-of-frame (SOF) shows up. The SOF is always 1 byte, with the two least significant bits set to ones (or hex 03). All computers on the LAN treat this frame as incoming data headed destination and source address are each 6 bytes.

These addresses are not the IP address in doted-decimal form; they are the physical addresses of the Ethernet LAN cards plugged into the computers on the LAN. In the destination and source address, the first 3 bytes identify the vendor of the interface hardware, and the vendor specifies the last 3 bytes. Then there is a 2-byte length field that tells the receiving computer the number of bytes of data to follow. The length of the data field is variable – from 46 bytes to 1,500 bytes long. Finally, the last field is 4 bytes of cyclic redundancy check (CRC). The CRC is calculated by the sending computer and again by the receiving computer, and their number must match to allow the use of the data. The frame must always be at least 64 bytes, so if the data field is very short, the sending computer must fill in with null data to pad the frame to its minimum length.

Now, remember that all computers are seeing this data on the LAN. They continuously listen for the alternate ones-and-zeros pattern that indicates there is a frame on its way. They all assume that this data belongs to them, so they start to read in the frame. All the computers read the frame until the completion of the destination address. Then only the computer with the address that matches the destination address continues to read the source, length, data, and frame check sequence fields. There is also a special frame called a Broadcast frame that is fully processed by all computers for address resolution.

Okay, you have a frame of data that has successfully reached your computer from another computer on the LAN. Your computer knows its dotted-decimal IP address,129.0.3.2, its physical address of the LAN interface card (hex something), and the physical address of the source computer (hex something), because it was embedded in the frame source address field. To make some sense out of all these numbers, you need a higher-level protocol to help out.

The ARP request asks for the physical address of the LAN interface card plugged into the destination computer. Again, remember that all the computers on the LAN are listening for the rumblings of the alternate 1s and 0s to wake them up for a possible response. The ARP broadcast request is received and processed by all of the computer and their networking software. Only the computer with your IP address of 129.0.3.2 will answer back with the hex address of its LAN interface card. Now the sending computer has the whole picture. It knows the IP address of itself and the other machine, and it knows the physical LAN interface card for itself and the other computer. Communications can now be established for these two machines.

But what if two computers begin to transmit data at the same time on a LAN segment? This is where the Carrier Sense Multiple Access with Collision Detection (CSMA / CD) comes into play. To begin with, the Ethernet interface card in the computer assists in sensing whether another computer is using the LAN to transmit data. It provides the computer and its network software with three status indications: transmit, receive, and collision. Collision is the detection of two of the ones and zeros strings that begin the Ethernet frame. The detection of this incoming data puts the interface in a "carrier sense" state, telling the computer that it cannot use the LAN just yet. After the packet of data has been received by the computer whose address matches the destination in the packet, the carrier sense signal is dropped and the computer is allowed access to the LAN.

To avoid stepping on the toes of the next packet that might be part of a long string of packets, the listening computer will not try to access the LAN for 9.6 microseconds after the carrier sense signal is dropped. The real genius behind this design is the timing of the bits and delays so that a continuous jamming of the LAN cannot occur by computers listening, waiting, and trying to send simultaneous data on the LAN. If there are 16 attempts to retransmit data and a collision results, the LAN interface reports that a network problem exists. This status is passed up to the network software on the colliding computers.

So what is all this about 10Base5, 10BaseT, broadband, baseband, MAUs and AUIs, and other terms like these in the Ethernet world? First, look at below the figure, and you will see how to read the designation for Ethernet types.

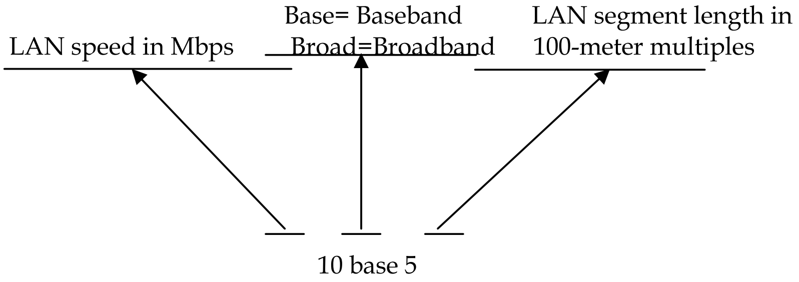

Ethernet IEEE 802.3 Physical Characteristics

The Ethernet physical types are designed to support different local area network speeds and topologies. Below the table shows the relative characteristics of the types. Ethernet is a baseband or broadband technology, depending on which cable type you use. Baseband means one channel of communications can exist at any time on a particular of data can exist on a particular Ethernet cable simultaneously.

Charact-eristic	10Base5	10Base2	10BaseT	10Broad36	10BaseFL
Data rate (Mpbs)	10	10	10	10	10
Signal method	Baseband	Baseband	Baseband	Baseband	Baseband
Maximum segment length	1,640	607 feet	328 feet	1,800 meters	1.2 miles
Media	50-ohm thick coax	50-ohm thin coax	Twisted-pair	Coax	Fiber optic
Topology	Bus	Bus	Star	Bus	Star

Ethernet physical characteristics

Ethernet cable comes in three basic types: thick 50-ohm coaxial cable, which is about 0.625 inch in diameter; thin 50-ohm coaxial cable, which is about 0.25 inch in diameter; and twisted-pair, which is basically telephone wire that has been twisted to canal out electromagnetic interference.

There are two possible topologies with Ethernet: bus and star. The bus topology is the simplest and most common of the two. You should use a bus topology for a large network with many users and longer segments. With repeaters or media converters, you can easily interconnect to other networks with different topologies. Standard Ethernet 10Base5 and Thin Ethernet 10Base2 are based on coaxial cable systems and therefore use the bus configuration. In this type of single cable LAN, all workstations are connected in succession on a single cable.

The star topology is used when you can or have to use twisted-pair cabling (10BaseT). This topology is god for locations where you already have twisted-pair cables in place. Star topology is also used for fiber-optic LANs (10BaseFl). Star topology means that there are hubs located in the midst of users. A hub allows a single bus connection to be connected in redial fashion to the users. The hub has a single, twisted-pair connection running to each user's computer, and there are no taps on the segment.

There are multitudes of Ethernet adapter boxes. You can get almost any kind of adapter you need. For instance, suppose you have two computing centers a few hundred feet apart and you need to run a LAN between them. You also have a concern because one of the centers has a tall communications tower that is one of the best lightning rods in town, and it gets hit often. The best way to avoid extending the lightning problem to the other center is to use glass fiber, because it is totally immune to electromagnetic fields and it cannot provide a ground loop for the lightning hits. At each end of the fiber you can use a 10BaseFL to 10BaseT repeater. The 10BaseT repeater can also be an expander, so it can handle groups of users at each site. Data connectivity is established, and lightning paths are not.

8.5 Token Ring

Token ring is second only to Ethernet in world popularity as a network medium. The token ring concept was born at IBM in the 1970s. The classic IBM local real network is composed of token ring technology. Like the Ethernet IEE standard 802.3, token ring sports is own IEEE standard designation of 802.5. As IBM continues to develop token ring, the IEEE specification 802.5 is modeled to fit. This is a case of the vendor driving the specification instead of the specification driving the vendor.

Like Ethernet, the token ring protocol provides services at the physical and data link layers of the OSI model. Token ring networks can be run at two different data rates, 4 Mbps or 16 Mbps.

Token ring gets its name from passing a token around a ring; it's that simple. Token ring uses the same concept as any packet-switched architecture, in that it has a frame of information with certain specific requirements and definitions. That is what makes token ring a network system. It is not just the way the cables re connected; there is also software involved to make it work. After all, if there were no networking software requirements for a network, it probably would not switch packets very efficiently. In the token ring environment, the owner of the token is king of the hill for the moment. In other words, when a computer has the token on a token ring LAN, it has the rights to transmit data. If a node that receives the token has nothing to do for the moment, it passes the token on to the next machine on the LAN. In this manner, every machine gets a chance to own the right to transmit when its turn come around the LAN loop. Of course, there is a maximum length of time that the token can be owned by any computer, to keep the game fair.

When a computer on the LAN gets the token and transmits a packet, it travels in one direction around the ring, passing all the other computers along the way. As with Earthnet, the packet is usually addressed to a singly station. As it passes the other computers on the LAN, and when it passes the machine that matches the address, the packet is fed to that computer's network software. The packet, however, continues to travel around the ring until it returns to the sending station, which removes it and sends a free token to the next station around the ring.

Token ring networks use what is called a ring topology. However, it is actually implemented in what can best be described as a collapsed ring that looks like a physical star topology.

In token ring LA Ns, each station is connected to a token ring wiring concentrator called a multistation access unit (MAU). In the token ring, the frame format is always one of two types: tokens and data/command frames. Data/command frames vary in length, based on the amount of information being transmitted in the data field. Command frames contain only commands and control information and have no data for the upper-layer protocols in the OSI model.

8.6 Fiber Distributed Data Interface

Fiber Distributed Data Interface is a high-speed data transport technology that is used when large amounts of data are being handled on a continuous basis. FDDI runs at an ample 100Mbps. The specification for FDDI originally required it to run on optical glass fiber cable. Today, due to the improvements in solid state technology, the FDDI specification includes unshielded twisted-pair as an alternate medium. The twisted-pair for FDDI is a higher quality specification than the IBM Type 1 for token ring or Ethernet 10BaseT. Fiber is the preferred medium for FDDI because it can be run over much longer distances than its twisted-pair counterpart.

FDDI is similar to token ring in concept. Like token ring, FDDI passes a token to control the access of machines to the LAN. Also like token ring, FDDI uses the same topology of collapsed ring, physical star. FDDI is a backbone technology at the local LAN level. This is not the same as a backbone at the global level, where gigabit networking is accomplished with light beams. FDDI is useful, for instance, in connecting buildings on a college campus.

FIDI uses single-mode and multimode fibers. Modes are light waves that are applied to one of the glass fiber at a particular angle. In single-mode fiber, there is only one bundle of light waves, traveling at different angles. Multi-mode fiber can have more independent data paths at once, but each path has to operate at a lower speed. The ultimate limit is the glass fiber and how much light it can propagate from one end to the other. Generally, single-mode fiber is used for longer distances, and it uses a laser as its source of light. The multimode fiber is used for shorter distances and multiple paths, and it usually uses a light-emitting diode for its light source.

FDDI is specified with four definitions

- Media Access Control (MAC) defines the access to the medium, including the frame format, token handling, addressing, algorithm for cyclic redundancy check, and the error recovery algorithms.

- Physical Layer Protocol(PHY) defines the data encoding and decoding process, the clock for the link, the framing, and other functions.

- Physical Layer Medium (PMD) defines the transmission medium, which includes the fiber-optic link, power level of the light beams, bit error rates, and the optical transmitting and receiving devices.

- Station Management (SMT) defines the FDDI computer configuration, configuration of the FDDI ring, and its control requirements. It also includes how to insert and remove a fiber from the network, initialization of the sequences, fault handling and isolation, and most important but often left out, the collection of statistics for network management.

Asynchronous Transfer Mode

Asynchronous Transfer Mode is radically different from Ethernet, token ring, and FDDI. ATM is a hardware-implemented, packet-switching protocol-and that is it. There is no cable definition for the ATM Forum, which consisted of more than 500 member companies, established ATM. ATM in 1994 as a result of three years of debate. The ATM specification is a global standard as a result of this forum.

ATM uses a 53-byte, fixed-length packet, called a cell. These cells are considerably smaller than the Ethernet 1500-byte frame, the FDDI 4500-byte frame, and the token ring 9000-byte frame. The key to ATM is its small frame or cell size. The small size allows simpler switching and short packet delays for transmission of the data ATM is an excellent fit for transmitting intense data for real-time voice and video.

ATM does not care what it is carrying for the sending machine. In its 53-byte frame, it will just have a bunch of bits that it gets from one end and drops off at the other. Since it has a small frame format, it is quick and limber. Following figure the format of the ATM frame.

5 bytes			48 bytes	
VCI Label 3 bytes	Control 1 byte	Header Checksum 1 byte	Optional Adoption Layer 4 bytes	Payload 44 or 48 bytes

The ATM frame format

The cell consists of 48 bytes of the user's data, which can be voice, telephone, video, computer data, or anything else that can be represented with binary. It has a 5-byte header attached, which contains the addressing. Note that the addressing is more than just destination and source. Within the 48-byte data field, there might be a 4-byte field for adaptation information. Adaptation bytes are used to fragment the original packets and reassemble them when they arrive at the destination. When this occurs, a control bit is sent in the header information to specify that adaptation is in use and the receiving end will have to handle putting it all back together again. There are also bits in the control field to indicate whether the cell if for flow control or data. Flow control cells are used to prioritize the cell and to contain information about bandwidth congestion.

ATM uses three protocol layers:

- Adaptation layer – Ensures that the service characteristics are satisfied, and is responsible for dividing the data into 48-byte payloads.

- ATM layer – Adds the 5 bytes of header information to the payload and then checks to see that all cells are sent to the proper virtual circuit on the far end.

- Physical layer – Defines the network interface.

ATM must establish a virtual circuit between the sender and the receiver for ATM communication to be successful. The sending machine negotiates for a virtual circuit with the network by specifying the type, synchronous or asynchronous, speed, and other criteria such as priority. ATM then established a profile for this circuit request in order to set up the necessary path as specified by the sending machine. Once the circuit has been realized, cells with the same virtual circuit identifier label arrive at the destination in the same order in which they were sent form the source.

8.7 Fast Ethernet

As the user needs continue to climb to higher and higher data rates at the local LAN level, the networking world looks for new, faster, and less expensive methods to transport the packets. Fast Ethernet is one answer to this dilemma. Fast Ethernet is defined as 100BaseT. It is a technology that can be blended into the existing 10BaseT LAN whenever and additional capacity is needed. Fast Ethernet provides the additional bandwidth without major restructuring or replacement of the existing LAN topology. Today's trend is to locate the high-powered servers in a data center where they can be secured and administered, and to use 100Mbps networking as a local connection between the data center and the user concentration.

Fast Ethernet uses the same Carrier Sense multiple Access/Collision Detection (SCMA/CD) protocol as its lower-speed cousin 10Base. Because of this compatibility, 100BaseT Ethernet is considered to be standard Ethernet, except that it is 10 times faster.

The cabling infrastructure for Fasts Ethernet 1000BastT supports three types:

- 100BaseTX is a two-pair system for data grade universal twisted-pair cabling.
- 100BaseT4 is a four-pair system for both voice and data grade universal twisted-pair cabling.
- 100BaseFX is a standard two-strand optical glass fiber cable system.

The range of these cable alternatives allows the migration of a 10Mbps LAN to 100Mbps, conveniently and economically. All these cable systems can be intermixed through a hub, with excellent results.

Frame Relay: The Cloud Topology

Frame relay is a protocol that was designed to be used with Integrated Services Digital Network (ISDN). As it turns out, Frame Relay became a network infrastructure all its own.

Frame Relay is a packet-switching data communications structure that is applied at the interface between user devices and network equipment – for instance, among routers, or computers, and the network switching nodes. One of the major public network carriers or can provide the network that provides the Frame Relay interface by a privately owned network. The interface for Frame Relay is the same as X.25, but it is different from X.25 in function. Frame Relay is a slim protocol that provides high efficiency and excellent performance.

Like ATM, the frame relay specification allows permanent virtual circuits and switched virtual circuits. PVCs are the most common and the lowest cost option.

The local management interface extensions are

- virtual circuit status messages [provide communication and synchronization between the network and the user equipment connected to the network. It reports the existence of new private virtual connections. It also provides information about private virtual connection integrity. If a private virtual connection is removed from the far end of a frame relay circuit, the virtual circuit status messages report that it is no longer present so that data is not dumped into the bit bucket at the far end. In frame relay, it is assumed that the backbone carrying the frame relay cells will be adequate and the receiving end will take the data. There is not status reported back to the transmitting end to say the data has been received.

- Multicasting – Allows a transmitting machine to send a single frame and have it delivered by the network to multiple end nodes – like standing in front of a crowd of people and shouting "Hello to all of you". The multicasting local management interface provides the routing directions to the physical network infrastructure and the address resolution. Otherwise, the sending machine would have to send the data to each receiving site individually – one right after the other. Multicasting is therefore good for reducing the overall network traffic.

- Global addressing – Gives connection identifiers global significance. It allows the Frame Relay network to look like an extension of the local area network as far as addressing machines is concerned. It is global addressing that makes the address resolution protocols (ARPs) perform over Frame Relay exactly as they do on the local area network.

- ♣ Simple flow control – Provides XON / XOFF flow control. A long-standing interface between data terminal equipment (DTE) and data communications equipment is the XON/XOFF commands in the RS-232 standard. If simple flow control is not implemented in the vendor's Frame Relay interface, you must be certain that your sending machines do not send out data faster or more dense, than the receiving machines on the far end can take in. After all, you wouldn't want to go to the bank teller machine and transfer your money from one bank to another only to have it leave one bank and drop into the bit bucket at the other bank. It could happen if handshaking is not implemented. *Handshaking* is the generic term for what two devices do with each other at the data interface point to control the flow of data.

Tomorrow's Component of Network Infrastructure

The Telecommunications Act of 1996 that was signed into law is beginning to push the envelope of technology for data communications. The enactment of the bill has opened the door for telephone companies to compete in Internet, video on demand, and other multimedia markets. This has already given birth to new technologies such as Asymmetric Digital Subscriber Line, virtual LANs, personal communications systems, and Gigabit Ethernet. These are discussed in the following paragraphs.

8.8 Asymmetric Digital Subscriber Line

Asymmetric Digital Subscriber Line (ADSL) service is the latest and greatest technology for getting high-speed data to the individual subscriber. ADSL has been tested in many locations in North America with excellent results. The technology is so new that commercial availability announcements are just being made at the time of this writing. There have been recent news releases that commercial ADSL interfaces are now available as standalone interfaces and personal computer interface cards.

ADSL provides high-speed data transfer, and it is modem technology that uses the existing phone line to a subscriber's home for connection to the subscriber's phone company. It is very much removed from the standard analog modem technology that brings the user data at rates up to 28.8 or even 34Mbps today. The eye-opener is that ADSL increases the analog modem rates by factors of 200 or more without the need to replace the existing telephone wires serving the subscriber's premises.

How does an 8Mbps download of a Web page sound, using the same wires that used to service your incredibly fast mid-1996 analog modem? That is what ADSL provides to transport data from the source to the subscriber's computer. Going in the other direction, from the user's computer to the source – for instance, requesting the download of a Web page from the Internet – the speed is still a whopping 640Kbps. There is no doubt that ADSL will soon spread through the marketplace and as it does, the cost of installation and use will plummet.

The use of ADSL is gaining popularity for two reasons. First, there is the incredible speed, and second, it uses the existing telephone industry architecture. It is estimated that the telephone companies cannot replace their embedded twisted copper lines faster than about 4 percent per year over the growth of new installations. The number of subscriber lines in the world is estimated to be about 640 million. Therefore it would require a quarter century to establish a new infrastructure capable of the high speed that ADSL offers. The driving force for ADSL in North America will not be the Internet. While many of us who dedicate our lives to the computer screen might think otherwise, the ability to support movie video on demand will most likely drive the cost of ADSL downward far quicker than accessing Web pages.

The asymmetric part of the technology means that the signals are different depending on the direction of transport. The downstream rate is more than 10 times the speed of the upstream rate. Some of the logic behind this is cost and some is technical. When new technology is implemented, you can be certain that economics will be the driving force. Totally unlike the original ARPANET, which had research-driven funding, the ADSL development is based on the gleam of wealth from video services-oh, and by the way, you can also use it for the Internet and connecting to your business. The success of ADSL hinges on the multimedia communications in our new age of information and connectivity.

Another key for ADSL is security and data privacy. ADSL is secure from the standpoint that it operates over a single telephone circuit that belongs to the user. This provides a unique, digital, point-to-point connection. ADSL is a highly complex signal structure that would be extremely difficult to wiretap.

Asymmetric Digital Subscriber Line requires an ADSL modem on each end of the circuit. What results is a combination of three services. Much like Integrated Services Digital Network, ADSL provides a high-speed downstream channel (about 8Mbps), a medium-speed duplex channel(about 640Kbps), and a plain old telephone service (POTS). In the design, the POTS channel is filtered and separated from the high-speed digital channels this allows your telephone to continue working if the digital service fails.

8.9 Virtual Local Area Network

Almost every major LAN vendor has recently introduced the virtual LAN (VLAN) concept. Virtual LANs will probably need to wait for the established world of routers and bridges to live out their productive years before the LAN base moves to Ethernet and token ring switched topology. Virtual LANs require a switching topology to exist. The virtual LAN seeks to minimize the amount of administrative time necessary to maintain the network. It does this through the automatic configuration that makes it virtual. The great thing about VLAN is that the network manager is supposedly released from the drudgery of LAN administration so that more attention can be applied to establishment of policy and structure.

Today's LAN switches are becoming less costly per switching port, allowing single user-per-port configurations. This makes the Virtual LAN an ideal fit. VLAN is not being deployed with great enthusiasm yet due to no universal across-the-board standardization; there is still a cost associated with administering the VLAN and it presents difficulties in enabling full, high-performance access to centralized servers.

A VLAN is essentially a broadcast domain in concept – in other words, a LAN in which the machines on the outer fringes can communicate with each other as if they were located on the same segment of Ethernet or token ring.

To be effective, VLAN has to be self-configuring based on a set of policies set by the system administrator and on the applications that the user runs. VLANs are proposed to be the network administrator's cure of most ills, but not all. Administrators spend about three-quarters of their time maintaining the infrastructure of the network and handling the movement of personnel and their machines into and out of domains, permissions, and file sets. Just the moving and changes alone is a particularly time-consuming task that does nothing to enhance the network or the network administrator's talents.

One of the major reasons for implementing a VLAN is to eliminate the use of routers to contain broadcast. Routing will always be necessary in the network but Ethernet or token ring switches instead of routers more effectively handle broadcast. Switches are also more cost-effective than routers on a per port basis. Routers are more complex and configuration is often a time-consuming and frustrating task. The virtual LAN is supposed to offer more network administrator – friendly methods of managing the domain and user population.

The VLAN offers a greater security. Network managers can define the applications, access, and ports on the VLAN switches in lieu of installing router-based firewalls that are more complex and more expensive. In addition to IP address and application data type, the administrator can also define the port. In single-user, per-port configurations, you have control as to who can access what. It is essentially the same as opening the network between the user and the machine so that the user is not allowed to have a connection to that machine.

There are some concerns about the VLAN automation of network administration and management. There is always a cost/benefit line where more automation results in more complexity and more complexity offsets the cost savings from automation. The more major concerns are that the virtual connectivity layered on top of physical connectivity makes troubleshooting more difficult, maintaining a VLAN across a maze of switches is confusing and difficult, and network traffic analysis is complicated by the presence of the VLAN software for the management task. Management software should be flexible and graphical. Preferably, even graphics of the LAN topology, which are generated from a database of devices and port connections.

8.10 Short Summary

Ethernet uses a high-level protocol called Address Resolution Protocol (ARP), ARP uses the same technique as standing in a large crowd with a megaphone and saying “Will Joe Blow please raise his hand?” It is that simple. ARP sends out a broadcast request onto the LAN, asking for the computer that owns the dotted-decimal IP address being sought to answer back.

Frame relay uses local management interface extensions. The LMI extensions are in addition to Frame Relay’s basic protocol for transferring data. The extensions are intended to make the implementation of complex and even global inter network easier to manage and control. In the specification, some of the LMI extensions are common to all requirements and the network expects them to be implemented by the user equipment. Other LMI extensions are defined as optional and as such, cannot be depended on to be present. Each vendor of frame relay will implement a subset of the optional extension they feel is important to their line of networking equipment.

8.11 Brain Storm

1. What is the Internet infrastructure?

2. Differentiate the today and tomorrow's component of network infrastructure.

3. Explain the concepts of FDDI.

Lecture 9

Internet Service Provider

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Internet Service Provider.
- Describe the needs of an ISP.
- Describe Basic connections.
- Describe about Bulletin board systems.
- Describe Not-so-basic connection and it needs.

Coverage Plan

Lecture 9

9.1 Snap Shot

In this session we discuss about Internet service provider and its needs. In ISP we have four classes global, national, regional and local. With the advent of internet protocols, modems, and web browsers, virtually anyone can gain access to the internet. ISP is important to note that not all ISPs are alike; you need to find out their capabilities.

9.2 ISP

The ISP is your only method of accessing the Internet. There are four classes of ISP; global, national, regional, and local. All regional Internet service providers must route traffic through national Internet service providers. Similarly, all local Internet service providers or they might use a national ISP. National (United States) ISPs can and re global with overseas services. So there are tiers of service and infrastructure are global with overseas services. So there are tiers of service and infrastructure when it comes to moving packets of data around the globe. "IP Addressing and the DNS", you can see that there must be an orderly assignment of IP addresses to maintain the integrity of the Internet. This is called registering your domain. Depending on what level of service you obtain, you may or may not have to register your domain. You might not even have a permanent domain to register.

By having the hierarchy of national, regional, and local ISPs, we are assured that from the top down, the IP addressing is correctly laid out, much like a directory tree on your computer. How do you find an ISP? There are a couple thousand to choose from. The key is to narrow the field quickly to a few select prospects and then start to be detailed in your study. If you are an individual who just wants to be connected, you will still find what you need in the following parts of this chapter.

9.3 Need of an ISP

Four classes of Internet service providers: global, national, regional, and local. Regional providers might offer connections in a state or a few states or perhaps half of the United States. Local providers might cover a large metropolitan area or a few geographically close cities. Another new entry into the market is the rural provider that is serving the larger, sparsely populated areas. Regardless of where or how much territory these providers service, they have to get their service from the big backbone carriers at the national or global level.

The Internet that started out a quarter century ago and grew up was dependent on the little pieces in the beginning. Now it is the little pieces that must depend on the big boys to provide any real value as a connection of global significance.

If you are facing the connection of a 5,000- or 10,000- employee corporation to the Internet, you could be in the national or regional market. If you are an individual who has a burning desire for a personal Web page just for the heck of it, you are probably in the local market. However, even a corporate-level interest can be quite satisfied with a local-level ISP. It really depends on what you want to do with your tap into the networked world.

Regardless of the size of the ISP, there is only a fixed set of services that can be purchased through an ISP. That is because Internet functionality is universal. No matter what connection speed or type, or how many miles of network or what the corporate organization looks like, you have a fixed bag of choices. The choices we are talking about here are the services, not the cables or connections or speed of the data. The bottom line is that if you want to connect to the national Internet backbone, you need to start seeking an Internet service provider because you cannot connect without one.

You will find an excellent source of Internet service providers on the World Wide Web. Point your browser to http://thelist.iworld.com/ and you will gain instant access to thousands of providers around the globe.

9.4 The Basic Connection

For the individual seeking a personal connection or the corporate Internet project manager wanting to try some sites to get the feel of the what this is all about, the following are some of the most basic ways to get started.

UNIX Shell Accounts

The absolute basic of basics is the UNIX shell account. This service is really the core deliverable for most local Internet service providers. The reason for this is the popularity of the UNIX operating system and the Internet-related tools that are available. After all, UNIX was handling the Internet long before Microsoft developed the first Windows version 3.0. With a UNIX shell account, you only need to have a standard personal computer like a 486 with a modem and a terminal program that will give you something like a DEC V 100 emulation. The DEC V 100 is vintage, but it is one of the most widely emulated terminal standards known. Using the terminal program, you simply use the directions provided by your ISP to connect to a UNIX shell account that the ISP establishes for your use. Once connected, you need a handy UNIX command-line reference and patience. With this simple setup you will be able to download files, send and read e-mail (usually Pine mail), and access newsgroups. Do not expect any densely colored screens or multimedia. Once you have used something like Netscape or Microsoft Internet Explorer, going back to a UNIX shell account is like canceling your hotel reservation and getting out the tent. It works and it works very well, and usually this type of account is less expensive although more difficult to find these days.

9.5 Bulletin Board Systems

Next in the basics line are the bulletin board systems (BBSs). Using the same 486 with modem and terminal program discussed earlier, you can access a BBS. Almost every city of any size has one or many. The BBS is operated sometimes by a kind soul who just enjoys having to look after the thing. Usually if the BBS is of any value, there will be some minimal charge. Many of the BBSs offer a connection to the Internet, and this part of the service might cost extra. BBSs come and go, so if you are looking for a serious Internet connection, keep reading.

Commercial Online Services

On the next tier of the basics is the commercial online service. The most prominent are the ones you get in the mail at least once each week. Prodigy, CompuServe, America Online, and Delphi are the most widely offered, even distributing their products with magazines at the newsstand. If you are seeking a fast connection to the Internet – and by fast we are talking about opening your mailbox, making a phone call to log on, and starting to surf – all you need is your credit card. For the novice, or the family entertainment and let's-see-what-this-Internet-is-all-about connection, these services really cannot be beat. With these services there is often a monthly minimum charge for a certain number of connected hours. After the minimum, the charge is by the hour.

If you plan to really investigate the depths of the Internet, you will want to secure an unlimited account. Unlimited means just that: You can log on and begin to plunder and loot the Internet 24 hours a day, all month long, and the charge is still the same.

If you are seeking a corporate ISP for this thing called the Internet, and you have not had the opportunity to experience the bliss that Internet connectivity can provide, you might consider using one of the commercial offerings to gain some familiarity.

The Plain Vanilla ISP

Enter the Internet service provider for the not-so-faint-of-heart. These are services like the AT&T WorldNet, BellSouth Net, Pacific Bell Internet, MCI One, and the slew of new entries making connectivity possible. As with the commercial online services, you can obtain a user account from most of these that will give you unlimited access to the Internet.

The primary difference between these ISPs and the commercial online services is content. Content is what you have access to when you connect. For example, many new customers leave a commercial service like America Online and become frustrated to find out that their chat rooms are nowhere to be found on WorldNet. Chat rooms are where simultaneous text discussions take place and all who participate see the same discussion. This perceived shortcoming is because AT&T offers three basic services: newsgroups, e-mail, and Internet access – and that's all, folks. Recently WorldNet and America Online have struck an agreement allowing access to AOL through WorldNet.

The dial-in ISP offerings are springing up everywhere. Some, such as Pacific Bell, have already begun ISDN service for the individual.

9.6 The Not-So-Basic Connection

If you need a mega-link to the Internet for your corporate-sized volume of traffic, you will need to get into the serious class of Internet service providers. The commercial online services also offer more serious and robust connections for the small business or corporate user. When you are considering the Internet as a means of revenue, more though and investigation are in order (especially if you report to company management that expects real results out of that handsome salary you are being paid). If you are faced with making this serious selection, you will want to consider securing the services of an expert in the field. Another key to success is to establish an evaluation team of your sharpest individuals but look for a mix of skills and strengths that complement each other. A team full of C+ + analysts will not see the broad picture that a team of analysts, engineers, and, above all, users will. Nothing makes a project more successful than involving the users up front to obtain their buy-in.

What Will You Use It For?

Develop a short document defining what is in the scope of the project and what is not in the scope. This will not prevent you from experiencing what is known as scope creep, but it may help you at least identify that it is beginning to happen. Identify your purpose.

Having an Internet connection also means having an Internet connection administrator. This person is charged with the responsibility of overseeing the health of your connection to the outside world. This is one of the many overlooked requirements of a successful Internet installation.

It is critical that you and your management are up front about the manpower that this Internet thing is going to require. The worst thing you can do is to assume that someone is going to magically step in and click here and there to straighten out all your woes. They just will not appear. Oh, of course there will be the self-appointed Internet expert, who has had one of these connections at home and can whip some freeware or shareware out of his desk drawer and leap to the rescue.

The self-proclaimed expert might be able to get something running, but it might be at the risk of losing valuable business data or, even worse, setting up your corporation for a software license violation. This will certainly get you in the limelight from the management viewpoint.

A couple administrators are needed because one person cannot work all the time. You might have only one full-time administrator, but you should appoint someone to be the backup and expect him to be knowledgeable of the basics. He should at least know how to restore user accounts, add user accounts, back up the system, press the reboot switch, and perform other higher-level functions. He can leave the really gutsy stuff to the main administrator. At the very least, you should have available an individual who can

- Add or delete e-mail accounts
- Add or delete e-mail aliases (forwarding addresses)
- Upload or download files with anonymous FTP
- Maintain HTML files if you have a Web server
- Install and maintain common gateway interface (CGI) programs
- Maintain server log files and security.

How Big Is It?

After you have determined what processes will benefit from the Internet connection, you should attempt to size the volume of traffic that the Internet connection will sustain.
The amount of bandwidth you need will obviously vary, depending on your users' expected use of the Web. Bandwidth is the amount of continuous data traffic that a given copper or glass fiber circuit can reliably carry. It's like water through a hose. A larger-diameter hose allows more water through than a smaller hose, given the same push at the faucet end. Internet circuits are no different; in fact all communications circuits are specified to somelevel of bandwidth. Here we see rating of Kbps (thousand bits per second), Mbps (million bits per second), and Gbps (billion bits per second), which is for the really serious user. That's really moving it through the pipe. Of course with more speed, there is more cost.

Calculating the amount of bandwidth required for a Web server is not at all an exact science. However, one can make estimates based on assumptions of the number of simultaneous requests, average size of data transfer, and amount of time users are willing to wait.

What Services Are Available?

Services from an Internet service provider can range from "here is your pair of wires" to "sit down here and click this mouse button". That is to say that there are as many and as few services as you ask for, if your provider offers them. Some services are required and others make life easier for the folks who are charged with maintaining this connection and seeing that they proactively respond to growth. That means watching the Internet connection usage and having the upgrade ready to be installed just before the CEO of your company angrily advises you that the Internet stock report is too slow. Such phone calls are deemed to be career limiting and should be avoided at all costs. This is entirely possible with the reporting tools and services that most corporate-quality ISPs offer. Following is a list of services that you can find available from most providers:

- Domain name services
- Web site development and specialized browsers
- Hardware and software
- Account management
- Flat-rate pricing.

- ☏ Training, tutorials, and support
- ☏ User-friendly documentation
- ☏ Network operations center
- ☏ Network usage audit
- ☏ Metropolitan area network

9.7 Need of Not-So-Basic Connection

You need to begin by identifying the requirements. This is absolutely the number one step, even though it can be performed in parallel with requesting high-level service offerings from some providers that seem to be obvious choices up front. Identifying the requirements is often a difficult task, especially when your used population has no clue what it would do with an Internet connection.

If your users are green on the subject, you will have to begin a little familiarization training. Since you are seeking an Internet connection, you might not have a connection to even a commercial online service. As pointed out earlier, you should not have to look far to find one of those freebie sign-up disks or CDs. So get it, plug in a fairly robust personal computer-Pentium class if you have one-and get started. All you have to do is follow to do is follow the instructions, and – shazam! – You will be online in no time at all.

Your team will need to formulate a list of the business processes that need to access the Internet. By developing which business process can benefit from Internet access and which need to have it, you can remove the desires of individuals who want to be on the Internet. There is difference between wanting and needing, and sometimes the intangible benefits get really gray (or soft, depending on your terminology) when personal agendas enter subtly into the picture.

9.8 Short Summary

ISPs are local, regional, or national Internet providers. Local ISPs are connected to the Internet through larger ISPs or through regional and national ISPs have multiple locations. Most Internet users connect to the Internet using a service provider. Each user has a TCP/IP address.

9.9 Brain Storm

1. What is ISP?
2. Give some examples for ISP?
3. Differentiate the needs of basic connection and not-so-basic connection?
4. Explain about bulletin board system.

Lecture 10

Domain Name Services

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the services of Internet.
- Describe the Domain name services.
- Describe web site development and specialized browsers.
- Describe about account management.
- Describe the network operations center.
- Describe about connection type and security.

Coverage Plan

Lecture 10

10.1 Snap Shot

In this session we discuss about the services of Internet and domain name services. Auditors love to pull the access logs to check on the security levels and see if the system administrator is doing his job. There are volumes and volumes of logs maintained by the ISP. What you buy with account management is the logical ordering and filtering of these logs so there is information and not just data. Firewalls are a critical component of any network security architecture. It takes good firewalls to provide protection without becoming a nuisance to the user. Firewalls have improved significantly in a short period of time because of the urgency created by network intrusion.

10.2 Services of Internet

Services from an Internet service provider can range from “here is your pair of wires” to “sit down here and click this mouse button”. That is to say that there are as many and as few services as you ask for, if your provider offers them. Some services are required and others make life easier for the folks who are charged with maintaining this connection usage and having the upgrade ready to be installed just before the CEO of your company angrily advises you that the Internet stock report is too slow. Such phones calls are deemed to be career limiting and should be avoided at all costs. This is entirely possible with the reporting tools and services that most corporate – quality ISPs offer. Following is a list of services that you can find available from most providers:

- Domain name services
- Web site development and specialized browsers
- Hardware and software
- Account management
- Flat-rate pricing
- Training, tutorials, and support
- User-friendly documentation
- Network operations center
- Network usage audit
- Metropolitan area network

10.3 Domain Name Services

The domain database is maintained by InterNIC. Currently, InterNIC is operated by a firm called Network Solutions, Inc. The process requires you to fill out a form and apply for your very own domain name.

You cannot be connected to the Internet without domain registration. This prevents duplicate domain addresses from existing. Once you are connected and have your own dotted-decimal notation and your chosen English site name, you will be ready to go.

If you are not into this level of frustration with the InterNIC, you should select an ISP that will register your domain name for you. This may cost a little more, but it is probably worth it to keep your blood pressure stable. With the onslaught of domain registration requests, the InterNIC is waging an upstream battle to stay afloat, and delays or errors should be anticipated.

The larger ISP will offer a domain registration service that provides more than just filling out the form and sending it to InterNIC. These services usually include working with you to establish a scalable host-naming scheme that is consistent with Internet guidelines and your needs. The InterNIC also requires a secondary DNS address on its form. In many cases the ISP provides its address as the secondary machine. Otherwise, you will have to provide two machines with different network addresses.

If you use one of these services, you only have to e-mail the ISP contact person when you add another machine to your network that affects the naming scheme. The ISP takes care of the details, usually in batches, perhaps on a weekly basis. There is another factor in favor of this approach. If your site is made up of non-UNIX machines, you may have to install one UNIX machine for DNS if you do not allow the ISP to handle DNS. This is because large ISPs are more likely to be UNIX based.

10.4 Web Site Development And Specialized Browsers

You can develop your own HTML Web pages and establish your Web server on your internal network. This demands that you address the issue of skills and manpower up front, just like having a system administrator. Of course the same individual can do both functions, and unless there is some juicy Web work to do, the system administrator will quickly become bored with simple account administration. You should also notice that there is a machine involved. The Web server requires a machine running Windows NT 4.0 with Internet

Information Server (IIS). You will also get a Gopher server and an FTP server. When you get to this point, you have crossed over the threshold in to the realm of intranet. You will also have set yourself up for a small support organization to handle the needs of the company and the server. It really depends on how far you want to go with implementation.

Your other alternative is to let the ISP handle your Web site. This usually includes a machine or services on a machine owned by the ISP and located at the ISP's headquarters. You can also purchase the level of service you desire (as always), but three that are usually included are

- ♣ Web communications service – This provides the infrastructure to set up your Web presence. It includes a Web server and the network connections.

- ♣ Web publishing service – This provides the trained arts and crafts types for content design as well as the HTML experts to make it all hum. You may have a choice of which browser-oriented publishing system you want to use.

- ♣ Web merchant service – This provides pricing and inventory tools for online commerce applications and it supports the processing and billing services like charges to credit cards. You will also need this service to connect to your internal products databases and perhaps your shipping and billing systems.

In addition, you can get service-level agreements when the ISP owns the machine. For instance, these services will probably be backed by a guarantee that they will be available to the spending public more than 99.5 percent of the time. That means you will be able to sell via the Internet about 51 of 52 weeks per year. The lost week is the 0.5 percent that the system is allowed to be out of service.

Another key feature that is offered is the specialized browser function. This means that you can have Netscape or Microsoft Internet Explorer customized with your logo and corporate colors – even a glossy photo of your CEO with a big grin.

10.5 Account Management

If you are small company, you probably will not need all the account management services that major ISPs offer. It will depend on what you really want in the way of information about what is going on with your Internet connection. That is what account management means. You can get reports of who, what, when, and how in as much detail as your paranoia requires. Just remember that the more reporting you get, the more you will have to look over, and that takes time unless you also use threshold accounting.

Whether you know it or not, when you access an FTP site, or call up a Web page, or Telnet to somewhere, or just log in to your service, your actions are identified and are usually recorded and kept for some period of time. Auditors love to pull the access logs to check on the security levels and see if the system administrator is doing his job. There are volumes and volumes of logs maintained by the ISP. What you buy with account management is the logical ordering and filtering of these logs so there is information and not just data. Some of the reports you can get are

♠ FTP usage summary – Total FTP usage, top FTP users by files transferred, top FTP users by bytes transferred, total puts versus gets (put and get re HTML commands).

♠ HTTP by hour summary – Produces a list of the number of HTTP accesses per hour.

♠ Mail report – Produces a summary of mail usage and you can even get it to just pick on the boss as he or she spends the day e-mailing friends and relatives.

♠ Smut report – Yes, even that. This report can list users who have visited or tried to visit any site with a keyword listed in your smut file. You might have to develop the smut file of obscene words or words that imply a nasty triple X-rated site or you may be able to use one of the ISP's standard files full of smut words.

♠ Telnet report – Telnet usage by user, connect time, or number of bytes transferred and reports of the top users. For your internal bean counters, you can usually get this report with accounting for Telnet charge-back purposes.

♠ Top surfers report – Produces a list of the access by selected user with time of day to pinpoint the abuse of the Internet connection by a particular individual.

Even if you have your own Web server, you can produce these same types of reports by logging all traffic with the server's monitoring tools. If your selected server does not produce specific reports, you can write a quick program with a product like Microsoft Access to query text log files and format the results into something that is unique. Be careful, though; if the world gets out that you are snooping (that is what the users will call it), you will certainly lose face in the social scene.

Flat-Rate Pricing

Many ISPs, such as the commercial online services, charge by the amount of connect time and / or the number of packets transacted over some minimum. You will need to evaluate your usage patterns and determine which will be the best route for your application. Even if you start out with a leased and dedicated circuit, you should watch your usage to see if switching the billing options will be a good idea.

Training, Tutorials, and Support

Sometimes it is not the connection speed or other services that matter as much (at least in the beginning) as training. The connection will not be much benefit if you do not know how to use it. The ISP may offer some amount of training with the connection cost. You should determine what this will be before signing the lease. One option you can seek if you are not in a position to train away is to get a training pool, which is a pool of dollars designated for training that you can use whenever you get people or time. This usually has an expiration time limit and you will not want it to lapse without getting the benefit of the classwork.

10.6 Network Operations Center

A network operations center is an absolute must for the serious connection. Networks are complex animals that can go astray with just a simple typo like a numeral in a DNS table. When that happens, you cannot connect no matte how hard you try. After all, garbage-in, garbage-out. The network operations center must to operate 24 hours per day, 7 days per week. That is called 24* 7 coverage, and anything less will surely leave you stuck. Murphy's Law dictates that the Net connection shall go down the very hour that the network operations center closes. You may be offered 24*5 or 12*7. These mean 24 hours per day Monday through Friday and 12 hours per day 7 days per week, respectively. If you are serious about your Internet connection, you will not take less than full coverage. However, if your connection is primarily for your outgoing use Monday through Friday, it makes sense to pay less and not

cover the weekend. Just be sure that come Monday morning, the network center is back up, and test your circuit before your employees arrive at work to get started.
The network operations center (NOC) is the place you call when your users call you. It will not be long after the connection is lost your phone (or your system administrator's phone) starts ringing. If you are fortunate, you can use the intercom to let everyone know you are working on it. When you or someone calls the NOC to report trouble, you should get a trouble number so you can check on their progress. After all, you are not their only customer and the trouble may be more widespread than just your connection. When you report the problem, it may just be intermittent. Then later, it may happen all the time, leaving you stranded. You will need to call the NOC and tell them to escalate the trouble ticket for higher-priority repair.

When you lease a network from a major ISP, you will probably have the option to include routers and termination equipment at the same time. You can often make an economic decision to lease or buy the equipment and still have it included in the realm of control of the NOC. If you buy the equipment, you will have to coordinate the vendor, type, and model with the ISP if you want them to oversee its health. The reason for this is the Router Information Protocol (RIP) that is used for remote testing and even remote programming.

Network Usage Audit

You will want to obtain usage audits from your ISP. These tell you how to plan for future growth of your connection. Audit reports are usually tabular, but if you can get graphical outputs, these will be a big help in visualizing the load on your network. You can probably get a report that includes incoming and outgoing packet counts by the hour. At least you should be able to get minimum, average, and peak usage defined as a percentage of the circuit capacity. A word on circuit capacity: you can get a 45Mbps T3 circuit, but only purchase a portion of its total capability and save the rest for growth. You will want your usage report to represent the amount of your purchased capacity, not its ultimate capacity.

10.7 Connection Types, Speeds, and Costs

Today's ISP that is business oriented offers a wide range of service and connection types. Any connection can be installed with a contractual limit on the data transfer capability (bandwidth). Many sell asynchronous services through a dial-up connection. Dial-up is good for mobile users and remote office locations. ISDN is probably the current best buy for the budget-minded startup interface.

Dedicated circuits start at 56Kbps and are offered at higher speeds in increments of 56Kbps or 64Kbps to T1 service that is 1.54Mbps. The next increment is a T3 circuit. The leap from T1 to T3 is significant in speed and cost. T3 is capable of 45Mbps, almost 30 times the capacity of a T1 circuit. You will need to be a very large user to require the full capacity of T3.

Average monthly costs are variable and while the trend is downward, temporary fluctuations in the market may increase your cost over the last customer with the same capabilities. consider leasing circuits as being similar to the fluctuating prices of computer memory, and you will understand what we mean.

Within each of the connection types you can find various speeds with various costs. For instance, a tiered T3 is a T3-capable circuit with a contractual cap of different data rates and costs. One advantage if you need this much to start out with is that it can easily be upgraded. It will probably only require a contract amendment to bump it up in 3Mbps increments to wherever you want to take it, upto 45Mbps.

When you contract for your connection, be clear about who provides the equipment to connect to the circuit on your end. Devices such as the router and CSU/DSU are required and may or may not be included in the circuit cost. It is best to draw the circuit on paper and let the ISP put the dotted line where its responsibility stops and yours begins.

10.8 Security

Internal business network is safe and secure until you connect it to the outside world. Once connected, your network is vulnerable to all levels access, from pranks to serious industrial espionage. The Cold War may be a thing of the past, but clandestine plunder through your network is still and always will be a threat. Computer security is like securing your home against a burglar. It comes down to your locks being tougher than his lock-breaking tools. With computer security, you can take care of today’s threat, but the forces of evil are continuously seeking access with different methods and the legitimate community must continuously seek new ways to defeat the access.

There are two general methods you can use to secure your network against these dreaded forces: firewall and encryption. A firewall involves using a computer between the Internet and your business's internal network. Encryption involves using algorithms to scramble the data inside the data packets so that they are unintelligible to outsiders.

Firewalls are a critical component of any network security architecture. It takes good firewalls to provide protection without becoming a nuisance to the user. Firewalls have improved significantly in a short period of time because of the urgency created by network intrusion. Today's firewalls provide audit trails of illegal access and administrative control over secure Internet and intranet environments. Firewalls are not in any way simple devices. If you intend to operate your own servers and just lease the connection from the ISP, you will need to learn great deal about the firewall and what is required to administer its applications. A firewall is just another computer with special programs to deny anything except what is permitted. This is an easier task than denying what is not permitted and allowing everything else.

Encryption scrambles the data packets before they leave your network, headed for the ISP. This is done with software at your end of the circuit. A security algorithm scrambles the data packet based on the contents of a software key. The software key is a file of unintelligible binary bits. At the ISP end, the same algorithm uses the same key to unscramble the bits. In between, someone is sniffing the circuit gets files full of bits that do not mean anything.

There is also a methodology called authentication. Authentication is the process of someone filling out a form with his personal or financial information in the presence of someone who can be trusted. In all methods of security, no matte how complex or how secret, there is always trust in another human being required. When an authentication application is completed, it is entered into a database of trusted souls and then authentication keys are generated. The keys are distributed to organizations that will use the authentication and to the individual who filled out the form. Then when any type of transaction takes place, it is the key that informs the organization on the receiving end that this must indeed be who the data packet says he is.

10.9 Short Summary

There are two general methods you can use to secure your network against these dreaded forces: firewall and encryption.

A firewall involves using a computer between the Internet and your business's internal network.

Encryption involves using algorithms to scramble the data inside the data packets so that they are unintelligible to outsiders.

Authentication is the process of someone filling out a form with his personal or financial information in the presence of someone who can be trusted.

10.10 Brain Storm

1. Explain the domain name services.

2. Explain the security.

3. Explain the services of Internet.

Lecture 11

Virtual Server

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the virtual server concepts.
- Describe the evolution of ISP.
- Describe quality of service.
- Describe about ISP backbone.
- Describe the value added services.

Coverage Plan

Lecture 11

11.1 Snap Shot

11.2 Virtual Server

11.3 Evaluation of ISP

11.4 Quality of service

11.5 ISP backbone

11.6 Value added services

11.7 Short Summary

11.8 Brain Storm

11.1 Snap Shot

In this session we discuss about virtual server and ISP backbone.The virtual server uses the idea of sharing a server and a high-speed Internet connection with other companies and, therefore significantly reducing the cost of establishing an Internet presence. Most companies that have high-speed connections do not use the full potential bandwidth of their Internet connection continuously.

11.2 Virtual Server

A virtual server allows you to establish an Internet presence with a high-speed connection at a fraction of the cost. With a virtual server you get the Internet services you wan without the worries and costs of an Internet connection, and your customers and uses will not be able to tell the difference.

The virtual server uses the idea of sharing a server and a high-speed Internet connection with other companies and, therefore significantly reducing the cost of establishing an Internet presence. Most companies that have high-speed connections do not use the full potential bandwidth of their Internet connection continuously. However, to avoid delays to their customers, they will purchase the fastest Internet connection they can afford. For smaller companies this is usually a 56Kbps or fractional T1 frame-relay (time-shared) connection. With a virtual server and your customers will not be able to tell the difference.

The virtual server might not be for everyone. It allows a company to establish a high-speed Internet presence or test the potential of Internet marketing without incurring the high costs of bringing a high-speed connection and server to the office. In other words, it allows a company to test the Internet waters before committing high startup costs. If your company, after a period of time of being on a virtual server, has developed a very popular Web page, you will want to consider getting your own Internet connection and/or server to better support your customers' needs.

11.3 Evaluation of ISP

ISPs come in all sizes and orientations, and making the best selection require some analysis. Remember that this is a communication service that will carry critical information to your business, so choose accordingly. The following sections describe some of the factors to consider when making your decision.

Orientation

Many Internet providers target their services at consumers and individual users rather than at commercial use. Be sure to seek a provider focused on the needs of businesses. These needs include many of the criteria that follow.

Point of presence

The closer a network's point of presence is to your site the less expensive it will be to connect to the network from your site. Thus, it is advantages to use a provider with a large national even international base of POPs.

It is important to look at how a given Internet provider is connected to other components of the net. A provider that is directly connected to the high-speed Internet backbone circuits will offer better service, lower delays, and probably a lower cost than a provider that relies on other ISPs to reach a backbone access point. It makes sense to look for a provider with several direct connections to national and international components of the internet.

Network topology

A high-priority criterion to consider is the topology of the ISP's network. You will need to be familiar with network concepts to make a reasonable assessment of the topology. What you are interested in are weak spots in their network and what capacity is left at peak loading periods. The ISP should have no problem with showing you maps of its network. Just be sure the maps are of their physical network and not a virtual network. Virtual networks are created by the router switching and can use physical paths that are not owned by the ISP in question. An ISP that can explain to you how its network is physically configured in terms that are understandable is one worth considering closely.

11.4 Quality of Service

The Internet will carry your business's mission-critical traffic. It must be reliable and available, and deliver high performance. Let's look at each of these requirements in a bit more detail:

♠ Reliability – Recall that the Internet consists of multiple networks connected together. Your Internet provider is actually one component of the whole, but it is the component through which all your traffic will ride. Therefore, the level of reliability built into your provider's piece of the network is of direct relevance to you. Look for a provider with redundant equipment at all major switching hubs, and redundant backbone links so that no single failure will isolate part of the network. Since the NOC is such a critical element in any provider's network reliability, it should be backed up with an interruptible power supply (UPS), including a self contained (gas or diesel) generator.

♠ Availability – This is a measure of the percentage of time you can actually get to the network and get your information through. Getting to the network is easy if you have a dedicated connection, such as a leased line or a frame relay link. If you are dialing in, you wouldn't want to get a busy signal very often. There is an actual measurement that describes this phenomenon: p-grade of service. You should only consider networks with a p-grade of no more than p.05 – meaning that no more than 5 out of 100 calls result in a busy signal. Be sure to ask about it.

♠ Performance – Your customers, employees, dealers, and business partners will be communicating with via your Internet link. The last thing you want is a service with delays and low throughput. You can determine the degree of performance a given network will deliver by carefully examining a network diagram. Things to look for include the speed of the backbone, the speed at which large nodes connect to the backbone, and the speed at which smaller nodes connect to the backbone. A high-speed, high-capacity back-bone ensures a service with minimal delays and the ability to transmit bandwidth-hungry multimedia information with ease.

11.5 ISP Backbone

Depending on the size of your provider, it may or may not have a backbone, but it must have access to a backbone. Backbone speeds that are based on fiber-optic cable range from OC-1 at 52Mbps to OC-48 at 2500Mbps. There are four NAPs, or network access points, in North America. Network access points are the National Science Foundation – supported backbone Internet. The national ISPs like AT&T, MCI, and Sprint typically connect to all four NAPs. What about the ISP you are evaluating? You should be knowledgeable about network classifications and how they relate to the speed in bits per second. Your connection to the ISP could be offered as a T3 at 45Mbps, but its connection to a backbone could only be a T1, or 1.54Mbps. Your connection to the backbone national Internet will be no better than the T1 speed although you may be paying for a T3 speed up front. The better ISP will have many connections to the backbone that will result in enhanced reliability for your.

If the ISP says it has a backbone, check to make sure it is running at backbone speed now or if it plans to run at backbone speed. For instance, if glass fiber is the backbone, it is capable of virtually any speed of data throughput that you can think of up to the speed of light, and that is pretty fast. What limits the data throughput on glass is that electronics take data and convert it to light and put the general practice is to pull in glass fiber with gobs of capability and start out with just enough electronics to meet the need for a short period into the future. One reason for this is the cost of the electronics is decreasing and the quality and speed is increasing with time.

11.6 Value-Added Services

While a flexible set of access options is essential, businesses often require additional services to enable the Internet to be more friendly, useful, or secure. These are called value-added capabilities, and they tend to fall into five categories:

- Security options – Choose a provider who can accommodate your security needs, whatever they may be. Some applications require no security at all, while others may need a firewall, encryption, authorization, or some combination of these. You don't want to have to shop around for these options; a single vendor should be able to provide them to you, along with a consulting service for needs assessment and security policy development.

- Turnkey Web services – A Web site be an invaluable tool to your business. If you have no experience with the Web, or even if you do, you will want to select a provider who can build and / or host a high-impact site using raw information you provide.

- Domain names – Your company will be identified on the Internet by its address, or *domain name.* Make sure your provider can give you a domain name that is business oriented and reflects your corporate identity, rather than one that is cryptic, technology oriented, or associated with your provider's identity rather than your own.

- Network operation center – In case anything goes wrong, your provider of choice should operate a network operation center that is staffed 24 hours a day, 7 days a week, every day of the year. A courteous and responsive customer service staff must be available to respond to your calls. If you are a business user, you'll also want to make sure those representatives aren't the same ones tending to the high-volume time-intensive needs of consumer users. In case you desire training or specialized consulting services, those services should be available as well.

- Experience – How long has the provider you're considering been in the Internet services business? Is it the provider's main business, or a side element? How large is its customer base, and are those customers happy? The answers to these questions will help you separate those providers with real expertise from those just beginning in the complex world of providing Internet services.

11.7 Short Summary

Whether you are an individual who has an interest in connecting to the Internet or representing a 10,000-employee organization with a multi-megabit connection requirement, you are faced with the same decision; which Internet service provider? One thing is fairly certain. Once you have an Internet connection, no matter which you represent, it will be difficult to give it up. You will become dependent on the Internet to move your data just you depend on your automobile to move you about.

The Internet has become a new business tool that enables communications, especially inter-enterprise communications, to be used as a strategic competitive weapon. As they say, information is power, and when you know how to obtain and to use your Internet connection, you are heading for more power.

11.8 Brain Storm

1. Explain the virtual server concepts.
2. How to evaluate an ISP.
3. Explain the value added services.
4. Explain ISP Backbone.

Lecture 12

Internet Services

Objectives

After completing this lesson, you should be able to do the following

- Discuss about Internet services
- Describe the secure services
- Describe about web client and web server security issues.
- Describe about remote access.
- Discuss about real time conference services and administrative services.

Coverage Plan

Lecture 12

12.1 Snap Shot

In this session we discuss about the services of Internet. This lecture briefly summarizes the major Internet services may be used to interest. This lecture lists every Internet service. Such a list would incomplete as soon as it was finished and would include services of interest only to a few sites in the world.

12.2 Internet Services

There are a number of standard Internet services that users want and that most sites try to support. There are important reasons to use these services; indeed, without them, there is little reason to be connected to the Internet at all. But there are also potential security problems with each of them.

What services do you want to support at your site? Which ones can you support security? Every site is different. Every site has its own security policy and its own working environment. For example, do all your users need electronic mail? Do they all need to transfer files to sites outside your organization? How about downloading files from sites outside the organization's own network? What information do you need to make available for the public on the Web? What sort of control do you want over web browsing from within your site? Who should be able to log in remotely from another location over the Internet?

- ♠ World Wide Web access (HTTP)
- ♠ Electronic mail (SMTP)
- ♠ File transfer (FTP)
- ♠ Remote terminal access (Telnet or preferable SSH)
- ♠ Hostname / address lookup (DNS): Users generally don't use this service directly, but it underlines the other four services by translating Internet hostnames to IP addresses and vice versa.

All five of these services can be safely provided in a number of different ways.

12.3 Secure Services

Secure Services and Safe Services

You will occasionally hear people talk about "secure services". They are referring to services that give two kinds of guarantees:

1. The service cannot be used for anything but its intended purpose, and/or
2. Other people can't read or falsify transactions with the service.

That doesn't actually mean that you can use the service to do anything what over and still be safe. For instance, you can use Secure HTML to download a file, and be sure that you re downloading exactly the file that the site intended you to download, and that nobody else has read it on the way past. But you have no guarantee that the file doesn't contain a virus or an evil program. Maybe somebody nasty runs the site.

It is also possible to use "insecure" services in secure ways – it just has to be done with more caution. For instance, electronic mail over Simple Mail Transfer Protocol (SMTP) is a classic example of an "insecure" service. However, if you carefully configure your mail servers and encrypt message bodies, you can achieve the goals mentioned previously.

When you evaluate the security implications to your environment in your intended configurations – whether or not it's "secure" or "safe" in the abstract is not of any great interest.

The World Wide Web

The Web:

The collection of HTTP servers on the Internet. The Web is responsible, in large part, for the recent explosion in Internet activity. The Internet Engineering Task Force (IETF) is currently responsible for maintaining the HTTP standard and the World Wide Web Consortium (W3C) is developing successors to HTML. Nobody "controls" the Web, however, much as nobody "controls" the Internet.

HTTP:

The primary application protocol that underlines the Web: it provides users access to the files that make up the Web. These files might be in many different formats (text, graphics, audio, video, etc.), but the format used to provide the links between files on the Web is the HyperText markup Language. (HTML)

HTML:

A standardized page description language for creating web pages. It provides basic document-formatting capabilities (including the ability to include graphics) and allows you to specify hypertext links to other servers and files.

12.4 Web Client Security Issues

Web browsers are fantastically popular and for good reason. They provide a rich graphical interface to an immense number of Internet resources. Information and services that were unavailable or expert-only before are now easily accessible. In Silicon Valley, you can use the Web to have dinner delivered without leaving your computer except to answer the door. It's hard to get a feel for the Web without experiencing it; it covers the full range of everything you can do with a computer, from the mundane to the sublime with a major side trip into the ridiculous.

Unfortunately, web browsers and servers are hard to secure. The usefulness of Web is large part based on its flexibility, but that flexibility makes control difficult. Just as it's easier to transfer and execute the right program from a web browser than from FTP, it's easier to transfer and execute a malicious one. Web browsers depend on external programs, generically called viewers, to deal with data types that the browsers themselves don't understand. Netscape and Explorer now support a mechanism that allows third parties to produce plug-ins that can be downloaded to become an integrated and seamless extension to the web browser. You should be very careful about which viewers and plug-ins you configure or download; you don't want something that can do dangerous things because it's going to be running on your computers, as if it were one of your users, taking commands from a n external source. You also want to warn users not to download plug-ins, add viewers, or change viewer configurations, based on advice from strangers.

In addition, most browsers also understand one or more extension systems. These systems make the browsers more powerful and more flexible, but they also introduce new problems. Whereas HTML is primarily a text-formatting language, with a few extensions for hypertext linking, the extension systems provide many more capabilities; they can do anything you can do with a traditional programming language. Their designers recognize that this creates security problems. Traditionally, when you get a new program you know that you are receiving a program, and you know where it came from and whether you trust it. If you buy a program at a computer store, you know that the company that produced it had to go to the trouble of printing up the packaging and convincing the computer store to buy it and put it up for sale. This is probably too much trouble for an attacker to go to, and it leaves a trail that's hard to cover up. If you decide to download a program, you don't have as much evidence about it, but you have some. If a program arrives on your machine invisibly when you decide to look at something else, you have almost no information about where it came from and what sort of trust you should give it.

The designers of JavaScript, VBScript, Java, and ActiveX took different approaches to this problem. JavaScript and VBScript are simply supposed to be unable to do anything dangerous; the languages do not have commands for writing files, for instance, or general-purpose extension mechanisms. Java uses what's called a "sandbox" approach. Java does contain commands that could be dangerous, and general-purpose extension mechanisms, but the Java interpreter is supposed to prevent an untrusted program from doing anything unfortunate, or at least ask you before it does anything dangerous. For instance, a Java program running inside the sandbox cannot write or read files without notification. Unfortunately, there have been implementation problems with Java, and various ways have been found to do operations that are supposed to be impossible.

In any case, a program that can't do anything dangerous has difficulty doing anything interesting. Children get tired of playing in a sand box relatively young, and so do programmers.

ActiveX, instead of trying to limit a program's abilities, tries to make sure that you know where the program comes from and can simply avoid running programs you don't trust. This is done via digital signatures; before an ActiveX program runs, a browser will display signature information that identifies the provider of the program, and you can decide whether or not you trust that provider. Unfortunately, it is difficult to make good decisions about whether or not to trust a program with nothing more than the name of the program's source. is "Jeffs Software Hut" trust worthy? Can you be sure that the program you got from them doesn't send them all the data on your hand disk?

As time goes by, people are providing newer, more flexible models of security that allow you to indicate different levels of trust for different sources. New versions of Java are introducing digital signatures and allowing you to decide that program with specific signatures can do specific unsafe operations. Similarly, new versions of ActiveX are allowing you to limit which ActiveX operations are available to programs. There is a long way to go before the two models come together, and there will be real problems even then. Even if you don't have to decide to trust Jeffs Software Hut completely or not at all, you still have to make a decision about what level of trust to give them, and you still won't have much data to make it with. What is Jeffs Software Hut is a vendor you've worked with for years, and suddenly something comes around from Jeffs Software House? Is that the same people, upgrading their image, or is that somebody using their reputation? Because programs in extension systems are generally embedded inside HTML documents, it is difficult for firewalls to filter them out without introducing other problems.

Because an HTML document can easily link to documents on other servers, it's easy for people to become confused about exactly who is responsible for a given document. "Frames" are particularly bad in this respect. New users may not notice when they go from internal documents at your site to external ones. This has two unfortunate consequences. First, they may trust external documents inappropriately (because they think they're internal documents). Second, they may blame the internal web maintainers for the sins of the world. People who understand the Web tend to find this hard to believe, but it's a common misconception: it's the dark side of having a very smooth transition between sites. Take care to educate users, and attempt to make clear what data is internal and what data is external.

Web Servers

Although the preciously described applications provide an important role for the Internet, it is the World Wide Web(WWW) that can be considered the engine that has pulled the use of the internet into an exponential increase. The WWW represents a general collection of servers that can be interconnected via the use of the Hypertext Transfer Protocol (HTTP) and uniform resource locators (UURLs) URLs represent a mechanism to reference files located on the same or different servers, whereas HTTP represents the protocol used to transmit Web documents.

Web documents are created using the Hypertext Markup Language (HTML), resulting in hypertext documents being placed on Web servers. Here the term hypertext represents a method for presenting information where selected words in the text of document can be associated with a link to another document or an application, resulting in the display of the document or initiation of the application pictures, digitized voice, and video to be associated with a document. Thus, through the World Wide Web, a variety of documents, conveying different types of information can be exchanged.

The WWW represents a client/server-computing environment. The client, in the form of a Web browser, accesses different Web servers by the browser user specifying a URI or selecting one from an associated link when viewing a document.

Table lists five common URL prefixes. Although all browsers support HTTP, many browsers are modular in design and require a user to specify the locations of an applicable client program to support access to their URLs. For example, clicking a link associated with a Telnet URL would require the activation of a Telnet client that is not built into most browsers. Similarly, the ability to hear voice and watch video clips commonly requires the installation of audio and video add into a browser.

Table 1.1 Common URL prefixes.

Prefix	Description
http://	HTTP server
ftp: //	FTP server
file://	Local HTML file
telnet://	Telnet server
gopher: //	Gopher server

The use of the World Wide Web is literally exploding as companies install Web servers as a mechanism to make people aware of their products and services, as well as provide a catalog shopping mechanism, obtain orders, and provide employees with browsers to comparison shop, perform research, and search the internet for information. From a few hundred servers connected to the internet a few years ago, recent estimates place the number of Web servers at more than 200,000-and it is still growing. From IBM to Playboy, Disney to Coca-Cola, and Chrysler to numerous electronic utilities, just about every Major Corporation government agency, and university has a Web site that provides information about the goods or services they offer. Thus, the importance of the World Wide Web to include client browsers that are extremely easy to use provides a driving force behind the growth in the use of the Internet.

12.5 Web Server Security Issues

When you run a web server, you are allowing anybody who can reach your machine to send commands to it. If the web server is configured to provide only HTML files, the commands it will obey are quite limited. However, they may still be more than you'd expect; for instance, many people assume that people can't see files unless there are explicit links to them, which is generally false. You should assume that if the web server program is capable of reading a file, it is capable of providing that file to a remote user. Files that should not be public should at least be protected by file permissions, and should, if possible, be placed outside of the web server's accessible area (preferably by moving them off the machine altogether).

Most web servers, however, provider services beyond merely handing out HTML files. For instance, many of them come with administrative servers, allowing you to reconfigure the server itself from a web browser. If you can configure the server from a web browser, so can anybody else who can reach it; be sure to do the initial configuration in a trusted environment. If you are building or installing a web server, be sure to read the installation instructions. It is worthwhile checking the security resources mentioned in Appendix A, Resources, for problems.

Web servers can also call external programs in a variety of ways. You can get external programs from vendors, either as programs that will run separately or as plug-ins that will run as part of the web server, and you can write your own programs in a variety of different languages and using a variety of different tools. These programs are relatively easy to write but very difficult to secure, because they can receive arbitrary commands from external people. You should treat all programs run from the web server, no matter who wrote them or what they're called, with the same caution you would treat a new server of any kind. The web server does not provide any significant protection to these programs. A large number of third-party server extensions originally ship with security flaws, generally caused by the assumption that input to them is always going to come from well behaved forms. This is not a safe assumption; there is no guarantee that people are going to use your forms and your web pages to access your web server. They can send any data they like to it.

A number of software products are now appearing with embedded web servers that provide a convenient graphical configuration interface. Outsiders should carefully configure these products if they are running on systems that can be accesses. In general, their default configurations are insecure.

Electronic Mail and News

Electronic Mail

Electronic mail is one of the most popular network services. It's relatively low risk, but that doesn't mean it's risk-free. Forging electronic mail is trivial, and forgeries facilitate two different types of attacks.

- ♠ Attacks against your reputation
- ♠ Social manipulation attacks.

Accepting electronic mail ties up computer time and disk space, opening you up to denial of service attacks, although with proper configuration, only the electronic mail service will be denied. Particularly with modern multimedia mail systems, people can send electronic mail containing programs that run with insufficient supervision and may turn out to be Trojan horses (programs that appear to do something interesting or useful but are actually concealing hostile operations).

Although people worry most about deliberate attacks, in practice the most common problems with electronic mail are inadvertent floods and people who put entirely inappropriate confidence in the confidentiality of electronic mail and send proprietary data via electronic mail across the Internet. However, as long as users are educated, and the mail service is isolated from other services so that inadvertent or purposeful denial of service attacks shut down as little as possible, electronic mail is reasonably safe.

One of the earliest applications placed on the Internet was the transfer of messages between computers. This messaging, or electronic mail transmission, was primarily used in an academic environment until the development of graphical user interface based programs facilitated its widespread use. The initial series of programs used to transfer messages were text-based applications that used mnemonics and special terms to send and retrieve messages. The development of GUI programs considerably facilitated the creation and transmission of messages. In addition, although the addressing structure of the Internet provides a challenge for novices to correctly address and compose messages the development if address books, online directories, and similar features converted electronic mail into a Point and click and type operation.

Currently, electronic mail represents one of the most commonly used applications on the Internet. Today businesses and individuals exchange ideas, place orders, check schedules, and even transfer video and audio files, however, the later tow operations require special encoding because they involve the transfer to binary files for which the original interconnection of gateways linking different networks to one another is prohibited. Although many gateways now support 8 bit bytes, the lowest common denominator is represented by gateways restricted to passing 7 bit bytes. To overcome this limitation, various encoding schemes such as uuencode, xxencode, and MIME were developed. Today many if not most, people connected to the Internet can exchange messages regardless of the contents of files embedded or attached to a message. The resulting encoding and decoding operations enable digitized invoices, pictures of automobile fender benders, and plain, text based messages to be exchanged.

Usenet News

While electronic mail allows people to communicate, it's most efficient as a way for one person to send a message to another person, or to a small list of people interested in a particular topic. Newsgroups are the Internet counterpart to bulletin boards and are designed for many-to-many communication. Mailing lists also support many-to-many communication but much less openly and efficiently, because there's no easy way to find out about all mailing lists, and every recipient has his own copy of every message. The largest discussion mailing lists have tens of thousands of subscribers; the most popular newsgroups have at least hundreds of thousands. Usenet news is rather like television; there's a lot going on, most of it has little socially redeeming value, and some of it is fantastically amusing or informative.

File Transfer, File sharing and Printing

Electronic mail transfers data from place to place, but it's designed for small files in human readable form. Electronic mail transfer protocols are allowed to make changes in a message that are acceptable to humans for instance, inserting before the word "Form" at the beginning of a line so the mailer doesn't get it confused with a header line but are unacceptable to programs.

File Transfer

File Transfer Protocol is the Internet standard protocol for file transfer. Most web browsers will support FTP as well as HTTP and will automatically use FTP to access locations with names that begin "ftp", so many people use FTP without ever being aware of it. In theory, allowing your users to bring in files is not an increase of risk over allowing electronic mail; in fact, some sites offer services allowing you to access FTP via electronic mail. FTP is also nearly interchangeable in risk with HTTP, yet another way of bringing in files. In practice, however, people do use FTP differently from the way they use HTTP and electronic mail, and may bring in more files and / or larger files. What makes these files undesirable? The primary worry at most sites is that users will bring in Trojan horse software. Although this can happen, actually the larger concern is that users will bring in compute games, pirated software, and pornographic pictures. Although these are not a direct security problem, they present a number of other problems, and they are often used as carriers for viruses. If you make sure to do the following then you can consider inbound FTP to be a reasonably safe service that eases access to important Internet resources.

- ♠ Educate your users to appropriately mistrust any software they bring in via FTP.
- ♠ Communicate to users your site's guidelines about sexual harassment policies and organizational resource usage.

You may have heard of other file transfer protocols. Trivial File Transport Protocol (TFTP) is a simplified FTP protocol that diskless machines use to transfer information. It's extremely simple so that it can built into hardware, and therefore supports no authentication. There's no reason to provide TFTP access outside of your network; ordinary users don't transfer files with TFTP. File transfer represents a mechanism to distribute information in large chunks, typically beyond the capability of electronic mail. The File Transfer Protocol included in the TCP/IP protocol suite was similar to early electronic mail programs in which originally FTP applications were text based. This required FTP users to note a variety of commands and processes required to establish a connection to a distant computer, change directories, and locate and transfer one or more files. Similar to electronic mail applications, a variety of GUI based FTP application programs were developed during the early 1990s that facilitate the used of this application.

FTP is a client/server application, with the server supporting two basic types of connections, anonymous and restricted. Providing any client with the ability to read, write, or read and write to one or more directories one the addressed server sets up an anonymous connection.

To obtain this access capability, the person using the client application enters the term anonymous as his or her user identification. By convention although not standardized, the person also enters this or her electronic mail address as the password associated with the anonymous account. In comparison, a restricted connection requires the user to have an account established on the server prior to being able to access the server.

Currently, FTP represents a popular mechanism for corporations to provide the general public with access to press releases, software patches, online documentation and similar information. In fact, most hardware and software developers maintain an anonymous FTP account on their inter net FTP server, allowing the general public read access to retrieve a variety of files. Many times, vendor manuals refer readers to the vendor's FTP server to obtain the latest version of a driver, software patches, or general information about other products. Thus, FTP has obtained a considerable role as mechanism for technical information dissemination.

File Sharing

Several protocols are available for file sharing. Which allow computers to use files that are physically located on disks attached to other computers. This is highly desirable, because it lets people use remote files without the overhead of transferring them back and forth and trying to keep multiple versions synchronized. However, file sharing is much more complicated to implement than file transfer. File sharing protocols need to provide transparency and rich access. These features are what make file sharing desirable for users, but the needs to be transparent puts limits on the sort of security that can be implemented, and the need to provide rich access makes the protocols complex to implement. More complexity inevitably leads to more vulnerability.

The most commonly used file sharing protocols are the Network File System (NFS) under Unix, the Common Internet File System (CIFS) under Microsoft Windows and AppleShare on the Macintosh. CIFS is part of a family of related protocols and has a complex heritage, involving Server Message Block (SMB), NetBIOS/NetBEUI, and LanManager. You will see all of these names, and some others, used to refer to file sharing protocols, sometimes with radical security implications, they are interrelated and, for the most part, interoperable, and at the highest level, their security implications are similar. In fact, at the highest level, all of the file sharing protocols have similar implications for firewalls; they are all insecure and difficult to use across the Internet.

NFS was designed for use in local area networks and assumes fast response, high reliability, time synchronization, and a high degree of trust between machines. There are some serious security problems with NFS. If you haven't properly configured NFS, an attacker may be able to simply NFS-mount your file systems. The way NFS works, client machines are allowed to read and change files stored on the server without having to log into the server or enter password. Because NFS doesn't log transactions, you might not even know that someone else has full access to your files.

NFS does provide a way for you to control which machines can access your files. A file called /etc/exports lets you specify which filesystems can be mounted and which machines can mount them. If you leave a filesystem out of /etc/exports, no machine can mount it. If you put it in /etc/exports, but don't specify what machines can mount it, you're allowing any machine to mount it.

A number of subtler attacks on NFS are also possible. For example, NFS has very weak client authentication, and an attacker may be to convince the NFS server that a request is coming from a client that's permitted in the exports file. There are also situations where an attacker can hijack an existing NFS mount.

These problems are mostly due to the fact that NFS uses host authentication, which is easily spoofed. Because NFS doesn't actually work well across the Internet in any case, there isn't much point in allowing it between your site and the Internet. It creates a security problem without adding functionality.

Printing Systems

Almost every operating system these days provides remote printing – via lp or lpr on Unix machines, SMB printing on Windows machines, or AppleTalk print services on Macintoshes. Remote printing allows a computer to print to a printer that is physically connected to a different computer or directly to the network. Obviously, this is highly desirable in a local area network; you shouldn't need as many printers as you have machines. However, all of the remote printing options are insecure and inefficient as ways to transfer data across the Internet. There is no reason to allow them. If you have a need to print at a site across the Internet or to allow another site to use your printers, it's possible to set up special mail aliases that print the mail on receipt. This is the method many companies use even across in-house wide area networks because it's considerable more reliable.

12.6 Remote Access

There are many situations in which you would like to run a program on a computer other than the one that you're in front of. For instance, you may be in front of a slow computer because you're travelling with a laptop, or your other computer is a supercomputer, or you're using "thin clients" – purposefully stupid computers – in order to lower maintenance costs and get economies of scale. Originally, remote access meant some form of remote terminal access, which allows you to use character-based application. These days, character-only access is rarely sufficient. Instead, you may need some form of remote graphics.
The general questions about remote access are the same for all methods.

- ♣ Are there appropriate controls on who can access the machine remotely? How are remote users authenticated?
- ♣ Can anybody take over a connection that's in progress?
- ♣ Can eavesdroppers pick up important information (particularly, authentication information)?

Remote Terminal Access

A third popular application that was a driving force behind the growth of the Internet is remote terminal access. Server applications were developed that enable remote users, to connect to hosts as if their computers were directly connected to the host instead of the two being geographically separated from each other. Two of the most, popular remote terminal access applications are Telnet and TN3270.

Telnet represents a TCP/ip application developed to enable a terminal device to include personal computers to remotely access an ASCII host. Since different hosts use different escape character sequences to position information on the screen, a Telnet application typically supports a number of terminal emulation, which in turn support the codes generated by a particular host. Common terminal emulation include, TTY, VT52, VT100, VT220,VT320, VT340,TB950,WYSE50, and WYSE60, which enable a PC to function as one of these terminals. In comparison, TN3270 represents a special type of terminal communications

program that is used for accessing IBM S/370 and S/390 hosts. Tn3270 emulates certain types of IBM terminals and enables a PC executing a
TN3270 emulation program to correctly interpret the screen codes generated by IBM hosts.

By obtaining the ability to remotely access a host computer, users are not restricted to their computer centers when they need mainframe-based information. Through the use of telnet or TN3270 they can use the internet to access the corporate host, check sales figures and delivery dates, obtain price quotations, and perform similar operations without having to return to the office. Thus, Telnet, provides a mechanism to increase employee productivity.

Remote Graphic Interfaces for Microsoft Operating Systems

Although Windows NT provides clients for most of the remote execution services described previously, and servers for many of them are available as part of the resource kits or third-party products, remote terminal services in general aren't very interesting on Windows NT. While there are character-oriented programs that will allow you to do many administrative tasks, most of the programs people want to use are graphical.

Microsoft provides remote graphical interfaces as part of Windows 2000 servers, in a package called Terminal Services. This is also available for Windows NT 4 as a special Terminal Server edition of the operating system. Terminal Services and Terminal Server both use a Microsoft-developed protocol called Remote Desktop Protocol (RDP) to communicate between clients and servers.

TCP/IP-based remote access is also available from almost every other remote access program in the Windows market, including LapLink, RemotelyPossible, and PcANYWHERE, to name only a few. There is also the controversial program BO2K, which is a freely available open source program that provides remote access. It is controversial because it is widely distributed as a tool for intruders, designed to provide remote access to outsiders; on the other hand, it is a full-featured and effective tool to provide legitimate remote access as well.

These programs differ widely in their security implications, although most of them are unfortunately insecure.

Network Window Systems

Most Unix machines currently provide window systems based on the X11 window system. X11 servers are also available as third-party applications for almost every other operating system, including all versions of Microsoft Windows and many versions of MacOS. X11 clients are rarer but are available for Windows NT. Network access is an important feature of X11. As more and more programs have graphical user interfaces, remote terminal access becomes less and less useful; you need graphics, not just text. X11 gives you remote graphics.

X11 servers are tempting targets for intruders. An intruder with access to an X11 server may be able to do any of the following types of damage:

Get Screen dumps

These are copies of whatever is shown on the users' screens

Read keystrokes

These may include users' passwords.
Inject keystrokes

They'll look just as if they were typed by the user. Imagine how dangerous this is could be in a window in which a user is running a root shell.

Originally, X11 primarily used authentication based on the address that connections came from, which is extremely weak and not suitable for use across the Internet. These days, most X11 servers implement more secure authentication mechanisms. However, just like Telnet, X11 is still vulnerable to hijacking and sniffing, even when the authentication is relatively secure, and solving the overall security problem requires that you encrypt the entire connection via SSH or a VPN solution.

12.7 Real-Time Conference Services

A number of different real-time conferencing services are available on the Internet, including talk, IRC, web chat rooms, and the various services provided over the Multicast backbone (MBONE). All of these services provide a way for people to interact with other people, as opposed to interacting with databases or information archives. Electronic mail and Usenet news are designed to facilitate asynchronous communications; they work even if the participants aren't currently logged in. the next time they log in, the email messages or news postings will be waiting for them. Real-time conferencing services, on the other hand, are designed for interactive use by online participants.

Internet Relay Chat (IRC) is sort of like Citizens Band (CB) radio on the Internet; it has its own little culture involving lots of people talking at each other. Users access IRC via dedicated IRC clients, or by using Telnet to access a site that provides public IRC client service. IRC servers provide hundreds (sometimes thousands) of named "channels" for users to join. These channels come and go (anyone can create a new channel, and a channel survives as long as there's any one on it), although some popular channels are more or less permanent. Unlike talk, which is limited to a pair of users, any number of people can participate on an IRC channel simultaneously. Some IRC clients allow a user to participate in multiple channels simultaneously (sort of like taking part in two different conversations at once at a party).

There are a number of security problems with IRC; most of the problems aren't with the protocol itself, but with the clients, and with who uses IRC and how. Many of the clients allow servers far more access to local resources than is wise; a malicious server can wreak havoc with a weak client. Further, many of the most frequent users of IRC are pranksters and crackers who use IRC to pass technical information among them and to try to trick other IRC users. Their idea of a fine time is to tell some neophyte IRC user "Hey, give this command to your IRC client so that I can show you this neat new toy I wrote". Then, when the unsuspecting user follows the prankster's directions the commands trash the system. Anyone using IRC needs a good client program and a healthy close of wariness and suspicion.

Purely web-based chat rooms have less vulnerability, but HTTP doesn't lend itself well to chatting, so these tend to be clunky and uncomfortable to use. People therefore have developed a number of hybrid solutions using plug-ins to HTTP clients. These provide much nicer interfaces but also introduce new vulnerabilities. Like IRC, they have many "bad neighbourhood" where people hand out looking for neophytes they can trick or attack. In addition, the protocols and the plug-ins themselves are often vulnerable.

More complicated systems allow richer conversations. As high-speed network connections become common, full-fledged video conferencing systems have become popular, even across the Internet. The most famous of those systems is Microsoft's NetMeeting and most other video conferencing systems in wide use are based on a set of international telecommunications union standards and protocols fir video conferencing. These protocols are extremely difficult to secure. They have almost every feature that makes a protocol different to protect, including using multiple data streams initiating data transfer from both ends of the conversation (instead of having a clearly defined client and server) using connectionless protocols, and dynamically assigning port numbers instead of using well known port numbers. While whey can be very useful, providing them securely requires an extremely specialized firewall. Because video conferencing involves large amounts of data, the firewall also needs good performance.

The MBONE is the source of a new set of services on the internet, focusing on expanding real time conference services beyond text based services like talk and IRC to include audio, video, and electronic whiteboard. The MBONE is used to send real time video of many technical conferences and programs over the internet (e.g internet engineering task force meetings, keynote sessions from USENIX conferences, space shuttle flight operations, and so on). At this point, the commonly used MBONE services appear to be reasonably secure. Although there are theoretical problems, the only reported attacks have been floods, which are easy to deal with. Theoretical problems have a way of eventually becoming actual problems, but these are extremely theoretical (nobody has verified that they are actually exploitable at all0 and not very threatening(if they were exploitable, they still wouldn't be catastrophic). Unintentional denial of service can be a real concern with the MBONE however, because audio and video can use so much band width.

Naming and directory services

A naming service translates between the names that people use and the numberical address that machines use. Different protocols use different naming services, the primary protocol used on the internet is the domain name system which converts between hostnames and IP address.

Using DNS internally and then relying on hostnames for authentication makes you vulnerable to an intruder who can install a deceitful DNS server. This can be handled by a combination of methods, including:

- Using IP addresses (rather than hostnames) for authentication of services that need to be more secure.

- Authenticating users instead of hosts on the most secure services because IP address can also be spoofed.

Windows 2000 networks use DNS in conjunction with the Active Directory service to locate resources. Clients access the active directory service via the lightweight directory access protocol, which is a widely used standard for access to directory information.

As WINS has evolved, the interrelationship between it and DNS has become ever more complex and confusing. WINS servers can consult DNS servers, and Microsoft DNS servers can consult WINS servers. The important things to remember about WINS are:

- WINS is designed as a purely internal protocol for a single organization.
- There are scaling issues using WINS on large and complex networks, even for a single organization.
- Microsoft is phasing out use of WINS in favor of DNS.
- WINS is less secure than DNS.

WINS has all the security issues that DNS has, and then some. First, WINS contains more information than DNS does. While DNS contains information, like hostnames, that you might not want an attacker to have, WINS contains information, like valid usernames and lists of running services, that you definitely don't want an attacker to have. Second, WINS is designed around dynamic registration; not only does it accept queries from hosts, it accepts new data from the network. This makes it much more vulnerable than DNS to hostile clients. Making WINS visible to the Internet is highly dangerous and not at all useful.

Authentication and Auditing Services

Another important (although often invisible) service is authentication. Authentication services take care of assigning a specific identify to an incoming connection. When you type a username and a password, something is using these to authenticate you – to attempt to determine that you are the user that you say you are Authentication may occur locally to a machine or may use a service across the network. Network services have the advantage of providing a centralized point of administration for multiple machines, and therefore a consistent level of trustworthiness.

A number of different services provide authentication services, sometimes combined with other functions. Under Unix, the most common authentication services are NIS and Kerberos. Windows NT normally uses NTLM, while Windows 2000 uses Kerberos by default, falling back to NTLM only for access to older servers. For various reasons, these protocols can be difficult to use across the Internet or for authenticating people who wish to connect over telephone lines, so two protocols have been developed for just this situation, RADIUS and TACACS.

12.8 Administrative Services

A variety of services are used to manage and maintain networks; these are services that most users don't use directly – indeed, that many of them have never even heard of – but they are very important tools for network managers.

System Management

Simple Network Management Protocol (SNMP) is a protocol designed to make it easy to centrally manage network devices. Originally, SNMP focused on devices that were purely network-oriented (routers, bridges, concentrators, and hubs, for instance). These days, SNMP agents may be found on almost anything that connects to a network, whether or not it's part of the network infrastructure. Many hosts have SNMP agents; large software packages, like databases, often have specialized SNMP agents; and even telephone switches and power systems have network interfaces with SNMP agents.

Modem SNMP agents often contain extremely sensitive data; the default SNMP agent for Windows NT includes the complete list of valid usernames on the machine and a list of currently running services, for instance. Many SNMP agents allow for machine reboots and other critical changes. Unfortunately, they are hardly secured at all. SNMP security currently relies on a clear text password, known as a community string, with a well-known and widely used default. Some SNMP agents implement additional levels of security, but these are still insufficient for extremely sensitive data. Allowing SNMP from the Internet is extremely dangerous.

With the introduction of SNMP v3, which provides better authentication and can encrypt data, it is becoming possible to run SNMP more securely: however, SNMP v3 is not yet widespread.

Routing

Routing protocols like RJP and OSPF are used to distribute information about where packets should be directed. Transactions on the Internet involve hosts distributed across the world, which are added, moved, and deleted, all without a single central authority to control them. The Domain Name System provides part of the information necessary to make this work, and another critical part is provided by routing services, which distribute information about which numbers are where and how to get to them.

The goods news is that routing information rarely needs to go to any significant number of hosts; in general, you will have at most a few routers that talk to the Internet, and those will be the only hosts that need to talk routing protocols to the Internet. In general, you will not need to pass routing protocols through firewalls, unless you are using internal firewalls inside a site.

Network Diagnostics

The two most common network management tools are ping and trace route (also known as tracert). Both are named after the Unix programs that were the first implementations, but both are now available in some form on almost all Internet capable platforms. They do not have their own protocols but make use of the same underlying protocol, the Internet Control Message Protocol (ICMP). Unlike most of the programs we've discussed, they are not clients of distinguishable servers. ICMP is implemented at a low level as a required part of the TCP/IP protocols all Internet hosts use.

ping simply tests reach ability; it tells you whether or not you can get a packet to and from a given host, and often additional information like how long it took the packet to make the round trip. Trace route tells you not only whether you can reach a given host, but also the route your packets take to get to the host; this is very useful in analyzing and debugging network trouble somewhere between you and some destination.

Time Service

Network Time Protocol (NTP), an Internet service that sets the clocks on your system with great precision, has clients on most operating systems. Synchronizing time among different machines is important in many ways. From a security point of view, examining the precise times noted on the log files of different machines may help in analyzing patterns of break-ins. Having synchronized clocks is also a requirement for preventing attackers from recording an interaction and then repeating it; if timestamps are encoded in the interaction, they will be incorrect the second time the transaction is replayed.

You do not have to use NTP across the Internet; it will synchronize clocks to each other within your site, if that's all you want. The reason that people use NTP from the Internet is that a number of hosts with extremely accurate clocks – radio clocks that receive the time signal from master atomic clocks or from the atomic clocks in the Global Positioning System (GPS) satellites – provide NTP service to make certain that your clocks are not only synchronous with each other but also correct. Without an external time service, you might find that all your computers have exactly the same wrong time. Accepting an external service makes you vulnerable to spoofing, but because NTP won't move the clocks very far very fast, a spoofed external clock is unlikely to make you vulnerable to a playback attack, although it could succeed in annoying you by running all your clocks slow or fast. Radio or GPS clocks suitable for use as NTP time sources are not terribly expensive, however, and if you are using NTP to synchronize clocks for an authentication protocol like Kerberos, you should by your own and provide all time service internally, instead of using an external reference.

Databases

For a long time, databases were relatively self-contained; most accesses to a database system were from the same machine that was running the software. These days, databases are very rarely self-contained. Instead, they are the data storage for larger, distributed systems, sales information systems, e-commerce systems, even large electronic mail systems all use databases and communicate with them over networks.

This makes secure remote communication with databases more important than ever. Unfortunately, database communication protocols tend to be propriety and different for each database manufacturer. Furthermore, they've only recently been designed with any concern for security. It is unwise to pass database transactions unprotected across the Internet.

Games

Games produce some special security challenges. Like multimedia protocols, they have characteristics that make them inherently difficult to secure; they're trying to make flexible, high-performance connections. Games also change frequently, are designed by people more interested in attractiveness then security, and are a favorite target of attackers. In general, you should avoid supporting game play through a firewall. There is no network security risk in running multiplayer games internal to a network.

12.9 Short Summary

Web servers can also call external programs in a variety of ways. You can get external programs from vendors, either as programs that will run separately or as plug-ins that will run as part of the web server, and you can write your own programs in a variety of different languages and using a variety of different tools

Transactions on the Internet involve hosts distributed across the world, which are added, moved, and deleted, all without a single central authority to control them. The Domain Name System provides part of the information necessary to make this work, and another critical part is provided by routing services, which distribute information about which numbers are where and how to get to them.

12.10 Brain Storm

1. Distinguish the web server and web client security issues.
2. Explain the services of Internet.
3. Explain the administrative services.
4. What is remote terminal access?

Lecture 13

Router Technology

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the router technology.
- Describe the network fundamentals.
- Describe Internet routing.
- Describe about routing protocols and routing software.
- Describe about mobile routing.

Coverage Plan

Lecture 13

13.1 Snap Shot

Prior to the late 1970s, centralized data processing prevailed, where a mainframe was connected via clusters of dumb terminals. The need for interhost communication was scarce, and the capability was very limited. In the IBM 360/370 world, data communication typically means an SNA network using multidrop lines, with X.25 as the remote communication network protocol.

The introduction of personal computers during the late 1970s created the need for sharing peripheral devices such as printers and disk space. The LAN enables PC users connected to the shared resources. Furthermore, the emergence of minicomputers drove the need for the distributed data processing.

During the 1980s, corporations across the world were beginning to invest in their own infrastructures. Information was regarded as an important asset. In order to enhance the corporate competitive edge, information must move freely among users within a company; to some extent this information infrastructure has been extended to customers and vendors. The Internet is the premier example of an inter-enterprise network. WAN and MAN tie the computers and their peripheral devices together.

13.2 Router Technology

Routed protocols of particular interest include IP, previously discussed. Commonly implemented router protocols can be classified into two categories: interior routing and exterior routing protocols. Interior routing protocols, such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Interior Gateway Routing Protocol (IGRP), and exterior routing protocols, such as Exterior Gateway Protocol (EGP) and Border Gateway Protocol (BGP), will be studied.

Data processing applications have become more sophisticated. The client/server model often means that a network of machines is connected to run communication protocols such as TCP/IP. In this age of multimedia, data objects (data records, video/audio files, graphics and images, and large files) must be delivered to remote machines regardless of the host computers' operating systems, hardware platform, and physical network connection. There has been an explosion in the use of the World Wide Web. In short Today's challenge for communications professionals is to mesh the global internet work into a seamless information repository for users so that they can "get it, move it, use it!"

13.3 Network Fundamentals (OSI Layers)

In order to simplify data movement among various end systems, it is pragmatic to modularize the communication mechanisms, so that each module behaves as a predefined communication standard. There are seven layers of hierarchical communication levels. Each layer provides a service to the layer above it. For example, the network layer provides a service to the transport layer, and the transport layer presents "data" to the internetwork system.

The network layer encapsulates the data, passed down from the transport layer, within a *network header*. The header contains information such as source and destination *logical* addresses.

The data link layer adds frame header and trailer to the encapsulated packet passed down from the network layer. The frame header contains the physical addresses.

The physical layer provides the physical transmission medium (such as Dsi, DS3, SONET) connections. Multiplexers, repeaters, and hubs are typical internetworking devices at the physical layer.

Hubs

The length of the interconnecting cable limits physical connections among hosts. The Distance between any two nodes must be within a cable-specific range or the signal will be diminished. A hub (or a multi-port repeater) is a physical layer device that receives, amplifies and retransmits the signals to the other end. Hubs generally are inexpensive devices. Because the hub amplifies the signals, it also copies the signal disturbance of errors. In a LAN environment, the number of repeaters placed between segments should be limited to three. Another alternative is to use bridges.

Bridges

As networking technology proliferated, users required solutions that would allow them to add more users to the network, extend network distances, and translate data between different LAN technologies. These requirements essentially created demand for the local bridge. A bridge is an internetwork component designed to interconnect LANs to perform the function of a single large network. Bridge technology supports layer 2 (data link layer) protocols, namely, logical link control (LLC) and medium access control (MAC)[2.] A bridge interconnects multiple networks, providing communications between networks. It examines the destination address of each frame received. If the source and destination are on the same network segment, the bridge filters the frame. If the destination is on a different network segment, the bridge forwards the frame. The protocols that support the bridge technology dictate the flow control, error handling, addressing, and media access algorithms. Examples of popular data link layer protocols include Ethernet, Token Ring, and FDDI (fiber distributed data interface).

The decision to keep traffic local or to forward the frame is made without regard to upper-layer protocols. Bridges can therefore be a cost-effective solution for physical-homogeneous, but protocol-heterogeneous, network environments. Bridges can also interconnect different physical layer media types such as coaxial cable, fiber-optic, or twisted-pair networks within the same data link layer.

Routers

The next technological challenge in the LAN arena involved connecting LANs to other geographically separate LANs over a wide area service. Routers provide connectivity by appropriately filtering and forwarding (routing) packets of information. Routers are more intelligent than bridges in that routers operate at the OSI model layer 3 (network layer). They are used to interconnect IP subnetworks or other subnetworks. Routers can also serve as the functions of hub (physical layer) and bridge (data link layer). Router technology is more sophisticated than bridge technology because routers can operate in bridging mode to handle LAN traffic and in routing mode to connect LANs to WANs.

Besides routing packets from source to destination, routers also can provide value-added features. The routing intelligence makes the value-added features feasible. These features include:

- Traffic sequencing
- Security filtering
- Network management
- Resource maintenance

Modern routers can parse the entire network layer header, thereby supporting sophisticated filtering (this was not always the case, since only a few fields could be examined until recently). Routers now forward an aggregate of 100,000 to 500,000 64-byte PDUs per second (fewer if the PDUs are of maximum data link layer level). Just the forwarding effort can place a burden on the router, but the addition of complex filtering further taxed the processor. The deep reading of the header and its storage translates into large amounts of RAM and flash memory on the router.

Typically, routers are required to support multiple protocol stacks each with its own routing protocols, and to allow these different environments to operate in parallel.

The router supports different data communication access media: FDDI, Token Ring, Ethernet, private and public networks, and switched LANs.

Routing algorithms are used to propagate connectivity information. While the particulars of each routing algorithm differ, there are several goals each routing algorithm attempts
To achieve. These include:

- Optimality. A routing algorithm should always choose the best routes, based on the route selection criteria.

- Rapid convergence. A routing algorithm must behave not only with robustness in the event of topology changes, it must also allow other routers to know the situation quickly so that a decision can be made.

- Robustness. If the network topology changes, it products the new best paths.
- Simplicity. It should minimize network bandwidth and router processing overhead.

Router algorithms can be classified by type, as in the following examples.

Static Versus Dynamic Routing

The decision process of routing packets from one router to another can be either static of dynamic. For the static routing, the path between source and destination is fixed. The static routing algorithm is typically adopted where network traffic is relatively predictable and network design is relatively simple. Static routing tables do not adjust to the network changes, so they should be used only when routes do not change or the destination host can be reached only through one route.

Static routing, by definition, is not much of a "routing algorithm" because of the fixed nature of the routing. Dynamic, on the other hand, adjusts to the changing network traffic or topology and makes the hopping decision in real time. Dynamic routing algorithms must be capable of adjusting their routing tables quickly when adverse circumstances (interrupted circuits, congestion etc.) occur. A router's dynamic routing table is built from the information exchanged by routing protocols. Routing protocols handle complex routing topology more quickly and accurately than the system administrators can.

Distance vector routing versus link state routing

Distance vector algorithm was the original form of dynamic routing. This algorithm requires that a router maintain a table of topology information that allows the algorithm to determine the direction (vector) and the distance to any link in the internet works.

The topology information indicates the numeric measurement (a metric) from any router to any other routers. The measurements can be a simple bandwidth, a hop count, or number of routers that must be traversed to arrive at the destination. Once a router calculates each of its distance vectors, it sends the information to its neighbors on a regular basis (e.g., once a minute). If network topology changes, the receiving router will modify its routing table and transmit it to each of its neighbors.

The distance vector routing algorithm does not allow a router to know the exact topology of an internet work. Because the node realizes only the cost of its adjacent routers, the cost of reaching remote routers is derived. Therefore for the distance vector routing encounters the slow convergence problem. This implies that information about rout failures has a propagation delay (e.g., 30 seconds). The consequence of this is that there may be a period of vulnerability.

Link-state routing is a newer form of dynamic routing. Each router maintains a database of the entire network topology. The information is somewhat analogous to a road map with you-are-here pointer showing a map reader's current location.

Each router sends out a link-state packet (LSP) that describes the current state of all its links. Sending the LSP is commonly referred to as route advertisement. An LSP is broadcast to all routers at a much longer interval than its counterpart in the distance vector routing algorithm.

Link-state routing eliminates many of the problem inherent in distance vector routing. Because each router knows the complete network topology, routing loops will not be formed. Link-state routing converges much faster than distance vector routing. Link –state routing sends routing updates only every half hour or when there is a change in network topology, whereas distance vector routing updates every minute or so. Link-state routing thus consumes much less bandwidth at the cost of increased route calculation processing.

13.4 Internet Routing

The internet technology is the result of research funded by the Advanced Projects Research Agency. The technology includes a set of network standards that specify the details of how computers communicate, as well as a set of conventions for inter connecting networks and routing traffic. Commonly referred to as TCP/IP, it can be used to communicate across any set of interconnected networks. Later, TCP/IP was included with the Barkeley Software Distribution (BSD) of UNIX.

The Internet protocol suite includes not only TCP and IP protocols, but also specifications for such common applications as mail, terminal emulation, and file transfer.

13.5 Routing Protocols

In order to route the packets from host A to host B, a set of rules must be established to coordinate the network topology and make intelligent routing decisions. Routing protocols are designed not only to switch to a backup route when the primary routes become severed, they are also designed to decide which route to a destination is “best”. On any network where there are multiple paths to the same destination, a routing protocol should be used.

Autonomous system

An Autonomous system (AS) is one particular group of networks under the same administrative authority and control. An autonomous system, sometimes referred to as a routing domain, is not merely an independent network. It is a collection of networks and gateways with its own internal mechanism for collecting routing information and passing it to other independent network systems. Within an autonomous system, a protocol must be used for route discovery, propagating, and validating routes. Those that operate within the same routing domain use called Interior Gateway Protocols (IGPs). The major IGPs include RIP, HELLO, and OSPF. A routing domain exchanges routing information with other domains using Exterior Gateway Protocol or Border Gateway Protocol. Cisco's proprietary interior routing protocol, Interior Gateway Routing Protocol (IGRP), has been in the marketplace the longest, has proved the most popular among router users, and can be used both within a domain and between domains.

Routing loops (bouncing effect)

If all routers in an internet-work do not have up-to-date, accurate information about the actual state of the network topology, incorrect routing information may be used to make a routing decision. Using incorrect information may cause packets to take less-than-optimum paths or paths that return the packets they have already visited.

Hold-down technique

Distance vector routing algorithms are self-correcting, but it takes a long time before the loop is detected. This problem can be avoided by using a technique called hold-down. Hold-down works as follows: When B realizes that N is not reachable, a hold timer is activated. If N becomes reachable again before the timer expires, B removes the timer and a notable update occurs. If an update arrives from either A or D (its neighbors) with a better cost than originally recorded, B removes the timer and updates the table-N is accessible! If A or D sends B a worse cost (in this case, the cost is 3) than B's original cost (it was 1 before the break), the update is ignored. When 30 seconds expires, B will inform A and D of the new disruption status. The loop, therefore, is avoided.

Interior routing protocols

As indicated in the previous section, there are two different ways of making decisions of route packets from source to destination for the routing protocols (1) distance vector algorithm and (2) link-state algorithm. In this section, we will be looking into three interior routing protocols: Routing Information Protocol, the Interior Gateway Routing Protocol, and the Open Shortest Path First. RIP is much simpler but less powerful than OSPF and IGRP. IGRP is the Cisco proprietary protocol, but given the fact that the IGRP has been in the marketplace the longest and Cisco has a significant router market share, IGRP is very well known within the routing industry.

Exterior routing protocols

Exterior Gateway Protocol (EGP): The EGP is an interdomain reachability protocol used in connecting internet backbone routers. EGP does not use routing metrics and therefore cannot make true routing decisions; it only exchanges reachability information. It reports which network is available through which routers.

EGP is the first exterior routing protocol. Internet has come a long way since the early years and EGP is the legacy of old ARPANet days. The new exterior protocol is Border Gateway Protocol (BGP).

Border Gateway Protocol (BGP): BGP represents an attempt to address the most serious of EGP’s problems. BGP is an interdomain routing protocol created for use in the Internet. Unlike EGP, BGP was designed to detect routing loops and exchange routing updates. BGP and other exterior protocols are replacing EGP in the Internet. BGP uses TCP as its transport mechanism.

BGP routing updates consist of a combination of network number and autonomous systems (AS) path pairings. The AS path is simply a list of all autonomous systems that must be traversed to reach a specific network. Because BGP lists the complete route to the destination, routing loops and slow convergence issues are avoided.

13.6 Routing Software

Another routing trend lies in the routing software development. Software routing (sometimes referred to as source routing) means that there is no need to have dedicated router hardware to store and update the addresses. The software may reside in a general-purpose PC or within the same host. The software gives the host machine the intelligence to calculate and make routing decision. Software routers choose the better routes more often because they typically discover all possible routes to the destination before the packet is actually sent.

Traditionally, the routing functionality, such as addressing, multicasting, and packet delivery capabilities, are handled in stand-alone hardware. There is a push to place those functionalities in general purpose computer software offerings.

The trend of going software can be seen by the alliance of the computer/router industries' major players: (1) computer vendor Compaq is developing the network product that will include Cisco System's routing technology; (2) Microsoft is teaming up with Bay Network's routing technology into Windows NT server and Cairo platform. Separately, Novell has been the software pioneer with a highly sophisticated routing product – NetWare Link Service Protocol (NLSP). Cisco also offers a software protocol named LAN2LAN.

There are a few drawbacks to software routing

- ♠ Some of the host routing is housed in the PC, which may not have the necessary processing power to handle the traffic.
- ♠ Software routing would not handle all or most of the routing protocols.
- ♠ Software routing may not be as configurable as the hardware router.

Another trend is to separate the routing decision from the forwarding decision. This model, evolving under the MultiProtocol Over ATM (MPOA) work of The ATM Forum uses a centralized route server, along with cut-through methods over ATM.

13.7 Mobile Routing

Currently there are some methods of allowing portable computer users to achieve mobile computing. They all involve static technology, which is unsuitable for the Internet routing. Mobile IP requirements include the following:

1. A mobile host should be capable of continuing to communicate, using the same IP address, after it has been disconnected from the Internet and reconnected at a different point.

2. A mobile host should be capable of interoperating with existing hosts, routers, and services.

3. No weakening of IP security should occur at the router level. If a legitimate mobile user can mover around use the same IP address, an intruder might be able to configure his or her address accordingly. Routers at the Internet access path must have a tight authentication process.

4. Multicasting should be possible while the mobile host machine is not stationary (i.e., maintain the location of the mobile hosts at the router level such that broadcasting messages will be delivered).

There are two major phases in the mobile IP protocol

1. Beaconing protocol. When the mobile host moves around, it can “realize” that it is moving by listening to beacons on the radio. At some point it will discover that the beacons from a new base are louder and clearer than those of the previous base; then it is time to switch.

2. Registration procedure. Through the beaconing protocol, the mobile host has discovered the IP and “media” address of the new base. It must now register with this new base and obtain its agreement to relay packets.

Currently, the IETF is researching this subject. There are several difficult technical tasks that are waiting to be resolved. In the mobile environment, throughput and delay may vary considerably. Roaming from one cell to another may cause temporary losses of connectivity. TCP will react by reducing its sending rate to a minimum. When the connectivity is reestablished, that sending rate will increase. In the cellular voice world, the quality-inconsistency problem may be acceptable. With the commercialization of Internet, the psychological and legal aspects have yet to be studied. For the technical part, the efforts will undoubtedly emphasize mobile transmission control and routing algorithm design.

13.8 Short Summary

This lecture briefly described routing technology. The reason for the internet-focused flavor is because the Internet has radically changed the way corporations do business. Voice communication technology is near maturity, and the major growth area will be in data communication.

Repeaters and hubs are physical layer relays in the OSI terminology; bridges are data link layer relays; and routers are network layer relays.

13.9 Brain Storm

1. Explain the router technology
2. Explain hub, router, and bridges.
3. Explain the routing software.
4. Explain the protocols for routing.

Lecture 14

Web Server

Objectives

After completing this lesson, you should be able to do the following

- Discuss about web's roots.
- Describe the web servers.
- Describe about transport issues.
- Describe about web access.
- Describe the connection modes and publishing web server.

Coverage Plan

Lecture 14

14.1 Snap Shot

14.2 Web’s roots

14.3 Web servers

14.4 Transport issues

14.5 Web access

14.6 Connection modes

14.7 Publishing web server

14.8 Short Summary

14.9 Brain Storm

14.1 Snap Shot

In this lecture we will discuss the web technology is providing an effective way to automate a number of business tasks and processes within the enterprise for use in intranets. The web is probably the most interesting use of truly distributed client/server technology available today. There are many different web server programs available. These servers are called HTTP daemon, or HTTP servers. The most prominent ones are available from the national center for supercomputing applications and CERN and Netscape Communications Corporation.

14.2 Web's Roots

Work on Web technology took place in the late 1980s as a way to simplify and systematic access to Internet-resident information.

In 1989, Tim Berners-Lee and Robert Calliau created the World Wide Web (also known as the Web, or W3) at CERN's Internet facility. The project was to be used as a means of transporting research and ideas effectively throughout the CERN organization. It merged the technologies of information retrieval and hypertext to make an easy but powerful global information system. This idea not only worked for the benefit of all. The Web is now the newest medium of expression. Its growth rate has been phenomenal.

The project's goal has been and continues to be to build a distributed hypermedia system. The Web is distributed, interactive, and dynamic.

The technology consists of client/server software that can understand many information retrieval protocols in use on the Internet (FTP, TELNET, NNTP, WAIS, Gopher, Finger, Rlogin etc.), as well as the data formats of those protocols (ASCII,GIF, Postscript, WAV, MPEG, etc.), and can provide a single, consistent, user interface to them all. This makes the technology of interest not only for the Internet but also for intranets applications.

The WWW is based on a client/server model where the client and server work independently of the other, both in a technological sense and in an administrative sense. The Web is made up of the client/ server, the URL, and protocols that HTML-based documents. There are a number of reasons for setting up a Webserver, even for companies not interested in the Internet. A Web server can handle many of the same operations as Lotus Notes; both allow the user to share information with widespread access. The difference lies in where the information is kept. With Lotus Notes, an organization must copy documents to a central repository; however, when using a Web server anyone can access the information from anywhere within a company. The World Wide Web can also allow a company to do away with traditional distributed databases, which carry a costly overhead. The Web gives a company the ability to incorporate graphics and audio.

14.3 Web Servers

A server is the basic part that differentiates a provider from a user. To provide information on the web you must either have your own server or rent space on a server. The most common platforms used are Microsoft Windows, UNIX, VMS, and Macintosh.

Each platform has a number of programs available to set up a server using different versions of HTTP. Since the server and client run independently, the server can provide information to other clients and servers on different platforms. Each platform and software has its advantages and disadvantages. For a small provider, however, whichever system is currently being used is usually efficient.

NCSA HTTPD was the most popular and widely used server in a recent survey. The poll surveyed 1722 servers. Of those surveyed, 54 percent stated they use a form of the NCSA HTTPD. This survey corresponds to what newsgroups on the Internet had already

Been reporting for sometime. NCSA HTTPD has different versions for Windows and UNIX. NCSA is a widely supported server on the Internet; therefore it is easy to gain information and expert help when working with the NCSA server. There are a number of experts that deal with only NCSA HTTPD, which is not the case with CERN. Another advantage to using NCSA is that it works well with other applications running on the same computer because it is native to the UNIX system. Some applications that coincide with HTTPD are Gopher, WAIS, and list servers. The third advantage is that it also integrates well with Perl, a programming language used with most CGI. Since Perl was developed to work with UNIX-based systems, NCSA and CGI work well together.

Operating system	Program	FTP availability
UNIX	NCSA httpd	Ftp.ncsa.uiuc.edu:/web/ncsa_httpd(free)
	CERN httpd	http://info.cern.ch/hypertext/WWW/daemon/Status.html (available from other sources, also, free)
	GN Gopher/HTTP server	http://hopf.math.nwu.edu/
	Perl(plexus)	http://bsdi.com/server/doc/plexus.html
VMS	CERN HTTP	http://delonline.cern.ch/disk$user/duns/doc/vms/ distribution.html (free)
	Region 6 threaded HTTP server	http://kcgll.eng.ohio-state.edu/ww/doc/ serverinfo.html
Microsoft windows and windows NT sever	NCSA httpd	ftp://ftp.ncsa.uiuc.edu/web/ncsa_httpd/ contrib./whtp11ab6.zip(free)
	HTTPS	ftp://emwac.edu.ac.uk/pub/https
	SerWeb	ftp://winftp.cica.indiana.edu:/pub/pc/ win3/winsock/serweb03.zip
Macintosh	MacHTTP	http://ww.uth.tmc.edu/mac_info/ machttp_info.html(shareware)

Internet addresses for server program

NCSA HTTPD for Windows is easy to set up and use. Therefore it makes a good server for first-time users to learn the basics of HTTP. It is beneficial for someone who wants to learn by setting up a small system before developing one for a large organization. There are some disadvantages with this server. When running on Windows, it is slower than most servers. NCSA also requires lots of memory and CPU power. NCSA does not have the capability to run as an HTTP proxy client server. NCSA has proved to be a powerful server for UNIX platforms, but is somewhat weak for Windows users.

The second most popular server in the aforementioned survey was CERN HTTP. CERN was developed for use on VMS. Although CERN will run on UNIX, it is not recommended. Configuration and setup is more difficult than NCSA for Windows, but it is no harder to learn than NCSA for Unix. This platform has the ability to serve as a proxy server, which requires less software to access the web. It also has a scripting language that works well with CGI. It is accessible from the Internet for free, as are most servers.

MacHTTP is one of the few servers available for Macintosh computers. MacHTTP requires a minimum of System 7 and MacTCP to support advanced features. It is a small program that allows almost anyone to establish a server and effectively experiment with the HTTP protocol. This makes it a popular server for first-time and small providers. This server does not have a large handling capacity like NCSA, therefore it is not recommended for large corporations. However, some people have been able to get around this by networking a number of computers, each of which runs a subset of the entire HTML document list from a separate server.

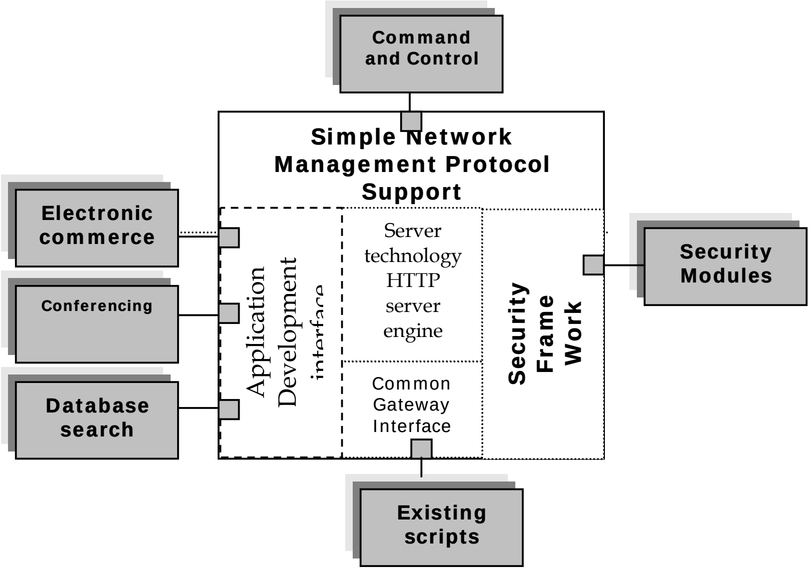

HTTP server engine

The data types and their quantities will affect the amount of disk space and maintenance required. If a company plans to distribute only text files, then the space requirement could be small. To utilize graphics or video, the data space will increase rapidly. How you decide to store the data types will affect the amount of CPU that is required. For example, if a company is providing access to a database, the CPU will use numerous resources to fulfill the requests. However, if you are simply providing lists of information via different links, the CPU requires many fewer resources.

The operating system that a company decides to use will influence the steps they need to take. Companies that have been using UNIX and are putting up a large server should have few problems. However, if a company is entering into Internet technology for the first time, it may be appropriate to hire a technical expert to determine the best equipment and to help install and maintain the server. On the other hand, a person who wishes to experiment on a hobby basis should not be discouraged from installing and running his or her own server.

The best way to set up the server us by starting from source code, so that the organization knows the system. However, you can also install a server with a precompiled binary software. Configuring a server can be done in a number of different ways, but there are some things that should be included in every server. For example, a server should be designed in a way that makes it easy to expand and up date as needs change. Therefore, it is a good idea to keep everything in one place in a tree structure so the information can be easily expanded or moved.

At a minimum, the configuration should specify the following information:

- ❖ The root (Server Root) of the directory tree.
- ❖ The user identify under which the server will run (for security reasons it is best to give the server as few privileges as possible)
- ❖ The path to the log files
- ❖ A mapping to executables on the server
- ❖ A mapping to HTML source files

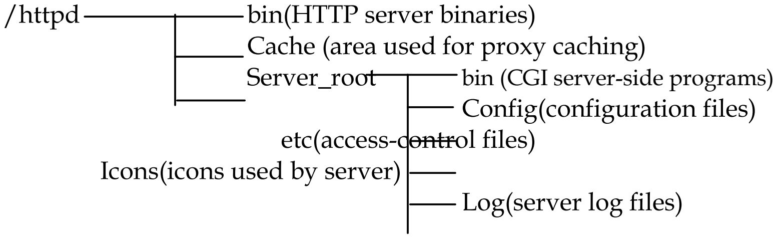
/httpd
bin(HTTP server binaries)
Cache (area used for proxy caching)
Server_root
bin (CGI server-side programs)
Config(configuration files)
etc(access-control files)
Icons(icons used by server)
Log(server log files)

Server directory tree

As mentioned earlier, there are a number of different protocols that the WWW can access. Each of these protocols has a special function. File transfer is the ability to transfer and locate information that is placed in a public domain. As the Internet grew, so did the needs of the users. Eventually this evolved into new services, such as electronic mail and bulletin boards. These services helped to encourage the growth of the Internet even more. Research then focused on ways to provide services on a higher level.

The basic advantage of using a Web server over another server is the language it supports and the ability to access many different places on the Internet from a single location. The Web server supports HTML, which allows for a variety of ways to provide information. With a Gopher server the user can access information only in textual format ad from lists. The Web server allows the user to insert links, graphics and sound, thereby making it more appealing and easier for people to use. Another advantage is the Web server's ability to allow links, which makes the server easier to maintain.

14.4 Transport Issues

HTTP is only a mechanism between sources and typically uses TCP/IP as a transport layer. HTTP allows the client to send a request for a hypertext document, then retrieves any items associated with the document. A user will often request a hypertext link to follow, which is often directed at the same server. TCP is the medium that establishes and transports this three-way handshake between the client and server.

TCP transfers the data in a series of PDUs. Every PDU of information does not have to be acknowledged before sending another packet. This form of control may seem to result in errors. There is a system that TCP uses, called slow start, to avoid/alleviate the errors. Slow start opens another window on the server side that keeps track of all received data. This window is opened every time a segment is received, and when a segment is not acknowledged the window is closed. When a window closes, TCP halts the data until the window receives the request it missed.

Analysis shows that HTTP spends more time waiting than receiving or sending data. HTTP is hindered by Slow Start because the URL is often longer than the maximum segment size. HTTP cannot respond until the entire URL has been received, which results in a delay between the initial request and the moment when information is sent. HTTP is also required to hold open resources for every connection that is established in a four-minute period. This level of resources is another major factor in the slowness of HTTP.

There are two main factors that affect the performance of a protocol the latency and bandwidth. Latency is a measure of the fixed overhead of a transition and does not change as documents get bigger or smaller. Bandwidth measures how long it takes to send data, this being the message length divided by the channel bandwidth. Bandwidth availability at the metropolitan and wide area level is being improved as time goes by. Faster modems and ISDN are two of the many ways that are increasing the speed of access; higher-speed facilities are also being introduced in the backbone. Unfortunately, reducing the amount of propagation delay is generally impossible. Therefore, latency due to propagation can become the dominating factor in some environments.

To combat the problems that have become associated with the early versions of http, a new improved version of HTTP has been proposed. Some has called the new version HTTP-NG.

One of the main differences between HTTP and HTTP-NG will be the basic model each uses when connecting. HTTP requires each request to open a new connection; each hyperlink/URL needs a separate TCP connection. HTTP –NG will allow multiple requests to be processed on one connection, thus increasing the speed of the transmission.

14.5 Web Access

Getting connected to the World Wide Web is relatively easy. There are actually a number of different ways to go about getting connected to the web. However, if you want to take advantage of everything on the web, then you need three things: a late-model computer, a minimum 14.4-Kbps modem, and access to a graphical web browser. You can get a free copy of the Netscape browser by contacting Netscape's home page.

An alternate way is to connect to a web browser called Lynx, which allows access to the web on a strictly textual basis. If you only want to see what is on the web without waiting for all the graphics, then Lynx is a possible way to become connected. All you need is an internet connection. If the internet service provider has Lynx, then you are ready to browse. However if the service does not offer Lynx, a user can get there from a number of sites that have Lynx available to the public. One place to try is www.law.cornell.edu. Because Lynx is a text-based browser, the user will miss out on the "real excitement of the web". Graphics range from full-color pictures to simple drawings scanned into home pages. A user will also miss out on the numerous audio clips available.

Gaining access to the graphics and audio clips requires better equipment, and this calls for more expensive equipment. This calls for more expensive equipment. This is not to say that it is difficult to get connected. In most cases, if users have a 486 IBM- compatible computer or a recent Performa or Power PC, they basically have everything they need. The only thing missing is a browser. Today there are literally thousands of Internet services that offer a complete package that allows to the web and other parts of the Internet. Many companies offer a free trial period.

All services have one thing in common: the way in which the connection is achieved. To be able to connect to the World Wide Web, a SLIP/PPP account must be account must be accessed. This account will allow the graphical browser software to access the hypertext documents on the Web. You can get a connection to a SLIP/PPP account without paying for on-line services. However, full access to the web in this case does require locating and installing a lot of software programs on your own.

The world wide web offers different things to different people and, depending on how much the user wants to do, determines the amount of equipment, time, and money it will take. Whether you wish to browse and enjoy the scenery or provide information to develop the growth of a company, the world wide web has something for everyone. For the most part it is easy to access and get your own site. Those who are not computer –knowledgeable can hire a company to put their information on the web. Or those who want to explore and try something new, there are programs designed for the beginner.

14.6 Connection Modes

There are several ways to connect to the Internet. You can dial up one of the commercial on-line services such as America Online, CompuServe, or Prodigy. Or, if you prefer, you can dial up an Internet service provider. If you are connected to a LAN, you need a network configuration program. On the other hand, if you calling an internet service provider to establish a SLIP or PPP connection, you need a dialer. In either case, you need additional tools to use the internet once you are connected.

To take full advantage of the internet, you need a number of pieces of software, a Winsock DLL, a Web browser, an e-mail reader, a newsgroup reader, an ftp client, and perhaps a gopher client and some other search tools.

14.7 Publishing Web Security Administration

In order to publish web pages on the Internet, it is desirable to have a dedicated computer that runs web service. In addition to the web server, a reasonably fast dedicated connection to the Internet is necessary. To be connected to the Internet, web servers can use TCP/IP over Ethernet. Web servers can also be installed on computers running multiuser operating systems such as Windows NT, Unix, or OS/2. In the past, web server technology was available only on a few platforms: Unix, Macintosh, VM, and VMS. Today, the platforms available also include, but are not limited to, OS/2, Windows NT, and windows. These systems can be set up to protect certain files from unauthorized access. Web servers can log activity such as the IP address, time, and request made for every connection. It is good practice to keep the logging on a computer separate from the Internet firewall computer. For some Internet web sites, the logging is kept on a machine separate from the web server and from the Internet firewall machine. Servers can forward requests for information that neither the client nor the server can access directly to applications called gateways, described in the next paragraph. Gateway support, logging, and user authentication are important features to look for when selecting a web server; logging is needed for both usage statistics and security. The person who administers the web server has come to be known as the Webmaster.

The web server waits for requests to come from browsers. When a request is made, the server locates the document and sends it back to the browser that originally requested it. Some requests may actually make the server run a script or program. These programs are called gateway scripts. The formal standard for these scripts is the CGI. A plain HTML document that the web server retrieves is static it is a text file that does not change. A CGI program is executed in real time, so it can output dynamic information.

14.8 Short Summary

A server is the basic part that differentiates a provider from a user. To provide information on the web you must either have your own server or rent space on a server.

There are several ways to connect to the Internet. You can dial up one of the commercial on-line services such as America Online, CompuServe, or Prodigy.

Latency is a measure of the fixed overhead of a transition and does not change as documents get bigger or smaller.

Bandwidth measures how long it takes to send data, this being the message length divided by the channel bandwidth.

14.9 Brain Storm

1. What is web server?
2. Explain various types of web server?
3. Explain about web access?
4. what is web ´s root?

Lecture 15

Internet Security

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Internet security.
- Describe the firewall concepts.
- Describe about securing Internet applications.
- Describe about virus checking and scanning.
- Describe the security administration.

Coverage Plan

Lecture 15

15.1 Snap Shot

A computer security policy cannot anticipate all scenarios of violations, but it can provide a proactive methodology to deal with most of them. After all, it is only a matter of setting up some access control lists and using the Internet security account manager to take care of them. All that is fine for getting it to work.

Network administrators worry about the security of their networks when they expose their organization's private data to Internet crackers. The level of protection is a function of more than access control lists and firewalls. An organization needs a security policy regarding unauthorized user access to resources on the private network. This chapter discusses the reasons you need security, and some of the firewall and network configurations that can help prevent outsiders from accessing your network.

15.2 Internet Security

When your internal network stops at the front door of your building, you do not have to worry about access from unauthorized individuals. You just have to watch out for the latest and greatest game that comes in via a disk in someone's briefcase – you know, the one he downloaded from the Internet last night that just happens to be infected with the latest virus. When you download some files at home and they wipe out your hard disk boot sectors, you are frustrated. You might even lose your checkbook accounting data for this year. So you get out all your program disks and start the laborious task of reloading everything. Oh well, it's just a few hours out of your life.

Computer access violations have become one of the primary concerns in today's interconnected environment. When an organization connects its private network to the Internet, its users have access to Internet services such as the World Wide Web, Internet mail, Telnet, and FTP (file transfer protocol). There is also the rage of having the company's home page on the World Wide Web, where it can also handle credit transactions from customers.

What to protect: Risk Analysis

One of the foremost reasons for creating a computer security policy is to ensure that efforts spent on your security yield cost-effective benefits. Although this might seem obvious, you would be surprised about where the effort is needed. For example, there is a great deal of publicity about crackers on computer systems; for most organizations, however, the actual loss from “insiders” is as great (or greater) a concern.

You need to do a risk analysis to decide what you need to protect, from whom you need to protect it, and how to protect it. Risk analysis is the process of looking at all of your risks and ranking those risks by the level of impact they may have on your operation if they were to occur. This process involves making cost-effective decisions about what you want to protect. Keep in mind that you do not want to spend more to lock up your information than it is worth. In doing your risk analysis, you will want to identify both the value of the assets and the risks.

Which Assets Need Protecting?

The first step in risk analysis is to identify all the things that need to be protected. Some things are obvious, like all the various pieces of hardware, but the not so obvious are the people who are normal users. The essential point is to list all things that could be affected by a security problem. The following sections detail some categories to consider.

Systems

Systems represent the physical component of your risk assessment. This category consists of

Hardware: CPUs, boards, keyboards, terminals, workstations, personal computers, printers, disk drives, communication lines, terminal servers, routers.

Software: Source programs, object programs, utilities, diagnostic programs, operating systems, communication programs.

Information

While the systems or physical component is critical, the content is the driving force for the risk assessment. The content risk consists of

Data: during execution, stored online, archived offline, backups, audit logs, databases, in transit over communication media.

Documentation: on programs, hardware, systems, local administrative.

People

Systems and content are amazingly secure when the people component is removed from the picture. Since this is usually not possible, the people component of the risk assessment consists of

Users : people needed to run systems.
Crackers : people who have no business in your business.

What Are your Threats?

Once the assets are identified, you need to identify threats to those assets and then look at the threats to determine what potential for loss or damage exists in the event of security breach. Think about which threats can affect which of your assets. The following sections explain some threats you need to consider.

Unauthorized Access

Denying unauthorized access to computers requires unrelenting efforts by systems administrators. Unauthorized access comes in many flavors. The most simple form is the use of another user's account to gain access to a system. Legally speaking, any computer resource that is accessed without prior permission can be considered unauthorized access to computing facilities. The seriousness of an authorized access can vary. So Sally give Joe her username and password so he can get some data while she is off tomorrow. It seems innocent enough, but for some sites, the simple act of granting access to an unauthorized user could be a major offense. An unauthorized access can point the way to the other security threats. In addition, some sites might be more frequent targets than others, making the risk from unauthorized access different from site to site. You will have to decide the risk from unauthorized access different from site to site. You will have to decide what is right for your site and incorporate those guidelines in your security polices.

Information Disclosure

A common threat is disclosure of information. How valuable or sensitive is the data stored on your computers? Disclosure of a password file might allow for future unauthorized accesses, whether it is from inside or outside of the company. Even a simple glimpse of a proposal could give your competitor an advantage that could severely impact your future as a company. Information is power. Management of information is time, and time is money.

Denied Permissions

Computers and networks have a way of building dependence on their being accessible when you need them. People quickly come to reply on these services in order to perform their jobs efficiently. Whenever these services are not available when called on, a loss in productivity results. Manual processes that were used months ago somehow cannot be Reconstructed when the new, technological approach will not work. Denial of permission comes in many forms and can affect users in a number of ways. A network can be shut down by a cracker's intrusion, jamming the flow of business traffic. A virus might slow down or cripple a computer system. You need to determine what your site has in terms of critical services; then you need to determine the effect on your site if one or more of those services were to become disabled.

15.3 Firewalls

You have seen those thick brick or masonry walls that separate sections of a building. They are called firewalls, and they are there to keep a fire on one side from spreading to the other side. The same theory applies to networks and the computing equipment that makes up a network firewall.

An Internet firewall is a system or group of systems that enforces a security policy. This policy is one that you establish between your company's network and the Internet. The firewall determines which inside services may be accessed from the outside, which outsiders are given access to the permitted inside services, and which outside services may be accessed by insiders.

An effective firewall requires that all traffic to and from the Internet must pass through the firewall, Every packet of data must pass through the firewall and pass inspection before it can go on to its destination. You might think that this is a potential bottleneck. You are correct; it will be if the firewall is not robust enough to handle all the traffic that the network is transporting. You will know when your firewall is reaching its limit when your World Wide Web users begin to get access denials due to firewall inadequacies.

A Perimeter of Defense

The area being protected lies within a perimeter of defense.

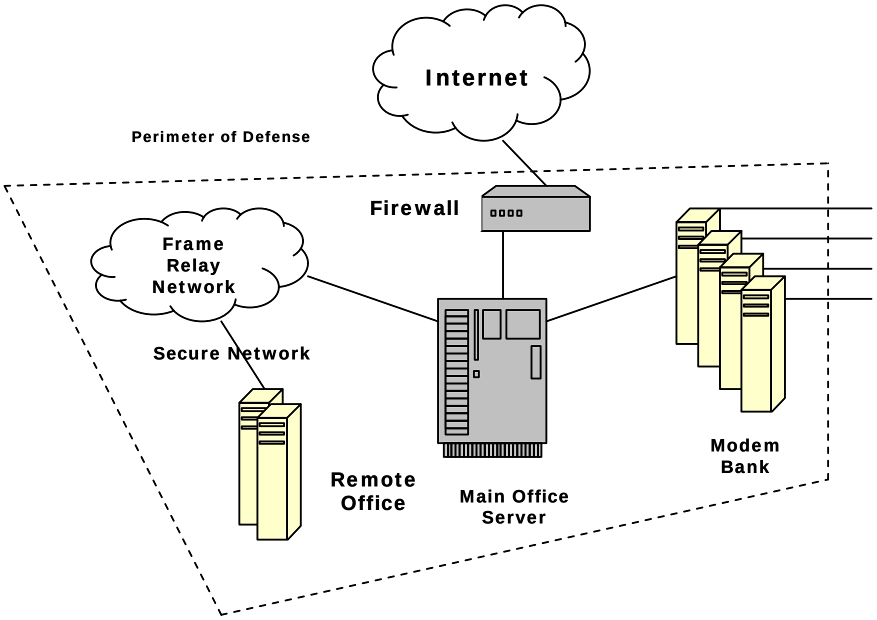

At the perimeter of defense, all traffic is stopped and checked for integrity. A firewall must allow only authorized packets to pass, and the firewall itself must be self protecting. Like the fire that has engulfed the building on both sides of a firewall, a network firewall is worthless once a cracker has gotten through or around it.

An Internet firewall made up of just a router, a bastion host, or a combination of devices that provides security for a network will not provide protection against the actions of your own company personnel. The firewall is just a part of the security that you establish. The purpose of your perimeter defense design is to secure the data resources of your company, but this requires more than network-monitoring equipment.

A major component of your security is a commitment from your company's management on security policy. Your security policy must include published security guidelines to inform users, because all users have responsibilities when it comes to security. Company policies must prescribe network and service access, local and remote user authentication, dial-in and dial-out, disk and data encryption, and virus-protection measures. None of this will work without employee training in security-related issues.

When to Protect

Internet firewalls are access managers. They are designed to manage user access between the Internet and your private network. Without a firewall, every computer on your private network is exposed to attacks from other computers on the Internet. In this case, your username and password will be your only protection against a cracker.

Internet firewalls are a mechanism that gives the network administrator a place to define a security policy. A firewall policy is the administrators definition of that type of data packets will be allowed to pass through the firewall,. Firewall policies also provide protection from various types of routing attacks. Routing attacks are attempts to revise the routing algorithms in a packet router constraining the passing of packets to and from your users. An internet firewall simplifies security management instead of making it more complex. With your network security being located at one place, you will be able it manage it more efficiently. It is much easier to oversee a firewall system than to run around to tens or hundreds of user machines to manage individual security structures.

The firewall offers a single point at which Internet security violations can be monitored. It can provide a report of violations in one location rather than your having to scan your user community to see what is happening on your network. Your network administrator should audit the logs of traffic through the firewall. You can set up limit criteria to filter out the normal masses of logging so your violations will not be lost in the midst of traffic.

An internet firewall can provide you with the necessary data to justify a need for higher-speed internet connections. You can also use the data to allocate the cost of your internet to the different user groups in your company. In this manner, the cost of the connection is shared, and the heavy user pays the largest share of the cost. In any case, you need to know what is happening with your internet connection. The statistical information produced by the firewall can help you plan your network upgrades.

Keeping the Back Door Closed

Power users are rebellious about restrictions imposed by a firewall proxy server (otherwise known as a bastion host). They feel they are smart enough to not download some corrupting package. They are often irritated at the network administrator who has control over their computers against their will. A method of covert retaliation is for a user to circumvent the security system by using his private, direct SLIP or PPP connection to an ISP. Because these types of connections bypass the security provided by the most carefully constructed firewall, they create a significant potential for back-door attacks.

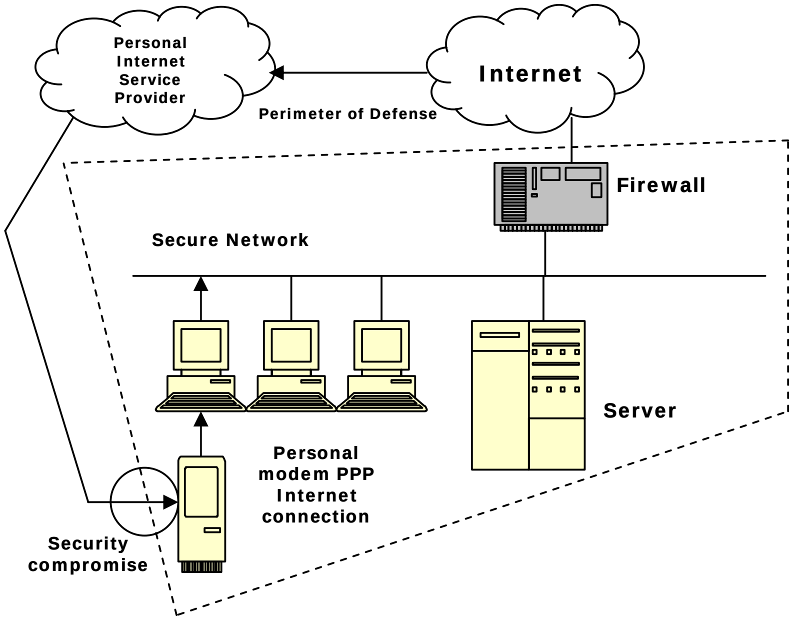

15.4 Securing Internet Applications

The following some time-tested tips for setting up Internet-related applications in a bastion host circuit-level gateway application.

Domain Name Service

These are some considerations for implementing domain name services in a bastion host environment:

- Set up an external DNS server on a bastion host for the outside world to access
- Configure your DNS to hide all host information.
- Use double reverse lookups to avoid spoofing.

- Consider hiding all internal DNS data, forwarding, and fake records.
- Don’t allow domain name service zone transfers.

Gopher

These are some considerations for implementing Gopher services in a bastion host environment:

- Use a dedicated bastion host
- Control the external programs your Gopher server can access
- Use proxies on Gopher wherever possible.
- Watch out for users reconfiguring Gopher clients.
- Use Microsoft Internet Explorer as your Gopher client.

Wide Area Information Server

These are some considerations for implementing a WAIS in a bastion host environment:

- Run WAIS on your bastion server on the standard port 210.
- Restrict your WAIS client to port 210 or use proxies
- Use Microsoft Internet Explorer as your WAIZ client.

Hypertext Transfer Protocol

These are some considerations for HTTP services in a bastion host environment:

- Use a dedicated bastion host for your HTTP server
- Control the external programs your HTTP server can access
- Use the standard HTTP port 80 or use proxies
- Use a caching proxy server.
- Watch out for users reconfiguring HTTP clients.
- Use Microsoft Internet Explorer as your HTTP client.

Network News Transport Protocol

These are some considerations for implementing NNTP in a bastion host environment:

- Do not use a bastion host as a news server
- Do allow automated newsgroup creation
- Allow external NNTP with just the sites you exchange news with.
- Use packet filtering or proxies to connect with trusted external NNTP servers.

Telnet

These are some considerations for implementing Telnet in a bastion host environment:

- Restrict incoming Telnet
- Use packet filtering and proxies for outgoing Telnet.
- Find an encrypted Telnet package.

File Transfer Protocol

These are some considerations for implementing FTP in a bastion host environment:

- Allow outgoing FTP only from and to port 1023.
- Use an FTP proxy server.
- Allow incoming FTP to come only to your bastion host.
- Be sure your FTP server is up to date.
- Prevent third-party transfers of files in the incoming area.
- Educate your users.
- Enable user authentication for non-anonymous FTP.

Simple Mail Transport Protocol

These are some considerations for implementing SMTP in a bastion host environment:

- Use normal store and forward techniques.
- Use packet filtering and restrict outgoing packets to the bastion host.
- Use packet filtering to restrict SMTP connections from the bastion host to specific internal mail hosts.
- Keep your mail server software current.
- Educate your users.

15.5 Virus Checking and Scanning

A virus is a piece of software that is executed when its host program is run. That means the virus is just sitting there waiting for you to crank up something innocent. When the host program is run, the virus is executed and does its damage.

A virus can be either destructive or just slightly annoying. Both varieties can multiply through your network faster than you can stop them. Viruses can be destructive as soon as they enter a system, or they can lie quietly until activated by a trigger such as a date and time.

Some "Viruses" Are Not Viruses

Not every malicious piece of code is a virus. Some kinds of non-virus but malicious code are worms, Trojan horses, and logic bombs.

Worms are designed to infiltrate common business applications and infect or destroy the data. Unlike what is termed as a virus, the worm does not normally replicate itself and spread to other systems. An example of a worm might be code that sends funds from your checking account to another bank.

Trojan horses are destructive program that are concealed in a standard piece of software, such as a spreadsheet program. Trojan horses also do not spread like a virus. Trojan horses are very likely to be found in game software that can be downloaded from the Internet. They are also found embedded in graphics files that are available in many newsgroups.

Logic bombs are similar to Trojan horses in destruction capability. They are classed as logic bombs because they use timing routines to go off at a particular date and time. The Michelangelo virus is embedded inside a logic bomb. The logic bomb is the favorite of the disgruntled former employee. Sharp employees can leave behind logic bombs that are triggered when they are deleted from systems where the employee had permissions for access.

Where Viruses Do Their Work

Different types of viruses attack different parts of a computer's environment, such as

- Boot-sector viruses
- File-infecting viruses
- Polymorphic viruses
- Stealth viruses
- Multi-partite viruses

The following sections review these types of viruses.

Boot-Sector Viruses

You will know you have a boot-sector virus when you get a message that says there is no valid drive or no operating system. Usually the boot-sector virus moves the master boot record to another location, thus rendering the disk useless. Because the boot record is the first thing that has to load, there is no way to detect the virus before it has done its job. Boot-sector viruses are loaded by booting the machine from a floppy disk that has been infected.

File-Infecting Viruses

Naturally, file-infecting viruses infect files. The most common operation of this type of virus is to infect all files that have the same extension, such as all .exe files. File viruses normally move a program's load code and replace it with their own. The resulting file size is usually changed, which helps determine that the file is infected.

Polymorphic Viruses

The polymorphic virus changes its appearance as it replicates itself throughout a network. This type of virus is difficult to detect because of the number of different binary patterns it can assume.

Stealth Viruses

As you might have guessed, the stealth virus attempts to hide from detecting software. This type of virus usually resides in memory and intercepts all attempts to use system calls to the operating system software. The memory must be examined and cleaned before any disk infection can be corrected or maybe even detected.

Multi-partite Viruses

Multi-partite viruses infect both the boot sectors and the executable files. They might even include the capability to replicate themselves with polymorphism techniques.

Viruses and Your System

The best virus protection is user education. It's that simple. Most viruses are still spread with floppy disks instead of network connections or dial-up services.

Viruses are designed to be spread. It does not bring much joy to a cracker to write a virus that is going to infect only once, where it will be discovered and destroyed. The cracker will want to get some mileage out of the malicious code he labors over. You should be especially conscious of opportunities that lead themselves to spreading viruses. The three that are most prevalent are field technicians, salespeople, and outside service people such as your friendly accountant.

Some rogue code is designed to be carried into a secure environment but not do any apparent harm. Its only purpose is to reconfigure your environment to allow entry from the outside to go undetected. This kind of code is known as a trapdoor.

A trapdoor is a way of getting to your data or network without entering usernames or passwords. Trapdoors are sometimes inadvertently left if newly developed code. A programmer might stick a secret routine in a product that lets him enter without the hassle of a username or password.

Virus-Proofing Your Network

There is no way to completely protect a computing environment from the attack of a malicious piece of code unless it is a dedicated process with no external access. This would have to be a small machine that has its programs stored in read-only memory with no access to external data transfer. Computer environments with random access memory and any type of magnetic storage media can be victims.

Your methods of protection will vary depending on the number of machines you have, the amount of information that is passed between them, how accessible they are to the outside world, and the pace of your operations. If your environment supports real-time memory with no access to external data transfers, you should be very concerned with virus detection. Note that is detection of viruses, not immunity. It is going to happen – there is no doubt about that. What you want to do is catch it before it has had an opportunity to spread or do a substantial amount of damage.

Every machine on your network should be equipped with virus-detection and correction software. The virus-detection software needs to be run frequently. If your operation has real-time, around-the-clock transactions, a virus scan once every hour might not be too frequent. You will have to decide what is right for your site. The key is to use the scheduler that almost every virus scan program includes so that you will not forget to scan for viruses. The program should be enabled to test for viruses in the boot sector and all executable files. Files with .zip as their extension are excellent candidates for hidden problem software. You should run full scans on all files, even though they take longer. You might want to scan for a subset more frequently and perform a full scan every few days. It will depend on your level of comfort and on your site's needs.

15.6 Security Administration

An integral part of developing an official site policy on computer security is to define your expectations of proper computer and network use. You should also define procedures to prevent and respond to breaches of security. You will need to take into account several aspects of your company's processes and organization.

Consider your company's business goals. Manufacturing pots and pans should have a security emphasis different from designing and building space shuttles.

The site security policy you develop must conform to existing policies, rules, regulations, and laws. It will be necessary to identify what impact these might have on your security policy.

Your policy should address issues of local security problems as well as problems occurring on remote systems as a result of a local host or user. Setting security policies and procedures means developing a plan for how to deal with computer security. Your plan should address

- What you are trying to protect
- From whom you need to protect it

- ❑ How likely the threats are
- ❑ How to protect your assets in a cost-effective manner

You must also review the process continuously. The environment changes often in the rapidly evolving world of today's technology. Your plan will need to be kept current if it is to be successful.

Policies

Policy creation must be a joint effort by technical and managerial personnel: technical people who understand the full ramifications of the proposed policy and its implementation, and managers who have the power to enforce the policy. A policy that is technically incorrect or has no teeth is useless.

You will need to establish your policy at a level in the organization that will be able to enforce the policy across department lines. You, as an administrator of networks, might have to cross a department line for technical purposes. The management chain on the other side of the line will need to be supportive of your efforts.

There are a number of issues that must be addressed when developing a security policy:

- ❖ Who is allowed to use the resources?
- ❖ What is the proper use of the resources?
- ❖ Who is authorized to grant access and approve usage?
- ❖ Who may have systems administration privileges?
- ❖ What are the users' rights and responsibilities?
- ❖ What are the rights and responsibilities of the systems administrator versus those of the user?
- ❖ What do you do with sensitive information?

Who is Allowed to Use the Resources?

You will need to define who is allowed to use your system and services. The policy should state who is authorized to use what resources.

What Is the Proper Use of the Resources?

The policy should state what is acceptable use as well as unacceptable use. It should also include types of use that may be restricted, such as downloading pornography from Internet sources.

Your acceptable-use policy needs to define the responsibility of individual users. They must know that their responsibility exists regardless of your security efforts. Following are some suggested topics:

- Breaking into accounts should not be allowed.
- Cracking passwords should not be allowed, even if it is for business purposes.
- Disrupting service is costly and should be reprimanded.
- Users should have a need-to-know basis to obtain information.
- Users should not modify files belonging to others.
- Sharing accounts is a no-no.

You should include a statement in your policies concerning copyrighted and licensed software. Licensing agreements with vendors require your assurance that the license is not violated. Users must understand that the copying of copyrighted software may be a violation of the copyright laws.

Who Will Administer the Systems?

One security decision that needs to be made up front is who will have access to systems administrator privileges and passwords on your systems. Obviously, the systems administrators will need access, but you will soon have power users who want special privileges. Your policy must balance restricting access with giving access to people who need these privileges. You should grant only enough privilege to accomplish the necessary tasks regardless of the perceived power of the user. Users holding special privileges should be accountable to some authority, and this should also be identified within the site's security policy.

Define the Users' Rights and Responsibilities

Your policy must be clear about the users' rights and responsibilities. It should be clearly stated that users are responsible for understanding and respecting the security rules of the systems. The following is a list of suggested topics:

- Guidelines for resource consumption
- What “abuse” means in terms of system performance
- A statement on sharing accounts and files
- The degree to which users need to keep their passwords secret
- Established password-expiration and minimum-length standards
- Backing up data files
- Disclosure of proprietary information
- Postings to mailing lists or discussion groups
- Policy on electronic mail communications.

The user has a right to privacy, but systems administrators need to gather sufficient information to diagnose problems. The policy must specify to what degree systems administrators can examine user files. A policy concerning the systems administrators’ obligation to maintaining the privacy of information viewed under these circumstances is a worthwhile effort.

15.7 Short Summary

Computer hardware is expensive, but hardware cost usually is significant compared to the development and ongoing support of the software. Software and the data created by the software are the real values in a computing environment. A capable computer technician can restore the hardware with a new operating system and hand it back to you clean as a whistle. You will then be faced with restoring the application programs and as much of the data as possible from the most recent backups.

Building a perimeter of defense is the first step. Then you will need to be constantly on the defensive against the ever-present danger of a system infection. Catching it before it has done much damage is the key. Frequently scan for viruses with software designed just for that purpose.

15.8 Brain Storm

1. Why we need security explain.
2. Explain about firewalls.

3. What are the applications we used for securing Internet explain?
4. What is virus?
5. How to scan for viruses explain.
6. Explain about security administration.

Lecture 16

Firewall

Objectives

After completing this lesson, you should be able to do the following

- Discuss about firewalls.
- Describe the firewall technologies.
- Describe about packet filtering.
- Describe about firewall architecture.
- Describe the firewall design.

Coverage Plan

Lecture 16

16.1 Snap Shot

In general, a proxy is something or someone who does something on somebody else's behalf. For instance you may give somebody the ability to vote for you by proxy in an election.

Proxy services are specialized application or server programs that take users requests for Internet services (such as FTP and Telnet) and forward them to the actual services. The proxies provide replacement connections and act as gateways to the services .For this reason proxies are sometimes known as application-level gateways. In this book when we are talking about proxy services we are specifically talking about proxies run for security purposes which are run on a firewall host either a dual –homed host with an interface on the internal network and one on the external network ,or some other bastion host that has access to the Internet and is accessible from the internal machines.

16.2 Firewall

First you need to determine what the firewall needs to do in detail. Yes you are trying to make your site secure, but how secure does it need to be?

Your first starting will be you security policy. If you don't have a security policy. Security policies for some suggestions on how to go about setting one up. You cant just do without a policy because a firewall is an enforcement device if you didn't have a policy before, you do one you have a firewall in place, and it may not be a policy that meets your needs.

Firewall

A components that restricts access between a protected network and the Internet or between other sets of networks.

16.3 Firewall Technologies

You may be familiar with some of the following firewall terms and some may be new to you. Some may seem familiar but they may be used in a way that is slightly different from what you're accustomed to (though we try to use terms that are as standard as possible). Unfortunately there is t completely consistent terminology for firewall architecture and components .Different people use terms in different—or worse still, conflicting ways. Also, there same terms come times have other meanings in other networking fields ; the following definitions are for a firewalls context.

Here are come very basic definitions; we describe these terms in greater detail elsewhere:

Host

A computer system attached to a network.

Bastion host

A computer system that must be highly secured because it is vulnerable to attach usually because it is exposed to the Internet and is a main point of contact for users of internal networks. It gets its mane from the highly fortified projections on the outer walls of medieval castles.

Dual-homed host

A general- purpose computer system that has at least two network interfaces.

Network address translation (NAT)

A procedure by which a router changes data in packets to modify the network addresses. This allows a router to conceal the addresses of network hosts on one side of it. This technique cab equable a large number of hosts to connect to the Internet using a small number of allocated addresses to connect to the Internet using valid addresses. It is not actually a security technique although it can provide a small amount of additional security. However , it generally runs on the same routers that make up part of the firewall.

Packet

The fundamental unit of communication on the Internet.

Perimeter network

A network added between a protected network and an external network , in order to provide an additional layer of security .A perimeter network is some times called a DMZ which stands for De-Militarized Zone (named after the zone separating North and South Korea)

Proxy

A program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy clients talk to proxy servers which relay approved client requests on to real servers and relay answers back to clients.

Virtual private network(VPN)

A network where packets that are internal to a private network pass across a public network, without this being obvious to hosts on the private network. In genitals , VPNs use encryption to protect the packets as they pass across the public network. VPN solutions are popular because it is often cheaper to connect two networks via public networks (for instance, getting them both Internet connections) than via private networks(like traditional leased-lime connections between the sites)

The next few sections briefly describe the major technologies associated with firewalls: packet filtering , proxy services. Networks address translation. And virtual private networks.

There are legitimate questions about now to distinguish between packet filtering and proxying , particularly when dealing with complex packet filtering systems and simple proxies . Many people believe that systems that pay attention to individual protocols and /or modify packets should not be considered packet filters and many even refer to these systems as transparent proxies. In fact, these systems don't behave much like older, simpler packet filtering systems, and it's a good idea not to apply generalization. About packet filtering to them blindly. On the other hand they don't behave much like proxying system, either.

Similarly, a number of proxying systems provide generic proxies which essentially function like packet filters, accepting all traffic to a given port without analyzing it. It's advisable to pay close attention to the individual technology a products users, without making assumptions based on whether it claims to be a packet filter or a proxy. However many systems still are clearly packet filters or clearly proxies, so it is worth understanding what these technologies are and how they work.

16.4 Packet Filtering

The action a device takes to selectively control the flow of data to and from a network . Packet filters allow or block packets, usually while routing them from one network to another (most often from the Internet to an internal network, and vice versa). To accomplish packet filtering you set up a set of rules that specify what types of packets(e,g. those to or from a particular IP address or port) are to be allowed and what types are to one blocked . Packet filtering may occur in a router in a bridge , or on an individual host. It is sometimes known as screening.

Packet filtering systems route packets between internal and external hosts, but they do it selectively . They allow or block certain types of packet in a way that reflects a sites own security policy, as shown in the below figure. The type of router used in a packet filtering is known as a screening router.

Packet filtering, every packet has a set of headers containing certain information. The main information is:

- IP source address
- IP destination address
- Protocol (whether the packet is TCP, UDP, or ICMP packet)
- TCP or UDP source port
- TCP or UDP destination port
- ICMP message type
- Packet size

The router can also look past the packet headers at data further on in the packet; this allows it, for instance to filter packets based on more detailed information (like the name of the web page that a somebody is requesting) and to verify that packets appear to be formatted as expected for their destination port. The router can also make sure that the packet is valid (it actually is the size that it claims to be and is a legal size for instance) which helps catch a number of denial of service attacks based on malformed packets.

In addition the router knows things about the packet that aren’t reflected in the packet itself, such as :

- The interface the packet arrives on
- The interface the packet will go out on

Finally, a router that deeps trick of packets it has seen knows some useful historical facts such as:

- Whether this packet appears to be a response to another packet(that is its source was the destination o a recent packet and its destination is the source of that other packet)
- How many other packets have recently been seen to or from the same host
- Whether this packet is identical to a recently seen packet
- If this packet is part of a larger packet that has been broken into parts (fragmented)

To understand how packet filtering works, let's look at the different between an ordinary router and a screening router.

An ordinary router simply looks at the destination address of each packet and picks the best way it knows to send that packet towards that destination, The decision about how to handle the packet is based solely on its destination . There are two possibilities : the router knows how to send the packet towards its destination , and it does so; or the router does not know how to send to packet towards its destination, and it forget about the packet and returns an ICMP "destination unreachable" message to the packet's source.

A screening router, on the other hand looks at packets more closely. In addition to determining whether or not it can route a packet towards its destination a screening router also determines whether or not it should. "Should" or "Should not" are determined by the site's security policy. Which the screening router has been configured to enforce.

Packet filtering may also be performed by devices that pay attention only to "should " and "Should not" and have to ability to router. Devices that do this at packet filtering bridges. They are rarer than packet filtering routers mostly because they are dedicated security devices that don't provide all the other functions that routers do. Most sites would rather add features to routers that they need anyway, instead of adding a dedicated device. However being a dedicated device provides advantages for packet filtering bridges ; in particular, they are harder to detect and attack than packet filtering routers. They provide the same general features that we discuss that we discuss for packet filtering routers.

Once it has looked at all the information , a straightforward packet filtering router can do any of the following things:

- Send the packet on to the destination it was bound for
- Drop the packet- just forget it, without notifying the sender
- Reject the packet – refuse to forward it, and return an error to the sender.

- Log information about the packet
- Set off an alarm to notify somebody about the packet immediately

More sophisticated routers might also be able to do one or more of these things:

- Modify the packet (for instance to do network address translation)
- Send the packet on to a destination other than the one thaw it was bound for (for instance, to force transaction through a proxy server or perform load a balancing)
- Modify the filtering rules (for instance to accept replies to a UDP packet of to deny all traffic from a site that has sent hostile packets)

The fact that servers for particular Internet services reside at certain port numbers lets the router block or allow certain types of connections simply by specifying the appropriate port number in the set of rules specified for packet filtering.

Here are some examples of ways in which you might program a screening router to selectively route packets to or from your site:

- Block all incoming connections from systems outside the internal network , except for incoming SMTP connections(so that you can receive electronic mail)
- Block all connections to or from certain systems you distrust
- Allow electronic mail and FTP services but block dangerous services like TFTP, the window System ,RPC and the "r" services (rlogin, rsh, scp, etc)

Proxy Services

You will also run into proxies that are primarily designed for network efficiency instead of for security these are caching proxies which keep copies of the information for each request that they proxy. The advantage of a caching proxies is that if multiple internal hosts request the same data, the data can be provided directly by the proxy .Caching proxies can significantly reduce the load on network connections. There are proxy servers that provide both security and caching; in general they are better at one purpose that the other.

Proxy services sit more or less transparently between a user on the inside (on the internal network) and a service on the outside (on the Internet) .Instead of talking to each other directly , each talks to a proxy. Proxies handle all the communication between users and Internet services behind the scenes.

Transparency is the major benefit of proxy services . It's essentially smoke and mirrors . To the user a proxy server presents the illusion that the user is dealing directly with the real server. To the real server , the proxy server presents the illusion that the real server is dealing directly with a user on the proxy host(as opposed to the user's real host)

How do proxy services work ? Let's look at the simplest case, where we add proxy services to a dual-homed host.

A proxy service requires two components : a proxy server and a proxy client. In this illustration the proxy server runs on the dual-homed host as proxy system there are other ways to set up a proxy server) A proxy client is a special version of a normal client programs can be used as proxy clients. The proxy server evaluates requests from the proxy client and decides which to approve and which to deny. If a request is approved the proxy server contacts the real server on behalf of the client and proceeds to relay requests from the proxy client to the real server and responses from the real server to the proxy client.

In some proxy systems instead of installing custom client proxy software, you'll use standard software but set up custom user procedures for using it.

There are also systems that provide a hybrid between packet filtering and proxying where a network device intercepts the connection and act as a proxy of redirects the connection to a proxy; this allows proxying without making changes to the clients or the user procedures.

The proxy server doesn't always just forward user requests onto the real Internet services. The proxy server can control what users do because it can make decisions about the requests it processes. Depending on your site's security policy, requests might be allowed or refused . For example. The FTP proxy might refuse to let users export files or it might allow users to import files only from certain sites. More sophisticated proxy services might allow different capabilities to different hosts, rather than enforcing the same restrictions on all hosts.

Some proxy servers do in fact just forward requests on, no matter what they are . These may be called generic proxies or port forwarders. Programs that do this are providing basically the same protections that you would get if you has a packet filter in place that was allowing

traffic on that port . You do not get and significant increase in security by replacing packet filters with provides that do exactly the same thing (you gain some protection against malformed packets but you lose by adding an stackable proxying program)
Some excellent software is available for proxying . SOCKS is a proxy construction toolkit , designed to make it easy to convert existing client/server application into proxy versions of those same applications. The trusted Information Systems Internet firewall Toolkit (TISFWTK) includes proxy servers for a number of common Internet protocols, including Telnet, FTP, HTTP, rlogin, X11 and others these proxy servers are designed to be used in conjunction with custom user procedures. See the discussion of these packages in proxy systems.

Many standard client and server programs, both commercial and freely available now come equipped with their own proxying capabilities or with support for generic proxy systems like SOCKS. These capabilities can be enabled at runtime or compile time .

Most proxy systems are used to control and optimize outbound connections they are controlled by the site where the clients are . It is also possible to use proxy systems to control and optimize inbound connections to servers(for instance to balance connections among multiple servers or to apply extra security).s This is sometimes called reverse proxying.

16.5 Firewall Architecture

Single –box architectures

In simplest firewall architectures have a single object that acts as the firewall. In general, the security advantage of single-box architectures is that they provide a single place that you can concentrate on and be sure that you have correctly configured, while the disadvantage is that your security is entirely dependent on a single place. There is no defense in depth, but on the other hand, you know exactly what your weakest link is and how weak it is, which is much harder with multiple layers.

Screening router

It is possible to use a packet, filtering system by itself as a firewall, using just a screening router to protect an entire network. This is a low cost system since you almost need a router to connect to the Internet anyway and you can simply configure packet filtering in that router. On the other hand, its not very flexible you can permit or deny protocols by port number, but it is hand to allow some operations while denying other in the same protocol, or to be sure that what coming in on a given port is actually the protocol you wanted to allow. In addition it gives you no depth of defense. If the router is compromised you have no further security.

Appropriate uses

A screening router is an appropriate firewall for a situation where

- The network being protected already has a high level of host security.
- The number of protocols being used is limited, and the protocols themselves are straightforward.
- You require maximum performance and redundancy.

Screening router is most useful for internal firewalls and for networks that are dedicated to providing services to the Internet. It's not uncommon for Internet service providers to use nothing but a screening router between their service hosts and the Internet, for instance.

Dual homed host

Dual homed host architecture is built around the dual homed host computer a computer that has at least two network interfaces. Such a host could act, as a router between the networks these interfaces are attached to it is capable of routing IP packets from one network to another. However, to use a dual homed host as a firewall, you disable this routing function. Thus, IP a packets form one network are not directly routed to the other network. Systems inside the fir3wall can communicate with the dual homed host, and systems outside the firewall can communicate with the dual homes host, but these systems can't communicate directly with each other. IP traffic between them is completely blocked.

Some variations on the dual homed host architecture use IP to the Internet and some other network protocol (for instance, netBeui) on the internal network. The helps to enforce the separation between the two networks, making it less likely that host misconfigurations will let traffic slip from one interface to another, and also reducing the chance that if this does happen there will be vulnerable clients. However, it does not make a significant difference to the overall security of the firewall.

The network architecture for a dual homed host firewall is pretty simple the dual homed host sits between and is connected to the Internet and the internal network.

Dual homed hosts can provide a very high level of control. If you are not allowing packet to go between external and internal networks at all, you can be sure that any packet on the internal network that has an external source is evidence of some kind of security problem.

On the other hand dual homed hosts are not high performance devoices. A dual home host has more work to do for each connection than a packet filter does and correspondingly needs more resources. A dual homed host would not support as much traffic as an equivalent packet filtering system.

Since a dual homed host is a single point of failure, it is important to make certain that its host security is absolutely impeccable. An attacker who can compromise the dual homed host has full access to your site (no matter what protocols you are running). An attacker who crashed the dual homed host has cut you off from the internet. This makes dual homed hosts inappropriate if being able to reach the internet is critical to your business.

You are particularly vulnerable to problems with the hosts IP implementation which can crash the machine or pass traffic though it. These problems exists with packet filtering routers as well, but they are less frequent and usually easier to fix. Architectures that involve multiple devices are usually more resilient because multiple different IP implementations are involved

Proxying is much better at supporting outbound services than inbound services. In a dual homed host configuration you will normally have to provide services to the internet by running them on the dual homed host. This is not usually advisable because providing services to the Internet is risky, and the dual homed hosts is
security critical machine that you don't want to put risky services on. It might be acceptable to put a minimally functional web server on the dual homed host (for instance, one that was only capable of providing HTML files and has no active content features additional

protocols, or forms processing) but it would clearly be extremely dangerous to provide a normal we server there.

The screened subnet architecture we describe in a later section offers some extra options for providing new, untru8sted or inbound services.

Appropriate uses

A dual homed host is an appropriate firewall for a situation where:

- Traffic to the Internet is small
- Traffic to the internet is not business critical
- No services are being provided to Internet based users.
- The network being protected does not contain extremely valuable data.

Multiple purpose boxes

Many single box firewalls actually provide some combination of proxying and packet filtering. This gives you many of the advantages of both you can allow some protocols at high speed while still having detailed control. It also gives you many of the disadvantages of both you are vulnerable to problems where protocols that you though were forced through the proxies are simply passed on by the packet filters. In additions, you have all a the normal risks of having only a single entity between you and the great outside world.

Appropriate uses

A single machine that does both proxying and packet filtering is appropriate for a situation where:

- The network to be protected is small.
- No services are being provided to the Internet.

Screened host architectures

Where a dual homed host architecture provides services from a host that's attached to multiple networks a screened host architecture provides services from a host that attached to only the internal network, using a separate router. In this architecture, the primary security is provided by packet filtering.

A simple version of a screened host architecture. The bastion host sits on the internal network. The packet filtering on the screening router is set up in such a way that the bastion host is the only system on the internal network that hosts on the internet can open connections to a (to deliver incoming email). Even then only certain types of connections are allowed. Any external system trying to access internal system or services will have to connect to this host. The bastion host thus needs to maintain a high level of host security.

Because this architecture allows packet to move from the internet to the internal networks, it may seem more risky than a dual homed host architecture, which is designed so that no external packet can reach the internal network. In practice, however, the dual homed host architecture is also prone to failures that let packets actually cross from the external network to the internal network. Furthermore it is easier to defend a router that it is to defend a host. For most purposes the screened host architecture provides both better security and better usability than the dual homed host architecture.

Compared to other architects, however, such as the screened subnet architecture, there are some disadvantage to the screened host architecture. The major one is that if an attacker manages to breaks in to the bastion host, nothing is left in the way of network security between the bastion host and the rest of the internal hosts. The router also presents a single point of failure; if the router is compromised, the entire network is available to an attacker. For this reason, the screened subnet architectures , discussed next,, has become increasingly popular.

Because the bastion host is a single point of failure, it is inappropriate to run high risk services like web servers on it. You need to provide the same level of protection to it that you would provide to a dual homed host that was the sole firewall for you site.

Appropriate uses

A screened host architecture is appropriate when:

❖ Few connections are coming from the internet (in particular it is not an appropriate architecture if the screened host is a public web server.)

- The network being protected has a relatively high level of host security.

Screened subnet architectures

The screened subnet architecture adds an extra layer of security to the screened host architecture by adding a perimeter network that further isolates the internal network from the Internet.

Why do this? By their nature, bastion hosts are the most vulnerable machines on your network. Despite you best efforts to protect them, they are the machines most likely to be attacked because they are the machines that can be attacked. If as in a screened host architecture, your internal network is wide open to attack from you bastion host, then you bastion host, then you bastion host is a very tempting target. No other defenses are between it and you other internal machines. If someone successfully breaks into the bastion hosts in a screened host architecture, that intruder has hit the jackpot. By isolating the bastion host on a perimeter network, you can reduce the impact of a breaks in on the bastion host. It is no longer an instantaneous jack pot it gives an intruder some access but not all.

With the simplest type of screened subnet architecture, there are two screening router, each connected to the perimeter net. One sits between the perimeter net and the internal network, and the other sits between the perimeter net and the external network. To break into the internal network with this type of architecture, an attacker would have to get past both router. Even if the attacker somehow broker in to the bastion host, he'd still have to get past the interiors router. There is not single venerable point that will compromise the internal network.

Bastion host

With the screened subnet architecture, you attach a bastion host to the perimeter net this host is the main point of contact for incoming connections form the outside.

- For incoming email session to deliver electronic mail to the site
- For incoming FTP connections to the site anonymous FTP server
- For incoming domain name system queries about the site and so on.

Outbound services are handled in either of these ways.

- Set up packet filtering on both the exterior and interior routers to allow internal clients to access external servers directly.

- Set up proxy servers to run on the bastion host to allow internal clients to access external servers indirectly. You would also set up packet filtering to allow the internal clients to talk to the proxy servers on the bastion host and vice versa, but to prohibit direct communications between internal clients and the outside world.

In either case, packet filtering allows the bastion host to connect to, and accept connections from, hosts on the Internet; which hosts, and for what servers, are dictated by the site's security policy.

16.6 Firewall Design

When you design a firewall you go through a process that you will then repeat over time as your needs change.

1. Define your needs
2. Evaluate the available products
3. Figure out how to assemble the products into a working firewall.

Define your needs

The first step in putting together a firewall is to figure out exactly what you need. You should do this before you start to look at firewall products, because other wise you risk being influenced more by advertising than by your own situation. This is inevitable and it has nothing to do with being gullible. If you don't know clearly what you need, the products that you look at will shape your decisions, no matter how suspicious you are.

You may need to re evaluate you needs if you find that there are no products on the market that can meet them of course, but at least you will have some idea of what you are aiming for.

What servers do you need to offer

You need to know what services are going to go between your site and the Internet. What will your users do on the Internet? Are you going to offer any services to users on the Internet. Are you going to let your users come into your site from the internet (if not, how are you providing your users with remote access)? Do you have special relationships with other companies that you are going to need to provide services for?

How secure do you need to be?

Many decisions have to do with relative levels of security. Are you trying to protect the world from destruction by protecting nuclear secrets, or do you want to keep from looking silly? Note that looking silly is not necessarily a trivial problem if you look silly on the front page of a major newspaper, it can be a real descants for the organization, at least. Many banks and financial institutions regard being above the fold as a significantly worse problem that losing money. One large organization if a small country fond that any time they appeared on the front page of the newspaper looking silly, their nations, currency dropped in value. You need to know what level of security you are aiming for .

How much usage will there be?

What kinds of network lines do you have? How many users will you have, and what will they do?

How much reliability do you need?

If you are cut off from the network what will happen? Will it be an inconvenience or a disaster?

What are you constraints?

Once you have determined what you need the firewall to do, your next job is to determine what the limits are.

What budget do you have available?

How much money can you spend, and what can you spend it on? Does personnel time count in the budget? How about consulting time? If you use a machine that you already own what does that do to your budget ? The budget is often the most visible constraint, but it tends to be then most flexible as well (as long as the organization you are building the firewall for actually has money somewhere.)

What personnel do you have available?

How many people do you have and what do they know? Personnel is much harder to change than budget even if you get agreement to hire people, you have to find them and interface them. Therefore, your first effort should be to fit the firewall to the available resources. If you have 46 windows NT administrators and one Unix person, start looking at windows NT based firewalls. If you have only one person to run the firewall and that is in additions to a full time job her or she is already doing get a commercial firewall and a consultant to install it.

What is you environment like?

Do you have political constraints? Are there forbidden operating systems or vendors or preferred ones? It is sometimes possible to work around these, but not always, for instance if you work for a company that sells firewalls, it is probably never going to be acceptable to run somebody else's firewall any where visible.
What country or countries are you going to needs to install the firewall in? firewalls often involve encryption technology, and laws about encryption and its export and import vary from country to country. If you are going to need to install multiple firewalls in different countries, you may need to use the lowest common denominator or develop an exception policy and strategy to deal with the situation.

16.7 Short Summary

A screening router is an appropriate firewall for a situation where

♣ The network being protected already has a high level of host security.

- The number of protocols being used is limited, and the protocols themselves are straightforward.

- You require maximum performance and redundancy.

- A dual homed host is an appropriate firewall for a situation where:

  - Traffic to the Internet is small
  - Traffic to the internet is not business critical
  - No services are being provided to Internet based users.
  - The network being protected does not contain extremely valuable data.

16.8 Brain Storm

1. What is firewall?
2. Explain the firewall technologies?
3. What is packet filtering?
4. Explain the architecture of firewall.
5. Explain the types of firewall architecture.

Lecture 17

Proxy Server

Objectives

After completing this lesson, you should be able to do the following

- Discuss about proxy server.
- Describe the services of proxy.
- Describe why we need proxy.
- Describe about how proxying works.
- Describe the terminology of proxy server.
- Discuss about proxying without a proxy server.

Coverage Plan

Lecture 17

17.1 Snap Shot

A proxy server allows the network administrator to implement a security policy that is more stringent than the packet-filtering router described later in this chapter.

Proxy servers manage the flow of Internet services through the firewall using special software called – what else? - a proxy service. A proxy service is loaded on the proxy server for each desired application. If you do not install a proxy service for a particular application, that service cannot be passed through the firewall. You can also configure a proxy service to allow only specific features of a particular application.

17.2 Proxy Server

A proxy server for a particular protocol or set of protocols runs on a dual-homed host or a bastion host: some host that the user can talk to, which can, in turn, talk to the outside world. The user's client program talks to this proxy server instead of directly to the "real" server out on the Internet. The proxy server evaluates requests from the client and decides which to pass on and which to disregard. If a request is approved, the proxy server talks to the real server on behalf of the client and proceeds to relay requests from the client to the real server, and to relay the real server's answers back to the client.

Proxy servers are sometime called bastion hosts. The bastion host takes in packets on one network connection and lets them out of another (your internal network). This way, the packets never flow directly from the Internet to your machines or vice versa. In the middle, the proxy software analyzes each packet and applies the network administrator's rule set. The rule set is the administrator's criteria for what is allowed to go through to the other side.

The bastion host hardware platform is a computer with two network ports. It uses a specially modified version of its operating system. One of the most popular bastion host operating systems is UNIX. If the bastion host is a UNIX platform, it executes a secure version of the UNIX operating system that is specifically designed to protect against outside access to the operating system. The proxy server is a small piece of code that is designed for network security. Proxy server code is kept small for speed and simplicity.

There will be a proxy routine for each service needed by the users. In other words, there can be a Gopher, mail, FTP, and Telnet proxy server. If you wanted to use a chat program, your packets would not be passed through the bastion host because there is no chat proxy.

Each proxy dumps detailed audit information by logging all traffic, each connection, and the duration of each connection. The audit log is the network administrator's list of clues for discovering and exposing access by unauthorized persons. You never know; one of those persons could be that mild-mannered reporter in the newsroom.

The bastion host can be programmed to require everything short of the user's thumbprint on floppy drive A. Anything goes here, from using remote authentication services to separate passwords and user IDs from two different individuals. It all depends on how the custom proxy code is written.

The proxies are usually written to support only a few functions or maybe even one function. The key design criteria are small, fast, and simple. Proxy code should not read the disk except when it is booted. Avoiding disk access greatly reduces the methods that a cracker can use to implant a bit of rogue code.

You might realize by now that the bastion host will require some love and care. If a proxy becomes the victim of some packet of intentional corruption, it will probably discontinue servicing the users. If this happens to the Gopher proxy, for example, the administrator will get mean shutting down the Gopher proxy and rebooting. Proxies are designed to be separate from each other to maintain independence and not crash the entire server. One proxy can be corrupt, and only that service – for instance, the Gopher service- will be affected. The remaining services – FTP, HTTP, and so on – will continue to perform their duties.

There are not many drawbacks to using the proxy server bastion host method. Perhaps the one of most consequence is the training needed to care for proxies on this setup. Of course, you can rely on your vendor to provide support, but you might want to make sure that its code is doing for you what you require.

As far as the user is concerned, talking to the proxy server is just like talking directly to the real server. As far as the real server is concerned, it's talking to a user on the host that is running the proxy server; it doesn't know that the user is really somewhere else.

Since the proxy server is the only machine that speaks to the outside world, it's the only machine that needs a valid IP address. This makes proxying an easy way for sites to economize on address space. Network address translation can also be used to achieve this end.

Proxying doesn't require any special hardware, but something somewhere has to make certain that the proxy server gets the connection. This might be done on the client end by telling it to connect to the proxy server, or it might be done by intercepting the connection without the client's knowledge and redirecting it to the proxy server.

17.3 Proxy Services

In general, a proxy is something or someone who does something on somebody else's Behalf. For instance, you may give somebody the ability to vote for you by proxy in an election.

Proxy services are specialized application or server programs that take users' requests for Internet services (such as FTP and Telnet) and forward them to the actual services. The proxies provide replacement connections and act as gateways to the services. For this reason, proxies are sometimes known as application –level gateways. In this book, when we are talking about proxy services, we are specifically talking about proxies run for security purposes, which are run on a firewall host: either a dual-homed host with an interface on the internal network and once on the external network, or some other bastion host that has access to the Internet and is accessible from the internal machines.

You will also run into proxies that are primarily designed for network efficiency instead of for security; these are caching proxies, which keep copies of the information for each request that they proxy. The advantage of a caching proxy is that if multiple internal hosts request the same data, the data can be provided directly by the proxy. Caching proxies can significantly reduce the load on network connections. There are proxy servers that provide both security and caching, in general, they are better at one purposing tan the other.

Proxy services sit, more or less transparently, between a user on the inside (on the internal network) and a service on the outside (on the Internet). Instead of taking to each other directly, each talks to a proxy. Proxies handle all the communication between users and Internet services behind the scenes.

Transparency is the major benefit of proxy services. It's essentially smoke and mirrors. To the user, a proxy server presents the illusion that the user is dealing directly with the real server. To the real server, the proxy server presents the illusion that the real server is dealing directly with a user on the proxy host (as opposed to the user's real host).

How do proxy services work? Let's look at the simplest case, where we add proxy services to a dual-homed host. A proxy service requires two components: a proxy server and a proxy client. In this illustration, the proxy server runs on dual-homed host. A Proxy client is a special version of a normal client program (e.g., a Telnet or FTP client) that talks to the proxy server rather than to the "real" server out on the Internet; in some configurations, normal client programs can be used as proxy clients. The proxy server evaluates requests from the proxy client and decides which to approve and which to deny.

If a request is approved, the proxy server contacts the real server on behalf of the client (thus the term proxy) and proceeds to relay requests from the proxy client to the real server, and responses from the real server to the proxy client.

In some proxy systems, instead of installing custom client proxy software, you'll use standard software but set up custom user procedures for using it.

There are also systems that provide a hybrid between packet filtering and proxying where a network device intercepts the connection and acts as a proxy or redirects the connection to a proxy; this allows proxying without making changes to the clients or the user procedures.

The proxy server doesn't always just forward user's requests on to the real Internet services. The proxy server can control what users do because it can make decisions about the requests it processes. Depending on your site's security policy, requests might be allowed or refused. For example, the FTP proxy might refuse to let users export files, or it might allow different capabilities to different hosts, rather than enforcing the same restrictions on all hosts.

Some proxy servers do in fact just forward requests on, no matter what they are. These may be called generic proxies or port forwarders. Programs that do this are providing basically the same protections that you would get if you had a packet filter in place that was allowing traffic on that port. You do not get any significant increase in security by replacing packet filter with proxies that do exactly the same thing (you gain some protection against malformed packets, but you lose by adding an attackable proxying program).

Some excellent software is available for proxying. SOCKS is a proxy construction toolkit, designed to make it easy to convert existing client/server applications into proxy versions of those same applications. The Trusted Information System Internet firewall Toolkit (TIS FWTK) include proxy servers for a number of common Internet protocols, including Telnet, FTP, HTTP, rlogin, XII, and others, these proxy servers are designed to be used in conjunction with custom user procedures.

Many standard client and server programs, both commercial and freely available, now come equipped with their own proxying capabilities or with support for generic prosy systems like SOCKS. These capabilities can be enabled at runtime or compile time.

Most proxy systems are used to control and optimize outbound connections, they are controlled by the site where the clients are. It is also possible to use proxy systems to control and optimize inbound connections to servers (for instance, to balance connections among multiple servers or to apply extra security). This is sometimes called reverse proxying.

17.4 Why Proxying?

There's no point in connecting to the Internet if your users can't access it. On the other hand, there's no safety in connecting to the Internet if there's free access between it and every host at your site. Some compromise has to be applied.

The most obvious compromise is to provide a single host with Internet access for all your users. However, this isn't a satisfactory solution because these hosts aren't transparent to users. Users who want to access network services can't do so directly. They have to log in to the dual-homed host, do all their work from there, and then somehow transfer the results of their work back to their own workstations. At best, this multiple-step process annoys users by forcing them to do multiple transfers and work without the customizations they're accustomed to.

The problem is worse at sites with multiple operating systems; if your native system is a Macintosh, and the dual-homed host is a Unix system, the Unix system will probably be completely foreign to you. You'll limited to using whatever tools are available on the dual-homed host, and these tools may be completely unlike the tools you use on your own system.

Dual-homed hosts configured without proxies therefore tend to annoy their users and significantly reduce the benefit people get from the Internet connection. Worse, they usually don't provide adequate security; it's almost impossible to adequately secure a machine with many users, particularly when those users are explicitly trying to get to the external universe. You can't effectively limit the available tools because your users can always transfer tools from internal machines that are the same type. For example, on a dual-homed host, you can't guarantee that all file transfers will be logged because people can use their own file transfer agents that don't do logging.

Proxy systems avoid user frustration and the insecurities of a dual-homed host. They deal with user frustration by automating the interaction with the dual-homed host. Instead of requiring users to deal directly with the dual-homed host, proxy systems allow all interaction to take place behind the scenes. The user has the illusion of dealing directly with the server on the Internet, with a minimum of direct interaction with the dual-homed host. Below the figure illustrates the difference between reality and illusion with proxy systems.

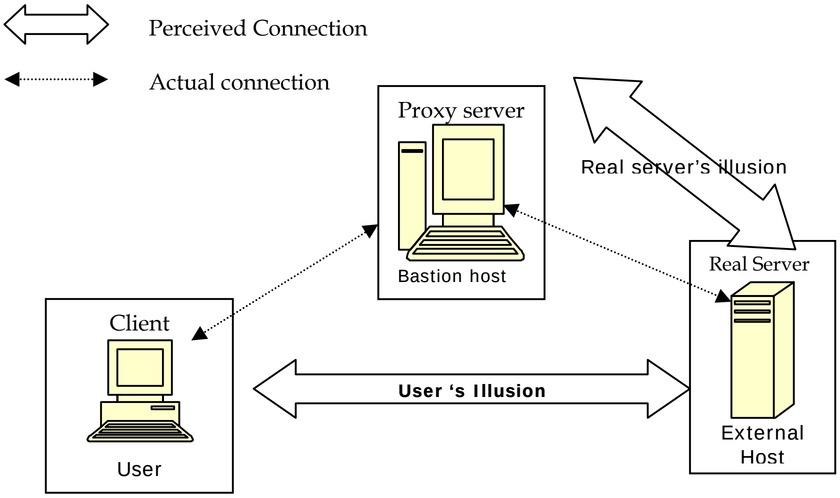

Proxy systems deal with the insecurity problems by avoiding user logins on the dual-homed host and by forcing connections through controlled software. Because the proxy software works without requiring user logins, the host it runs on is safe from the randomness of having multiple logins. It's also impossible for anybody to install uncontrolled software to reach the Internet; the proxy acts as a control point.

17.5 How Proxying Works

The details of how proxying works differ from service to service. Some services provide proxying easily or automatically; for the services, you set up proxying by making configuration changes to normal servers. For most services, however, proxying requires appropriate proxy server software on the server side. On the client side, it needs one of the following:

Proxy-aware application software

With this approach, the software must know how to contact the proxy server instead of the real server when a user makes a request (for example, for FTP or Telnet), and how to tell the proxy server what real server to connect to.

Proxy-aware operating system software

With this approach, the operating system that the client is running on is modified so that IP connections are checked to see if they should be sent to the proxy server. This mechanism usually depends on dynamic runtime linking (the ability to supply libraries when a program is run). This mechanism does not always work and can fall in ways that are not obvious to users.

Proxy-aware user procedures

With this approach, the user uses client software that doesn't understand proxying to talk to the proxy server and tells the proxy server to connect to the real server, instead of telling the client software to talk to the real server directly.

Proxy-aware router

With this approach, nothing on the client's end is modified, but a router intercepts the connection and redirects it to the proxy server or proxies the request. This requires an intelligent router in addition to the proxy software (although the routing and the proxying can co-exist on the same machine).

Using Proxy-Aware Application Software for Proxying

The first approach is to use proxy-aware application software for proxying. There are a few problems associated with this approach, but it is becoming easier as time goes on.
Appropriate proxy-aware application software is often available only for certain platforms. If it's not available for one of your platforms, your users are pretty much out of luck. For example, the *gateway* package from Sun is a proxy package for FTP and Telnet, but you can use it only on Sun machines because it provides only precompiled Sun binaries. If you're going to use proxy software, you obviously need to choose software that's available for the needed platforms.

Even if software is available for your platforms, it may not be software your users want. For sample, dozens of FTP client programs are on the Macintosh. Some of them have really impressive graphical user interfaces. Others have other useful features; for example, they allow you to automate transfers. You're out of luck if the particular client you want to use, for whatever reason, doesn't support your particular proxy server mechanism. In some cases, you may be able to modify clients to support your proxy server, but doing so requires that you have the source code for the client, as well as the tools and the ability to recompile it. Few client programs come with support for any form of proxying.

17.6 Proxy Server Terminology

This section describes a number of specific types of proxy servers.

Application-Level Versus Circuit-Level Proxies

An application-level proxy is one that knows about the particular application it is providing proxy services for; it understands and interprets the commands in the application protocol. A circuit-level proxy is one that creates a circuit between the client and the server without interpreting the application protocol. The most extreme version of an application-level proxy is an application like Sendmail, which implements a store-and-forward protocol. The most extreme version of a circuit-level proxy is an application like plug-gw, which accepts all data that it receives and forwards it to another destination.

The advantage of a circuit-level proxy is that it provides service for a wide variety of different protocols. Most circuit-level proxy servers are also generic proxy servers; they can be adapted to serve almost any protocol. Not every protocol can easily be handled by a circuit-level proxy, however. Protocols like FTP, which communicate port data from the client to the server, require some protocol-level intervention, and thus some application-level knowledge. The disadvantage of a circuit-level proxy. Like a packet filter, it controls connections on the basis of their source and destination and can't easily determine whether the commands going through it are safe or even in the expected protocol. Circuit-level proxies are easily fooled by servers set up at the port numbers assigned to other services.

In general, circuit-level proxies are functionally equivalent to packet filters. They do provide extra protection against problems with packet headers (as opposed to the data within the packets). In addition, some kinds of protections (protection against packet fragmentation problems, for instance) are automatically provided by even the most trivial circuit-level proxies but are available only from high-end packet filters.

Generic Versus Dedicated Proxies

Although "application-level" and "circuit-level" are frequently used terms in other documents, we more often distinguish between "dedicated" and "generic" proxy servers. A dedicated proxy server is one that serves a single protocol; a generic proxy server is one that serves multiple protocols. In practice, dedicated proxy servers are application-level, and generic proxy servers are circuit-level. Depending on how you argue about shades of meaning, it might be possible to produce a generic application-level proxy server or a dedicated circuit-level proxy server. Neither of these ever occur, however, so we use "dedicated" and "generic" merely because we find them somewhat more intuitive terms than "application-level" and "circuit-level".

Intelligent Proxy Servers

A proxy server can do a great deal more than simply relay requests; one that does is an intelligent proxy server. For example, almost all HTTP proxy servers cache data, so that multiple requests for the same data don't go out across the Internet. Proxy servers can provide better logging and access controls than those achieved through other methods, although few existing proxy servers take full advantage of the opportunities. As proxy servers mature, their abilities are increasing rapidly. Now that there are multiple proxy suites that provide basic functionally, they're beginning to compete by adding features. It's easier for a dedicated, application-level proxy server to be intelligent; a circuit-level proxy has limited abilities.

17.7 Proxying without a Proxy Server

Some services, such as SMTP, NNTP, and NTP, naturally support proxying. These services are all designed so that transactions (e-mail messages for SMTP, Usenet news postings for NNTP, and clock settings for NTP) move between servers, instead of going directly from a client to a final destination server. For SMTP, the messages are forwarded towards an email message's destination. NNTP forwards messages to all neighbor servers. NTP provides time updates when they're requested but supports a hierarchy of servers. With these schemes, each intermediate server is effectively acting as a proxy for the original sender or server.

If you examine the "Received:" headers of incoming Internet email (these headers trace a message's path through the network from sender to recipient), you quickly discover that very few messages travel directly from the sender's machine to the recipient's machine. It's far more common these days for the message to pass through at least four machines:

∗ The sender's machine
∗ The outgoing mail gateway at the sender's site (or the sender's Internet service provider)
∗ The incoming mail gateway at the recipient's site
∗ Finally, the recipient's machine

Each of the intermediate servers (the mail gateways) is acting as a proxy server for the sender, even though the sender may not be dealing with them directly.

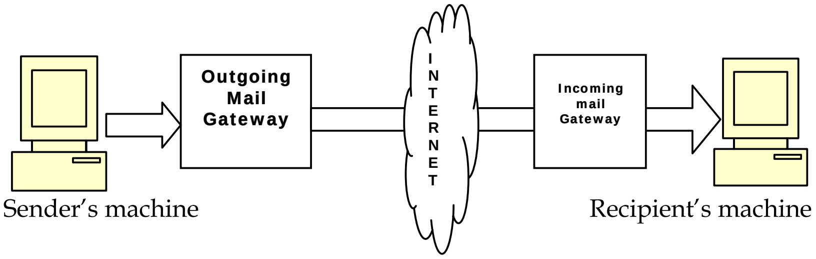

17.8 Short Summary

Proxy server bastion hosts have many benefits. They give the network administrator complete control over which services are permitted, have the capability to support strong user authentication, and provide detailed logging information. Also, the filtering rules for a proxy server are much easier to configure than for other technologies such as the packet-filtering router.

17.9 Brain Storm

1. Explain the types of proxy servers.
2. How proxy server works explain
3. What is proxying?
4. What software's is need for client side proxying?

Lecture 18

Packet Filtering

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the packet filtering.
- Describe the concepts of filtering by address.
- Describe the concepts of filtering by service.
- Describe about packet filtering router.
- Describe the rules of filtering.

Coverage Plan

Lecture 18

18.1 Snap Shot

If you put enough work into it, you can do anything you want to with packet filtering; all of the information that crosses the internet has to o into a packet at some point, after all. But some things are very much easier to do than others. For instance, operations that require detailed protocol knowledge or prolonged tracking of past events are easier to do in proxy systems. Operations that are simple but need to be done fast and on individual packets are easier to do in packet filtering systems.

18.2 Packet Filtering

Packet filtering is a network security mechanism that works by controlling what data can flow to and from a network. The basic device that interconnects IP networks is called a router. A router may be a dedicated piece of hardware that has no other purpose, or it may be apiece of software that runs on a general purpose compute running UNIX windows NT, or another operating system packets traversing an internet work travel from router to router until they reach their destination. The internet itself is sort of the granddaddy of internet works the ultimate “ network of networks”

A router has to make a routing decision about each packet it receives it has to decide how to send that packet on towards its ultimate destination. In general a packet carries no information to help the router in this destination. The packet tells the router where it wants to go but not how to get there. Routers communicate with each other using routing protocols such as the routing information protocol and open shortest path first to build routing tables in memory to determine how to get the packet to their destinations. When routing a packets a router compares the packet’s destination addresses to entire is in the routing table and sends packet onward as directed by the routing table. Often, there won’t be a specific route for a particular destination, and the router will use a default route generally, such a route directs the packet towards smarter or better connected routers.

Packet filtering devices that keep track of packets that they see are frequently called stateful packet filters (because they keep information about the state of transactions). They may also be called dynamic packet filters because they change their handling of packets dynamically depending on the traffic they see. Devices that look at the content of packets , rather than at just their headers are frequently called intelligent packet filters In practice almost all stateful packet filters also are campanile of looking at the contents of packets and many are also cable of modifying the contents of packets so you may see all these capabilities lumped together under the heading “stateful packet filtering” However something can legitimately be called a “stateful Packet filter” without having the ability to do advanced content filtering of modification.

A packet filtering system is also a logical place to provide virtual private network or network address translation services. Since the packet filter is already looking at all of the packets, it can easily identify packets that are intended for a destination that is part of the virtual private network encrypt those packets and encapsulate them in another packet bound for the appropriate destination.

ADVANTAGES OF PACKET FILTERING

Packet filtering has a number of advantages

One screening router can help protect an entire network**:** One of the key advantages of packet filtering is that a single, strategically placed packet filtering router can help potent an entire network If only one router connects your site to the Internet you gain tremendous leverage on network security regardless of the size of your site by doing packet filtering on that router.

Simple packet filtering is extremely efficient

Because simple packet filtering requires paying attention only to a few packet headers it can be done with very low overhead Proxying is a fairly time consuming operation and adding proxying means directing connections through another program usually on a machine that otherwise wouldn’t n be necessary to the rooting process Packet filtering takes place on a machine tea was already in the critical path , and introduces a much smaller delay.

However there is no free lunch; the more work your packet filters do the slower they will be . If your packet filters behave like proxies ,doing complicated datadriven operation that require keeping track of multiple packets they will tend to perform like proxies as well.

Packet filtering is widely available

Packet filtering cap anilities are available in many hardware and software routing products both commercial and freely available over the internet . Most sites already have packet filtering capabilities available in the routers they use.

Most commercial router products include packet filtering capabilities. Packet filtering capabilities are also available for a number of general-purpose computers. These are discussed further in the chapter, Packet Filtering.

18.3 Filtering by Address

The simplest, although not the most common form of packet filtering is filtering by address. Filtering in this way lets you restrict the flow of packets based on the source and or destination address of the packets without having to consider what protocols are involved. Such filtering can be used to allow certain external host a to talk to certain internal hosts, for example or to prevent an attacker from injecting forged packets(packets handcrafted so they appear to come form some where other than their true source) into your network

For example let's say that you want to block incoming packets with forged source addresses you would specify the following rule.

Rule	Direction	Source address	Dest. Address	Action
A	Inbound	Internal	Any	deny

Note that direction is relative to your internal network. In the router between your internal network and the internet. You could apply an inbound rule either to incoming packets on the internet interface or to outgoing packets on the internal interface .or to outgoing packets on the internal interface. Either way you will achieve the same results for the protected hosts. The difference is in what the router itself sees. If you filter outgoing packets, the router is not protecting itself.

Risks of filtering by source address

It is not necessarily safe to trust source addressees because source addresses can be forged. Unless you use some kind of cryptographic authentication between you and the host you want to talk to you won't know it you're really talking to that host , or to some other machine that is pretending to be that host. The filters we have discussed previously will help you if an external host is claiming to be an internal host but they won't do anything about an external host claiming to be a different external host.

There are two kinds of attacks that rely on forgery source address and main the middle.
In a basic source address forgery attack an attacker sends you packets that claim to before someone you trust in some way, hoping to get you to take some action based on that trust, without expecting to get any packets back form you. If the attacker doesn't care about getting packets back from you, it doesn't mater where the attacker is. In fact, your responses will go to who ever the attacker is pretending to be, not to the attacker. However if the attacker can predict your responses, it doesn't mater that they are going somewhere else. Many protocols are predictable enough for a skilled attacker to be successful at this. Plenty of attacks can be carried out without the attacker's needing to see the results directly. For example, suppose an attacker issues a command to your system that causes it to mail back your password file, if your system is going to send the attacker the password file in the mail there is no need to see it during the attack itself.

In many circumstances- particularly those involving TCP connections – the real machine (that the attacker is pretending to be)will react to your packets(packets that are attempting to carry on a conversation it knows nothing about) by trying to reset the bogus connection. Obviously , the attacker doesn't want this to happen. Therefore the attack must complete before the real machine gets the packets you're sending or before you get the reset packets from the real machine. There are a number of ways to ensure this – for example:

- Carrying out the attack while the real machine is down
- Crashing the real machine so the attack can be carried out
- Flooding the real machine while the attack is carried out
- Using an attack where only the first response packet is required, so that the reset doesn't matter.

Attacks of this kind used to be considered a theoretical problem with little real world effect but they are now common enough to be considered a serious threat.

The man in the middle forgery attack depends on being able to carry out a complete conversation while claiming to be the trusted host. In order to do this the attacking machine needs to be able to not only send you packets but also intercept the packets you reply with. To do this the attacker needs to do one of the following:

- Insinuate the attacking machine into the path between you and the real machine . This is easiest to do near the ends of the path, and most difficult to do somewhere in the middle, because given the nature o0f modern IP networks, the path through "the middle" can change at any second.

- Alter the path between the machines so it leads through the attacking machine. This may be very easy or difficult, depending on the network topology and routing system used by your network the remote network and the Internet service providers between those networks.

Although this kind of attack is called "man in the middle", its relatively rare for it to actually be carried out in the middle (external to the sites at each end) because nobody but a network provider is in a position to carry it out in that way,. And network providers are rarely compromised to that extent (people who compromise network providers tend to be working on quantity. Packet sniffing will give them many hosts rapidly, but man in the middle attacks give them only one site at a time) Theses attacks tend to be problems only if one of the involved sites has hostile users who have physical access to the network (for example this might be the case if one site is a university) .

So, whom can you trust? At the extreme, nobody, unless you trust the machines involved at both ends and the path between them. If you trust the machines but not the path you can use encryption and integrity protection to give you secure connection over an insecure path.

18.4 Filtering by Service

Blocking incoming forged packets, as discussed previously, is just about the only common use of filtering solely by address. Most other uses of packet filtering involve filtering by service, which is somewhat, more complicated.

Outbound Telnet Service

Let's look first at outbound Telnet Service in which a local client is talking to a remote server, We need to handle both outgoing and incoming packets show a simplified view of outbound Telnet.

The outgoing packets for this outbound service contain the user's keystrokes and have the following characteristics:

- ❖ The IP source address of the outgoing packets is the local host's IP address
- ❖ The destination address is the remote host's IP address.
- ❖ Telnet is a TCP-based service so the IP packet type is TCP
- ❖ The TCP destination port is 23 that's the well-known port number Telnet servers use.
- ❖ The TCP source port number (which we'll call "Y" in this example)is some seemingly random number greater than 1023.
- ❖ The first outgoing packet establishing the connection will not nave the ACK bit set the rest of the outgoing packets will.

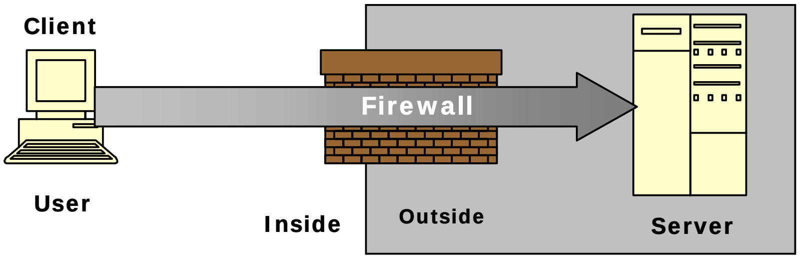

The incoming packets for this outbound service contain the data to be displayed on the user's screen (for example the "login:" prompt)and have the following characteristics:

The IP source address of the incoming packets is the remote host's IP address

- ❖ The IP destination address is the local host's IP address
- ❖ The IP packet type is TCP
- ❖ The TCP source port is23 that's the port the server uses
- ❖ The TCP destination port is the same "y" we used as the source port for the outgoing packets

- All incoming packets will have the ACK bit set(again, only the first packet establishing a connection, has the ACK bit off; in this example, that first packet was an outgoing packet ,not an incoming packet)

Note the similarities between the header fields of the outgoing and incoming packets for Telnet. The sane addresses and port numbers are used they're just exchanged between source and destination. If you compare an outgoing packet to an incoming packet, the source and destination addresses are exchanged and the source and destination port numbers are exchanged.

Inbound Telnet Service

Next let's look at inbound Telnet service in which a remote client communication with a local Telnet server Again we need to handle both incoming and outgoing packets:

The incoming packets for the inbound Telnet service contain the user's keystrokes and have the following characteristics:

- The IP source address of these packets is the remote host's address
- The IP destination address is the local host's address
- The IP packet type is TCP
- The TCP source port is some random port number greater than 1023(which we'll call "z" in this example)
- The TCP destination port is 23
- The TCP ACK bit will not be set on the very first inbound packet, establishing the connection but it will be set on all other inbound packets.

The outgoing packets for this inbound Telnet service contain the server responses and have the following characteristic:

- The IP source address is the local host's address
- The IP destination address is the remote host's address
- The IP packet type is TCP
- The TCP source port is 23
- The TCP destination port is the same random port "z" that was used as the source port for the inbound packets
- The TCP ACK bit will be set on outgoing packets

Again note the similarities between the relevant headers of the incoming and the outgoing packets: the source and destination addresses are exchanged and the source and destination ports are exchanged.

18.5 Packet Filtering Router

A number of packet filtering routers available, some good and some not so good. Almost every dedicated router supports packet filtering in some form. In addition, packet-filtering packages are available for many general-purpose Unix and PC platforms you might want to use as routers

How do you choose the best packet filtering router for your site? This section outlines the most important capabilities a filtering router should have. You should determine which of these capabilities are important to you and select a filtering system that offers at least those capabilities.

18.6 Rules of Filtering

- ❖ It should Have Good Enough Packet Filtering Performance for Your Needs

In addition, firewall performance depends on the complexity of packet filters. You should be sure that the speeds you are quoted are speeds with a reasonable filter ser (some manufactures quote the speed achieved with packet filtering enabled but no filters set for instance) Stateful packet filtering, intelligent packet filtering and reassembly of fragmented packets will all slow down performance.

A firewall with more than two connections may also have higher speed requirements. With two connections the maximum required speed is that of the slowest connection with three connections the required speed can rise. For example if you put a second Internet connection onto an external router, it now needs to drive both at full speed if it's not going to be a limiting factor. If you put two internal networks onto it, it's going to be a limiting factor. If you put two internal networks onto it, it's going to be achieve the higher speed of those not works to route between them.

- It Can Be a Single-Purpose Router Or a General-Purpose Computer

Don't expect a single device to serve as your packet filtering router and also to do something that's not part of your firewall. (You may have a device that's doing packet filtering and proxying or packet filtering and selected bastion host services or even all three) In a practical sense you should expect to be using a dedicated packet filtering router. This doesn't mean you have to buy a single-purpose router, however. You might choose to use either a traditional, single –purpose router or a general-purpose computer dedicated to routing. What are the pros and cons of each choice?

If you have a large number of networks or multiple protocols you will probably need a single-purpose router Routing packages for general –purposes computers usually do not have the speed or flexibility of single-purpose routers and you may find that you will need an inconveniently large machine to accommodate the necessary interface boards.
On the other hand if you are filtering a single Internet link you may not need to do any more than route IP packets between two Ethernets. This is will within the capabilities of a reasonable 486-based computer and such a machine will certainly be cheaper than a single purpose router (It may even be free if you already have one available within your organization) Routing and filtering packages ate available for Windows NT and many other Microsoft operating systems as well as most variants of Unix.

- It Should Allow Simple Specification of Rules

- It Should Allow Rules Based on Any Header Or Meta-Packet Criteria

- It Should Apply Rules in the Order Specified

- It Should Apply Rules Separately to Incoming and Outgoing Packets on a Per-Interface Basis

- It Should Be Able to Log Accepted And Dropped Packets

What Rules Should You Use?

Clearly most of the rules that you will put into your packet filtering system will be determined by the kinds of traffic you want to accept. There are certain rules you will almost always want to use, however

We've already discussed these rules in various places but here's a summary list of some standard protection that you should automatically apply unless you have a strong reason to do otherwise :

- Set up an explicit default deny so that you are sure that the default behavior is to reject packets
- Deny inbound traffic that appears to come from internal addresses (this is an indication of forged traffic or bad network configuration)
- Deny outbound traffic that does not appear to come from internal addresses (again such traffic is either forged or symptomatic of network misconfigurations)
- Deny all traffic with source or IP options set
- Deny ICMP traffic over a reasonable size
- Reassemble fragments into entire packets.

18.7 Short Summary

Packet filtering, every packet has a set of headers containing certain information. The main information is:

- IP source address
- IP destination address
- Protocol (whether the packet is TCP, UDP, or ICMP packet)
- TCP or UDP source port
- TCP or UDP destination port
- ICMP message type
- Packet size

18.8 Brain Storm

- What is packet filtering?
- How to filter the packet by address explain.
- What is the difference between filtering by address and filtering by service?
- Explain the rules of filtering.
- What is packet filtering router?

Lecture 19

Internet Application

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the application of Internet.
- Describe the Tools and utilities of Internet.
- Describe about electronic mail.
- Describe about Internet relay chat and browser.
- Describe the world wide web and uniform resource locator.

Coverage Plan

Lecture 19

19.1 Snap Shot

This chapter looks at a basic set of applications that most Internets use day in and day out and that will run on Windows NT 4.0. At first, you might not use all of them. Your first Internet roaming will most likely require the master gadget of them all, the browser. Your Internet provider might give you its customized browser, or you might be left to obtain your own. Either way, I'll discuss the basics of the browser and how to set it up in.

This chapter discusses the browser tool – specifically, Microsoft Internet Explorer 3.0. With this single tool you can accomplish virtually everything mentioned here. I will give you some understanding of what the browser is all about. First, I will give you some insight as to how the browser gets its job accomplished by showing you some of the individual applications like FTP, Veronica, Gopher, the World Wide Web, Chat, and a few others. All of the package described here are known to work with Microsoft NT Server 4.0.

19.2 Internet Application

Your second Internet experience will no doubt be your e-mail. How excited you will be when someone actually sends you mail over your new Internet connection! The world will be at your fingertips.

You will discover File Transfer Protocol (FTP) and the Archie search client, and then there will be "Look out, Internet – here I come". With FTP, you will daringly extract software from the thousands of available archives and sites.

Soon you will venture into the Usenet newsgroups. If you cannot find your favorite topic, you have not looked around enough, because it is out there. The day will come when you post your own message to the newsgroup. You will also be perfectly normal when you wait a while and then look to see that it is actually there for all to see. What a rush!

Immediately you will be somewhat overwhelmed by the vastness of this resource you have acquired. The archives of seemingly endless subdirectories are everywhere you look. You will learn about the search tools out of sheer necessity for your survival in the Internet jungle.

When you have your confidence sufficiently bolstered, you might even log onto Internet Relay Chat (IRC), and your pulse will quicken when you are discovered online and receive a welcome.

I will present a high-level look at these applications and leave it to you to plunder the depths of each as you desire. Internet applications come in four flavors: free, shareware, you-buy-it, and demos. You will find excellent free software (called freeware) on the Internet. Do not be afraid to use it; many have gone before you. The idea behind shareware is that you try the software and, if you like it, send the author a modest amount of compensation for his time. It is an honor system for most applications. The you-buy-it classification is clear - you pay and then you use, just like going to your favorite software outlet store. Demos just show you the product; they do not provide all the features, like print or save.

19.3 Tools And Utilities

Every trade has its tools – machine, mason, carpenter, draftsperson, librarian, and yours, whatever it may be. The Internet is no different when it comes to a set of tools and utilities to get the job done and done well. The tools described here are FTP, Telnet, Gopher, Veronica, Ping, and Archie. Remember, there are hundreds more but we are trying to limit the range to just the basic applications in this chapter.

File Transfer Protocol

File Transfer Protocol (FTP) is one of the all-time greats of networking software. With it you can fill a 2GB hard disk in no time at all and, if you are like the rest of us, that is exactly what you will do. That is part of the great learning experience that the Internet offers. You can download programs and try them and download programs and try them and … you will reach the point where something has to go. Then the cycle starts all over again. You will use the anonymous login to begin your plucking of the Internet chicken.

FTP gives you the capability to send or receive files with another computer on a network. In this case I am talking about the Internet when I say network. Anonymous login means that you do not have a personal account on the remote computer, so you have to log in anonymously. The remote servers that will allow anyone to connect and transact file transfers are called anonymous FTP servers. Like all servers, they too require a username and password to let you connect. Today, anonymous FTP servers are looking for a username of anonymous and a password of your e-mail address.

You can FTP with only a dial-up Internet account, a terminal program like the HyperTerm in Windows95 and NT, and a modem. A few years ago this was the only way to get started with browsing the Internet. Now, when you obtain access to the Internet, you generally receive some software to get you started, and what you receive depends entirely on your selected Internet service provider (ISP).

The OAK Software Repository (oak.oakland.edu) is a public service of Oakland University's Office of Computer and Information Services. OAK offers many collections of computer software and information to Internet users free of charge. To use basic FTP to obtain files from the OAK Repository, you start a terminal program, log in to your Internet service provider, and then type in the following:

```
ftp oak.oakland.edu
```

you will soon see a reply:
connected to oak.oakland.edu

Next you will see lots of stuff about the operating system and accepting your username and password, and then there will be a prompt:

```
ftp>
```

At this point you are ready to begin moving files with the FTP command set. The tricky thing about using this method of FTP is that while you are using a computer running a Windows-based operating system, the remote machine machine is most likely UNIX based. When you use a terminal program to access the remote machine, you need to use the syntax of that machine, not yours. That means you will need to know the UNIX-based FTP commands and that the slash in your DOS c:\is actually c:/ in UNIX.

The FTP command set is not difficult to master, but in this day and age it is not necessary unless you are just itching to step back in time and use some black-screen command-line methods.

There are several Windows-based FTP clients available on the Internet. One that is readily available and free (which makes it even better) is WS_FTP32 by JohnJunod.

With WS_FTP32, you just start the program. The first time you run it, you will have to enter your user ID and password, which will be anonymous and yourname@yourISP, respectively. Don't forget to click the Save button so you do not have to enter it again. Then you select the site you want to access and watch it all go. Your screen will show you your directory and the remote directory with its respective files. The program gives you complete control of the remote computer and its files, to the extent that the systems administrator has permitted.

ARCHIE

File searches on the Internet with a program called Archie. Archie sends queries to special Archie servers that are specifically designed to list files on anonymous FTP sites. The search can be done for a very specific file/program name or even on a fragment of a name. The server will return information listing anonymous FTP sites that carry the corresponding file. Archie servers are located in more than a dozen major countries around the world. Archie servers are usually backlogged with a large number of requests during the local prime-time hours. You can apply your grade school knowledge of world where the locals are not using their servers. The information available on different Archie servers is not necessarily the same, and sometimes a search of more than one Archie server is necessary.

The search program WSARCHIE32 and John Junod's file-retrieval program WS_FTP32 make a substantial pair because they integrate well together. Once WSARCHIE32 is configured with your user name and password (your e-mail address) and the location of WS_FTP32, you are set to go. You enter your search terms in the Archie client, and it will locate the file you are seeking. Then, when you right-click the mouse on the file. You will have the option to retrieve it. If you choose to do so, control is passed onto WS_FTP32, which brings the file home.

Archie searches can be a few characters of the filename or the full name of the file. WSHARCHIE32 allows complex queries beyond even these with domains, wildcards, and more.

An Archie search can be initiated after starting WSARCHEIS32 by typing in a search string in the Search for box. Select an Archie server from the Archie server dropdown list. Next, if desired, you may enter a match domain that is used to restrict the search to a subset of hosts within a particular domain; for example, .ac.uk will restrict the search to academic sites in the UK.

The type of search performed by the Archie client is determined using the search type radio buttons shown under the Search button. The choices have the following effects:

Substring: This is a simple, everyday Substring search. A match occurs if the file (or directory) name in the database being searched contains the user-given Substring.
Substring: This option is the same as a basic Substring search, but the case of the strings involved becomes significant.

Exact: This is the fastest search method of all. The restriction is that the user string has to exactly match the string in the database. This is provided for those who know just what they are looking for.

Regex: The Regex option searches the database with the user string.

Exact first: This is a check box that means the exact search method will be used first. If no matches are found, the method selected using the radio buttons will be used. This box is disabled if the exact search method is chosen.

Once you have determined the search method, you click the search button and WSARCHI32 goes to work. While the query is being performed, various status indications are given so that you can see the progress of the query. The indications are in the status bar and also in the title bar of the application.

When the search is finished and has located your files, the results are displayed in the Hosts, Directories, Files, and File Details sections of the WSARCHIE32 program screen. This program comes with a good help section where you will find details on query formats.

Telnet

Telnet is a program that lets you connect to another machine as if you were sitting at the machine's user terminal. You need three things to use Telnet: a network connection or modem to another computer, an account on that computer, and Telnet software on your computer. You can use Telnet to do neat things like connect to the Library of Congress or to other city and county libraries.

Many Telnet clients can connect by modem only. Better designs are able to connect via the network connection so you avoid the long distance charges that you would incur otherwise. There is also a version of Telnet called TN3270 that is better for connecting to an IBM site. TN3270 provides the keyboard mapping that emulates a standard IBM3270 terminal.

Telnet is started from the black-screen command line with a command like this:

```
telnet [host [port]]
```

If you are using your browser and have it configured with an external Telnet helper application, you can access a Telnet site by going to your browser's open dialog box and entering this:

telnet:// [host [port]] instead of http://www.something

When you Telnet to another computer, you are usually remotely accessing the operating system of a UNIX computer. Many of these computers are set up specifically for some application; for example, you can access university card catalogs using Telnet. To use Telnet, you need to know the name of the computer to which you want to connect.

Telnet does not understand Hypertext Markup Language (HTML), embedded images mice, or the other friendly features we associate with the World Wide Web. Also, note that Telnet is point-to-point communication between two machines, if more than two sites need to communicate, Telnet alone will not do the job.

Gopher

The Gopher client provides document-retrieval functions. It has tables of directory entries on Gopher servers anywhere in Internet space. Gopher is a prime example of distributed computing because its main purpose for existing is to be a distributed document-retrieval system. A Gopher server can store text, binary, image, and sound data. The Gopher system's linkage to other Gopher servers makes it an extremely powerful application for locating documents on the Internet.

Gopher has been around a while and, like the other applications, it can be run from the command line or in Windows versions.

If you want to try a command-line sample of Gopher, you need to use Telnet to log in as an anonymous user to a Gopher server:

Telnet gopher.mus.edu

Once you are connected and have the system prompt, you enter

Gopher

If you are using your browser, you can access a Gopher site by going to your browser's Open dialog box and entering.

Gopher ://gopher.msu.edu instead of http://www.something

Gopher provides an intuitive search tool that is easy to learn and use. It is especially handy for research on the Internet, and is a personal favorite over using the browser for such work. The advantage of Gopher is its small size, which allows it to run well on small and not-so-fast machines. WSGopher is about 370KB, whereas the newest browsers run from 5MB to over 10MB. Also, Gopher is quick because it does not have the overhead of the browser. WSGopher provides good font size and selection so you can print nicely formatted documents from your searches. Also, it allows you to bookmark new sites so you can build a research library quite easily. You will find selections for locating all the Gopher servers in the world in the Home Gopher directory. Gopher servers are organized geographically by region: Africa, Europe, Middle East, North America, Pacific, and South America. Under each of these choices you will find the individual countries in each region.

Veronica

Veronica is a system that queries titles in Gopher servers on the Internet. Veronica stands for "very easy rodent-oriented net-wide index to computerized archives." Veronica also includes references to resources provided by WWW servers, Usenet archives, and other Telnet-accessible information such as library card files.

You will not find a Veronica client to install on your computer. Veronica is accessed by Gopher client software. When you search a Gopher server for a keyword, it is Veronica that is doing the looking. Gopher just gives you Veronica's results. In fact, in the WSGopher Bookmarks area we discussed in the previous paragraphs, you will find the entry Search Gopher servers using Veronica. Double-click this entry, and you will be presented with a dialog box to enter the keyword you are seeking. Veronica will do the rest.

Ping

Ping is a computer equivalent of sonar. Ping clients are used to send a data packet to a server on the network to confirm that a good connection exists. By good connection, I mean the physical network and also the domain addressing or server name. A Ping client sends data packets to a host, which bounces a reply packet back. Packets are numbered so that any that do not make the round trip can be identified. The Ping client can calculate the time it takes for a round trip of the packets and count how many are missing. Using these parameters, you can check the quality of the connection.

19.4 Communication

One of the driving forces for the early evolution was the need for enhanced communications in the form of e-mail. Since those days, networked communications has grown to include other messaging systems. In this section we will look at sample o available text based communication applications. You can also find messaging applications that allow audible or voice exchanges between users of the network.

19.5 E-mail

E-mail can be sent to and received from any area of the Internet. To do so, you must have a mail gateway. E-mail is the easiest of the Internet connections to get. You can attach a file produced in your favorite word processing, graphics, sound, or spreadsheet file and it is as if you had handed it to him on a disk. It transports just like the mail, but it does it in seconds instead of days. It doesn't take long to see how the benefits of such an easy exchange can enhance your business or personal life.

There are lot of e-mail services on the Internet, and they differ in format from each other. To resolve this format issue, the Internet uses mail gateways between the major Internet service providers. The mail gateway takes in a format from one provider, coverts it to the correct format for the provider on the other side, and sends it on to the recipient.

The normal e-mail address looks like this:
username@server.(net, com, edu, org, gov, or mil)
You will sometimes see the country code on the end of the address, like this:
username@server.net.au for Australia

You can obtain information from your ISP on the format of its mail gateway service to other gateways. For instance, if you are on America Online, the following formats apply:

AOL to CompuServe: 23423, 321@cis
AOL to the Internet: jnutin@nowhere.edu
Internet to AOL :aoluser@aol.com

19.6 Internet Relay Chat

JarkkoOikarinen originally wrote Internet Relay Chat in 1988. Since then it has been used in over 60 countries around the world. It was designed as a replacement for the Talk program that was popular at that time.

IRC is a multiuser chat system. Chat is a system of messaging that lets all participants in a text-based discussion read what everyone else has typed on their machines. In chat, people meet on channels. A channel is a virtual room that has specific topic of conversation. It enables people to talk in groups or privately. There is no restriction to the number of people who can participate in a given discussion or the number of channels that can be established.

All IRC servers pass messages from user to user over the IRC network. One server can be connected to several other servers and to hundreds of clients. Several IRC networks exist, and the largest one, called Efnet (Eris Free net), can serve over 15,000 users. Smaller IRC networks like Undernet and Delnet are less busy and can be more available.

The mIRC32 client takes your input, runs it through some filtering, does some parsing, and then passes it on to the IRC server you have selected. The first time you run mIRC32, you have to fill in some information about yourself. You need to enter your real name, e-mail address, nickname, IP address, and local hostname under File / Setup / IRC Servers and Local Info, as well as the IRC server with which you want to connect.

To join conversations, send private messages, and handle and control mIRC32, you will need to learn some simple commands. In IRC, all commands start with a forward slash(/). mIRC32 assumes that text not begin with / is a message to someone and will send it to the current channel you are visiting or to the person you are chatting with a private chat. A list with the most-used commands on IRC to given in the mIRC32 help file.

To start talking, just type! When you're done saying what you have to say, press the Enter key. When you ready to leave c channel, type /part #channelname.

The channel window that opens on the right side when you join a channel will give you an alphabetical of people who are currently on the channel. The names with the @ in front of the name are the channel operators. Some regular channels have established rules of etiquette. The first person on a channel becomes the channel operator by default, and he or she can kick you off of the channel if you do not adhere to there rules.

Try to connect to a geographically close server because it takes time for the messages to migrate across networks. You can always ask for suggestions when you log in to an IRC channel and you can leave whenever you wish. Apply the same courteous etiquette you would use if you had walked up to a small group of people having a discussion and you should have no problem.

Usenet

The Usenet news system is the society page of the Internet. Usenet is thousands of electronic bulletin boards, each oriented to specific area or topic of interest, where users can post messages or respond to other postings. You can subscribe to specific groups that interest you so that your news program will retrieve all of the articles processed by those groups. Usenet news is accessed by logging in to your service provider's news or NNTP (Network News Transfer Protocol) server.

There are several newsreaders to choose from, but here I will discus the Free Agent newsreader by Forte. Free Agent is especially designed for offline reading to help reduce connection costs. It enables you to retrieve all the new headers in selected groups since the last time you connected. Once you retrieve the new headers, you can go offline and mark what you want to read. You then log back on to retrieve the corresponding bodies, and then log off again to read them. If you have an unlimited time Internet account, it is not necessary to log off and on; you can just use the program online to browse the newsgroups.

Free Agent is free to individuals for personal use at home and to students or staff of educational institutions and nonprofit organizations.

Free Agent has a lot of options, and the best thing to do is to just plunder in the toolbar and the drop-down menus. There is also a mass of information available by selecting Help on the menu of Free Agent.

19.7 World Wide Web

The World Wide Web is the newest and fastest-growing area of the Net. The WWW, as it is called, was started by Tim Berners-Lee while he was at CERN (the European Laboratory for Particle Physics). The WWW was created for internal use by CERN to allow researchers around the world to collaborate on research issues.

The WWW is much the same as Gopher in its setup, but it employs HTML, or Hypertext Markup Language. In this language, the action of clicking specific words takes you to another link on the Web. The other link might be a picture, a document, or even a sound file. Think of it as a huge relational database where information that has relativity is linked.

HTML allows such features as clicking a picture of a public figure and hearing that person speak. With today’s extensions to the original HTML, you can see a video of that person speaking as well. This is the type of thing you encounter on the WWW. It is set up on a system of home pages, which are nothing more than sites on the Net, similar to the root c:\directory being sort of a home directory on your computer.

19.8 Universal Resource Locators and Addressing

There is not much you can say about the Universal Resource Locator (URL) addressing scheme. That's good; it was supposed to be simple, and it is. In the URL you will find the type of protocol, the address of the site where the resource is located, the subdirectory location, and the name of the file if it is needed. Some URL also have the port number appended.

A typical URL for the World Wide Web looks like this:

http://www.ncsu.uiuc.edu/electric/software/loadflow.exe

This code line represents these components for a Hypertext Transfer Protocol (HTTP) Web browser:

http – The protocol that the browser is using to reach the information.
:// - Required by the URL specification.

www.ncsu.uiuc.edu - The domain name of the organization where the server is located.

/electric/software/ - The directory and subdirectory where the information is located on the server.

Loadflow.exe – The file that the browser is trying to retrieve from the server.

If the request requires a port number, it looks like this:

http://www.ncsu.uiuc.edu:80/electric/software/loadflow.exe

If you are using the Gopher program I discussed earlier, the Gopher server has addressing that looks like this:

Gopher.ncsu.uiuc.edu

When you want to reach a Gopher server with a Web browser, you type it in like this:
Gopher://gopher.ncsu.uiuc.edu

If you want to Telnet to a site with a browser, you type in this:
Telnet://ncsu.uiuc.edu

Finally, Usenet News can be reached with a URL that looks like this:
News:alt.business

Note that the double right slash is not used for this command.

19.9 The Browser

Hypertext, hypermedia, hyperlinks, hypertension…sorry, that last one does not fit. You can see www.something on anything from soup cans to funeral parlors. The World Wide Web has struck the nerve of society at its most critical point, and the reaction has been not less than explosive.

There are a lot of rich folks who bought browser stock early in the game. Browsing will probably replace the term “channel surfing” on television, especially when you can buy a television set that lets you browse the World Wide Web with a remote-control unit. Yes, it’s already here, so get out that loose change and get ready to purchase.

The World Wide Web is both a distributed hypertext and distributed hypermedia system. That means a document link can point to another document anywhere in the world of networks. It also means that the document that comes in response to your browser’s request can contain text, picture, music, voice, video, or smut (oops, but there are ways to block access to that last part from the kids).

19.10 Short Summary

File Transfer Protocol (FTP) is one of the all-time greats of networking software.

The World Wide Web is the newest and fastest-growing area of the Net.

A channel is a virtual room that has specific topic of conversation. It enables people to talk in groups or privately. There is no restriction to the number of people who can participate in a given discussion or the number of channels that can be established.

Usenet news is accessed by logging into your service provider’s news or NNTP server.

19.11 Brain Storm

1. Explain about application of Internet?
2. Explain the tools and utilities of Internet?
3. Explain about email?
4. What is WWW?
5. Explain the uses of browser?
6. What is URL?

Lecture 20

Internet Vs Intranet

Objectives

After completing this lesson, you should be able to do the following

- Discuss about Intranet versus Internet.
- Describe why we need intranet.
- Describe about intranet increase communication.
- Describe about user friendly.
- Describe the benefits of intranet.

Coverage Plan

Lecture 20

20.1 Snap Shot

In this lecture unit we can able to learn about intranet rationale. The fresh new approach to all this electronic madness is the intranet. An intranet is an internal, network-based infrastructure. In general the intranet is not made available to reviewers outside the security firewall that connects the Internet.

20.2 Intranet Versus Internet

The Internet has lit the fire of ingenuity in the business community. Developed for academic, government, and commercial researchers, the Internet has grown into a worldwide network with millions of participating internauts.

What has made the Internet so popular is standardization. Before the Windows graphical user interface(GUI) came into play, Internet access and use was confined to the masochistic few. Oh sure, the data was out there. The academic community had been filling servers with wonderful knowledge for years. But black screens and cryptic syntax did little to excite the masses and beckon them to participate. The standardization of the Internet and the World Wide Web made it possible for the rest of us to point-and-click our way across the nation and around the world. With the GUI in place, the knowledge can now be tapped population has made the discovery.

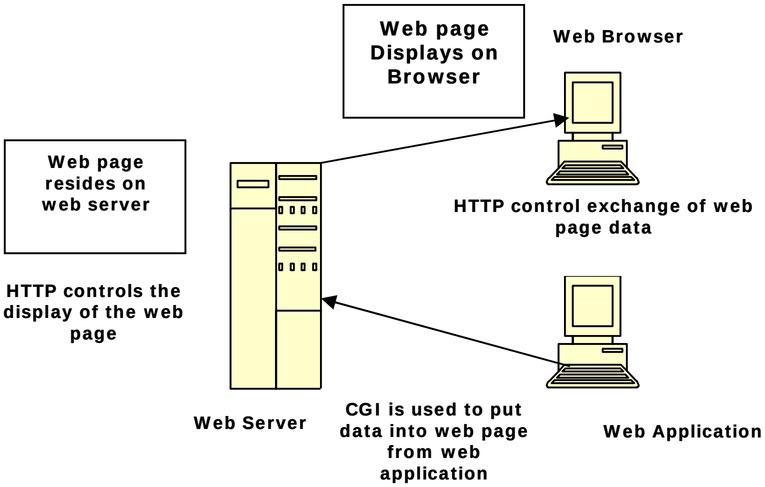
Web page Displays on Browser
Web Browser
Web page resides on web server
HTTP control exchange of web page data
HTTP controls the display of the web page
Web Server
CGI is used to put data into web page from web application
Web Application

The primary distinction between the Internet and the intranet is purpose and use, not technology. In fact, the technology is absolutely identical for both entities. To separate the two requires a security function like a firewall. This prevents outsiders from using inside intranet information yet allows insiders to use outside Internet information.

The Internet is a distributed cost environment in that no single entity owns the whole thing. The basic design philosophy of the Internet involves minimizing cost and maximizing information flow. An intranet is usually a wholly owned infrastructure – a network that has been put in place over time to handle a business environment. This infrastructure carries a high embedded cost per functional element that it supports. For instance, consider the cost of connecting a remote workstation to a centrally located mainframe for the conventional client/server application. This cost of equipment, network cable, expensive software, labor to install, and ongoing support is substantial.

Standards are the key. At the Internet level, it is the standards that make it possible for so many different computer systems to exchange information. The global acceptance and adoption of standard protocols have resulted in a profusion of easy-to-use tools and applications for the Web environment. Many tools are available freely on the Internet and perform as well as their expensive commercial counterparts. The bottom line is, if you can do it on the Internet, you can do it plus more on the intranet.

Are the Internet and intranet in competition? Yes, in a way. While they essentially complement each other, the intranet has caught the attention of business and industry. Money talks, and the market is currently being driven by the profusion of new intranet applications. While it is the Internet that has been in the driving seat for over a quarter of a century, it is now the intranet that has come of age and is taking the wheel.

20.3 Do you need to Intranet?

There is currently a shift in the control structure of most savvy companies. Progressive companies are releasing authority and responsibility to empower employees. In response, employees, who greatly outnumber the management structure, begin to develop ideas in a parallel manner. The empowered workplace becomes like a snowball that starts out the size of a baseball and grows rapidly as it rolls along. Effectively, the company expands its momentum by releasing control to numbers of individuals. With this comes the need for increased coordination of activities and monitoring or progress. The need to collaborate without the normal delays of document preparation, printing, and distribution becomes critical. The successful company will find methods to rapidly disseminate information among its people.

When documentation is relatively static, there is no problem with using the paper methods. For instance, your company phonebook might be published by a central organization like your human resources department. In past years, an annual or semiannual update and reprinting might have been sufficient. No doubt, though, days or weeks after printing you had names scratched out and new phone numbers added. Also consider that in the past employees generally made a company their home for life. This is not so in the current employment picture. People come and go, and your phonebook grows ragged with erasures and new entries. You are glad to get a fresh printing, and then it starts all over again.

Suppose your company has multiple sites and people who need timely access to human resources documents, policies, desktop documents such as phonebooks, specifications, and product data.

20.4 Intranets increase Communication

You cannot eliminate essential documentation and communication. The empowered workforce moves quickly. Increased communication is absolutely essential for the technology within companies. In today's competitive business arena, access to information is a critical element of success.

The use of technology to improve communication in the business environment has been met with various degrees of success. For example, there is e-mail that results in the unnecessary stuffing of employees' mailboxes, and client/server databases that require teams of support personnel. What is needed is an open environment where information and not just data is available on demand from the employees' end, not dumped down the pipe from management's end.

The open type of information environment is being provided by the conversion of common networks to intranets in businesses of all sizes. Intranets provide a highly effective, timely, and extensible communications platform. Using the technology developed for the Internet and the World Wide Web, new applications can be prototyped in days where it requires months in the distributed client/server environment develop – and without a forklift upgrade of network architecture.

In this manner, creating and utilizing an intranet via Internet methods is a scaleable architecture. In addition, your existing client platforms are probably Web compatible, because there is a Web browser implementation for all major client architectures. Your information will also be distributed by the organizations that are most familiar with the content and that can keep it more accurate in less time. There is no requirement to replace existing documentation and database architectures; existing content can be accessed directly or interpreted into the Web-oriented environment.

Intranets Enable the Mobile User

In the client/server company, a mobile user is hard pressed to establish a connection with all the data he might need while on the highway. Competition is creating the absolute requirement to cut the cost of doing business. One very effective way to do this is to arm your mobile force with laptop computers. Not just any laptops with a dull gray screen, but machines that rival the best of the desktop variety – laptops that have full multimedia capability to present your firm's products in interactive, full-motion video with stereo sound.

For the mobile user, connectivity to the office is essential for survival out there on the back roads. Using the Web techniques of the intranet via remote TCP/IP, the road warrior can reach any information anywhere on the company intranet. The same Web-based multimedia information he would use at his desktop is available in a hotel room where he can customize his presentations at the last minute for each customer site he visits. Thus the intranet puts a new slant on mobility.

20.5 Intranets are User Friendly

Your users are probably already aware of the Internet's user-friendly. HTML-based Web screens. This same GUI is the basis for intranet design, bringing with it navigation at the click of a home page link. Instead of paper that often is too late to suit its purpose, your users can view information with audio, video, or graphical content.

Every day, more of your prospective intranet users are either connecting to the Internet on a personal basis or have access to it in their workplace. Everyone who comes in contact with the

Internet receives a basic training course for the Web-based intranet environment. You essentially have in place a large part of what otherwise would be a training expense. The Web browser is a universal communications medium.

An intranet will allow your organization to publish critical information as it needs to be revised. There will be no need to accumulate changes to a volume until the point when printing a new release is cost-effective. Printing anything is not cost-effective these days, anyway; information changes too rapidly to commit it to paper and expect it to be relevant for more than a few weeks or months.

With intranet publishing, information becomes communication. What were words on paper can become graphical or audio content to make it appealing and understandable. Furthermore, it becomes communication because the information content can be updated and is instantly available to all users simultaneously. For instance, when a new benefits program is announced, the content on the Benefits page can define the new change. The company's internal home page can have a What's New icon pointing employees to the Benefits page, so employees have the new information at their fingertips. The content can be changed or updated to reflect new information at any time.

Intranets Reduce Publishing Cost

An intranet in your organization can reduce the costs of content development. Preparation of the materials and distribution costs are virtually eliminated. The publication process of today begins with creation of the desired content, iterating through several drafts and reviews, then moving the paper medium from the printer to the employee. The intranet publishing process is essentially two steps: creation of the content and migration of the content to the intranet technology. Even the review of the content can be performed with a browser and, in fact, reviews should be done online. That way, the content can be reviewed in its native environment, just as the employees will view it when completed.

Information in an intranet environment is kept in key database machines. This means that there is a single, master copy of the information (of course, a backup copy of the online data is kept offline for disaster recovery).

This intranet publishing methodology reduces the costs and the time involved. Most LAN environments can support Internet access and Web technology. Most intranet Web servers are designed to run on commonly available personal computing platforms. For an organization of several hundred employees, even the common Intel 486-class personal computer can handle the traffic. The key is simultaneous connections when considering whether you will need to include more computer power in your intranet decision.

20.6 The Benefits of the Intranet

Because of HTML, the Internet and now the intranet are fast becoming the new wave technologies of the 20th century. The volume of Web traffic exceeds all other traffic on the Internet and the increasing use of graphics, audio, video, and other data types on Web servers will continue to drive network growth. All of this is the result of the ease that the GUI and HTML provide.

The increased traffic from HTML-enabled personal computers has expanded the information resources. Consecutively, the expansion of information resources has furthered the development and deployment of browsers over many computing platforms.

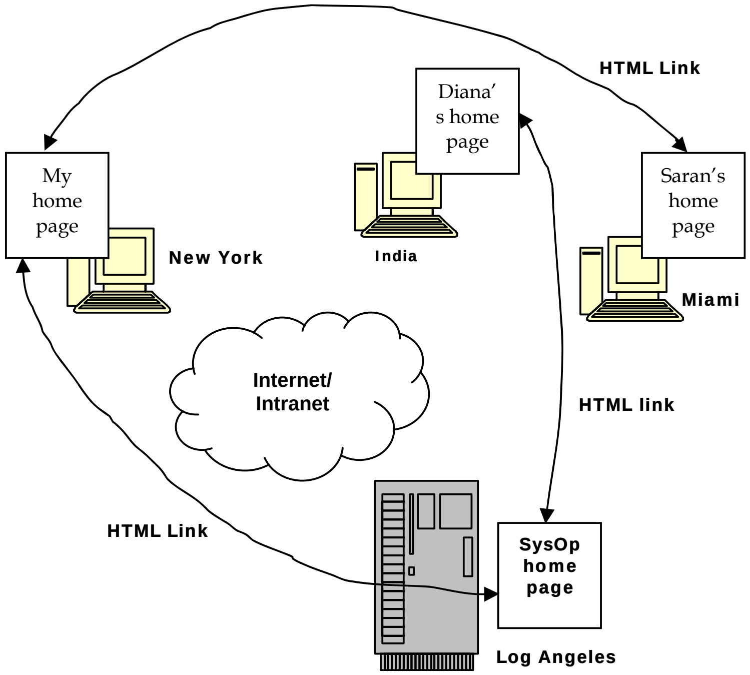

Just like the transition from the magnetic card typewriter, today’s office employee is becoming familiar with retrieving information through his Web browser. The Web browser is becoming the universal interface to all information types. It doesn’t matter what the locations of these information types are. A browser on a personal computer can find any InterNIC – registered domain in the connected world.

User Interface Is a Benefit

With the basic GUI that a Windows product provides, there is no way to apply the imagination and creativity of your employees. Windows gives you exactly what it was designed to do, standard ways to do standard things. No matter what the application is, you should be able to click essentially the same icons to achieve the same results in a windowed operating system.

With Web-based technology, you are unlimited in your imagination with Web page development. Of course, this also can allow for some creativity that is a little too extreme. With HTML you can build an interface with any functional product that supports TCP/IP. The interface can have any type of humanly recognizable feature. The intranet technology is extremely simple to apply. Hyperlinks in HTML can start applications on remote machines and return results to you or anyone else in the world that is connected to the network.

Because it is simple, you eliminate much of the support issues and cost to maintain specialized and dedicated functionality. Web page HTML code is easily readable without extensive education requirements. It does not require a FORTRAN or C programmer to write HTML.

Intranet technologies provide the standards. The standards are used to develop new approaches for meeting the problems of today’s businesses. Communication is the key to business success. Development of more intranet-based business tools is the key to extracting the total benefit from this exciting new life, giving application for the aging client/server mechanism.

20.7 Short Summary

Web based intranets become essential tools for your users when the content of the intranet goes beyond simple text. Users will expect a tool that gives them flexible management and retrieval of company-wide information. It is hard to imagine that we could need more functionality than is provided by today’s HTML standard and the infrastructure that it supports.

Internet technologies in an intranet setting can dramatically increase the speed of communication and the quality of information in the modern, progressive organization.

The intranet philosophy is gaining momentum rapidly, but we don’t have to leap in above our heads with a labor-intensive intranet installation.

20.8 Brain Storm

1. Distinguish the concepts of Internet and intranet?
2. Explain the benefits of intranet?
3. What is intranet?
4. Why we need intranet explain.

Lecture 21

Intranet Application

Objectives

After completing this lesson, you should be able to do the following

- Discuss about Intranet and Internet web site development.
- Describe about the application of intranet.
- Describe about organizational web pages.
- Describe the distributed strategy.
- Describe about the web master.

Coverage Plan

Lecture 21

21.1 Snap Shot

In this lecture we can learn about the development of Internet and intranet technologies. The main difference between the two is the scope of the network. While an intranet web site is accessible to only certain users in one or several organizations, an Internet web site accessible to millions of users all over the Internet.

21.2 Internet & Intranet Web Site Development

The anatomy of an Internet web site is very similar to that of an intranet web site. The main difference between the two is the scope of the network. While an intranet web site is accessible to only certain users in one or several organizations, an Internet web site is accessible to millions of users all over the Internet.

Internet Web-Site Development

Internet Web-site development is a constant balancing act between developing compelling and richly interactive content and addressing compatibility issues. There are two solutions to this problem – each with its own set of advantages and drawbacks.

- ❖ Standardize on HTML 2.0 – standardizing on HTML 2.0 ensures that virtually all Web browsers display the contents of your Web pages more or less the same way. However, standardizing on HTML 2.0 limits how well information can be presented to users in a visually appealing manner. Even the most conservative Web pages have begun to go beyond HTML 2.0 and use HTML attributes such as tables and background colors to present information in a visually appealing manner.

- ❖ Standardize on HTML 3.2 – the formatting and layout capabilities of HTML 2.0 are very limited. HTML 3.2 offers more powerful Web-page design and layout capabilities for creating visually engaging Web pages. Netscape Navigator and Microsoft Internet Explorer both support HTML 3.2 and account for over 95 percent of all Web browsers in use today. Although standardizing on anything higher than HTML 2.0 used to make your Web site effectively inaccessible to a large percentage of users browsing the Web, the situation is rapidly changing – especially with online services such as America Online and CompuServe adopting Microsoft Internet Explorer or Netscape Navigator as the default Web browser.

Let’s explore how bandwidth availability, platform compatibility, scope of audience, and security affect Internet Web-site development.

- ❖ Bandwidth availability – One major difference between Internet and intranet Web sites is the bandwidth available to users browsing your Web site. Most Internet users are connected to the Internet via relatively slow POTS (Plain Old Telephone Service) modem links. Information published on an Internet Web site should be optimized for transmittal over low bandwidth Internet connections.

- ❖ Server latency – Server latency must be addressed when deploying a Web site on the Internet. Web servers are no longer used exclusively to publish static content on the Internet. Increasingly, Web servers are using server side applications to create dynamic content. Although a 486DX2/66 computer can saturate a T1 connection with static content, the situation changes dramatically when the same computer has to process database queries and execute server-side applications to fulfill HTTP requests. Use Windows NT Performance Monitor to detect bottlenecks and other server latency issues. Although the performance of your server can be enhanced by adding more RAM, if CPU usage is high, you will have to upgrade your server to a multiprocessor server.

- ❖ Scope of network take into account the scope of your network when setting up your web server. Confidential information that should remain only within you organization should not be published on a web server that is accessible via the internet.

- ❖ Platform compatibility the internet consists of a wide variety of hardware platform and operating systems. When publishing information, platform compatibility should be taken into account to ensure that information published at your Web site is accessible to a wide variety of users. When users need to use a special helper applications that run on several platforms. At a minimum, Windows and Macintosh users should be able to view information published at your Web site, as should those who use widely used flavors of UNIX.

- ❖ Security – While an intranet Web site is accessible primarily to select individuals, an Internet Web site is accessible to millions of users all over the world. Never use clear-text passwords to protect sensitive information distributed to and from an Internet Web site.

When distributing sensitive information via the Internet, configure your Web server to encrypt the data.

- ❖ Scope of audience – The Internet consists of a very diverse group of users. When developing an Internet Web site, consider the scope of your audience. For example, if the targeted audience. For example, if the targeted audience of your Web site is not very technically inclined, do not assume your users use the latest version of Internet Explorer or Netscape Navigator, and don't assume they have helper applications installed on their systems.

Intranet Web-Site Development

Intranet Web sites are increasingly becoming the backbones of the information infrastructures within organizations. Prior to intranets, information in organizations was usually scattered across countless servers with obscure server names, user names, and passwords. To make matters more confusing, there was no uniform medium or format to access and view the information. Intranets solved this information-distribution problem by providing a cross-platform medium, Hypertext Transport Protocol (HTTP), and document format, Hypertext Markup Language (HTML).

- ∗ Bandwidth availability – Although Internet users are usually plagued with bandwidth-limitation issues, the same if not true for intranet users. Intranets are connected using high-bandwidth network connections (compared to the relatively low-bandwidth modem connections used by millions of users to access the global Internet). You can take advantage of high bandwidth network connections when developing an intranet Web site by offering multimedia-rich Web pages that are visually appealing and exciting to navigate.

- ∗ Server latency – Intranet Web servers are often bandwidth intensive they are used to publish constantly changing information such as sales reports and forecasts. Be sure your Web server is capable of serving HTTP requests in a timely manner. If your Web server is too slow, users may become frustrated and hit the stop button and resubmit anew HTTP request, thus compounding the situation. Use the Performance Monitor to monitor the performance of your Web server and detect server-latency problems.

- Web-browser compatibility – When developing content for an intranet Web site, you generally don't have to be concerned about Web-browser compatibility issues because most organizations standardize are Windows desktops, Internet Explorer can be used as the standard Web browser. The same is not true for an Internet Web site. While one user navigates the Web with the latest version of Internet Explorer, another might navigate your Web site with a less powerful Web browser (such as Mosaic).

- Platform compatibility – Microsoft and Netscape Internet client applications are available for Windows and Macintosh platforms. If you intend to use cutting-edge Web-publishing technologies such as Active Documents, your users should use either Windows 95 or Windows NT Workstation 4.0. Although Microsoft intends to release a version of Internet Explorer for Windows 3.x, it will lack some features of the 32-bit version of Internet Explorer 3.0. In case your users are still using Windows 3x, I strongly recommend that you upgrade them to Windows NT or Windows 95. The enhanced Internet connectivity features and stability offered by Windows NT and Windows 95 are well worth the cost of upgrading. If RAM prices have been holding back the deployment of Windows NT, you'll be happy to know they have recently come down significantly. Although new technologies are usually first implemented on Windows NT 4.0 and Windows 95, Microsoft is now aggressively developing Internet client applications for Macintosh users. By the time you read this, more Internet client applications will be available for the Macintosh platform.

- Security – Although system administrators have more control of an intranet Web site , even in small organizations certain information must be kept confidential and out of reach of those who are not authorized to have access to the information. Use Web-server encryption to ensure that sensitive information distributed via your intranet Web site does not fall into the wrong hands. If encryption is not used, a user with a protocol analyzer can intercept sensitive information, as well as user names and passwords distributed to and from your intranet Web server.

- Scope of audience – The audience of an intranet Web site is smaller than that of an Internet Web site. However, the intranet Web-site audience is very specialized. Intranet users navigate your Web site to find information that will help them get their work done more quickly and more efficiently. An intranet Web site should provide comprehensive and up-to-date information that helps users make informed and timely decisions. A good way to begin providing such information is by electronically publishing all the informative brochures and newsletters that are currently distributed in printed form in your organization.

21.3 Intranet Applications

Internet methods were put to work eliminating the paper machine in early implementations. Today, there is more of a trend to implement information sets that support the core goals or business direction of the company. Some that are applicable to all business types are human resources, company news, safety information, and even the weather forecast. Sales-oriented industries might include product information about their own products and that of their competitors. Financial information like 401(k) plan balances can be securely viewed by employees. You can include the president's wisdom of the day, which will surely make your intranet the center of your employees' interest.

The intranet generally begins with a corporate home page as a starting point. From there, *hot links* (pointers to other pages) take the user to the selection of his choice. Links such as financial communication, marketing data, manufacturing status, and human resources are common. A soft tone can be applied by including non-business announcements like outings, weddings, great vacation experiences, or births.

Sales Applications

Contract awards are favorite announcements for sales-related organizations. In the competitive environment of today's business, it is critical to inform your sales organization about the competition. The intranet Web page is ideal for rapidly changing communication about the opposition and their products, and how they compare with your own.

One method of sales in today's business is telemarketing. This is where dozens of employees use computerized mailing lists to call prospective customers and offer everything from magazine subscriptions to clothing. Your telemarketing organization might be only an order desk that takes orders from responses to television commercial or marketing channels. With an intranet in place, the telemarketing organization's employees can have instant access to the details about items you have for sale. They can have video or high-quality graphics to give them the information they need to answer the customers' questions. The order information can be input in the same screen that the product is displayed on using the browser's forms technology.

Human Resources Applications

Human resources departments in the corporate environment are famous for generating volumes of fine-print documents about policy, benefits, retirement, and more. Many human resources applications are available on the market today that use the intranet technology for communications with employees. Companies have converted their existing employee manuals to their intranets and have added additional functionality. There is the initial and obvious advantage of the point-and-click viewing of the company human resources documentation.

Generally, as a second or later phase, the company integrates the intranet with the legacy databases so employees can perform database queries. An employee can use his Web browser to securely log into his financial packages and instantly see his retirement or similar accounts. More recent advancements in this area allow an employee to interactively change items like his percentage of participation in a 401 (k) program or select different options in a cafeteria-type health program.

Technical Support / Help Desk

Intranets are employed in service organizations to provide faster technical information to internal users. The newest functionality in the lingua franca of the Web, Hypertext Markup Language (HTML), provides many techniques for forms-based information transfer. Forms are used for requisitioning materials and supplies, reporting problems, and obtaining software updates and new downloads. Considering the individual's input, servers with work order applications alert the proper response organizations.

For the service organization, a valuable benefit is the knowledge base. A knowledge base is a database that contains help requests to resolve problems as well as the resolution of those problems. Users can use a Web browser to submit a query on a particular problem. The return information consists of the same instances of these problems and their resolution. A knowledge base grows in value to an organization as more users feed it with methods of resolution.

Like the telemarketing application, the intranet in the service organization can give the telephone help disk visual aids. For instance, the help desk employee can request a video segment of an appliance showing how to use one of its attachments. The employee can discuss the use of the appliance with the customer as the video runs its course.

21.4 Organizational Web Pages

In the empowered business, the current high-speed culture change is breaking down organizational structure. Each internal organization is separating from the culture to pursue its own individual mission. In support of this, the Internet technology provides each organization with the opportunity to have its own hot link from the main corporation's home page. This results in the ideal medium to communicate current information at the level of the individual. Web search functions give easy access to information and answers to questions in support of the daily business environment.

Employee Feedback Applications

The hypertext markup language provides forms support. With a browser, you can access a Web page with full-in-the-blanks features. The possibilities of use are endless. For instance, your marketing department might design a survey for a product and test it on your employees. Putting the survey on the intranet home page makes it instantly accessible to every employee who has access to a corporate machine.

Secure access to employee personal information like material status, personal history, and home address can be input directly by the employee, eliminating the normal paper forms and delays. When the personal data needs to be changed, it is the employee who accesses his own data and makes the revision. The human resources personnel are not involved beyond receiving something like an e-mail that a record has been modified and by whom.

Retirement investment information is an excellent application for employee self management via the intranet. Current retirement options can be viewed on the intranet, and deduction selections can be changed at the click of a mouse. Stock exchange and mutual fund data can be extracted from the Internet, and fed to the financial update page on the intranet. This can give the employees a retirement fund management function wherein they can modify their purchases and conversions of financial vehicles to suit their current needs.

Software Distribution

In the traditional client/server installation, applications are installed by hand at each user’s location. With the intranet technology, software can be delivered to users with sufficient instructions to let most install it themselves. Like the videocassette recorder and the hand-held calculator before it, the personal computer is becoming a consumer device. New Age users are becoming to installing their own software.

New technologies such as Java will allow the creation and transparent distribution of objects when requested by the browser without user intervention. Java places functional code on the browser end of the browser-to-Web-server connection. The result is an increase in intelligence in the browser and a possible decrease of intelligence on the Web server end.

Mail and More

With the emergence of Web technology, applications such as Microsoft Exchange push the envelope with their functionality. In the past, the transfer of simple text was the extent of the e-mail facility. Modern intranet mail products provide standard, simple techniques for attaching and sending document files, sound, and video between individuals.

With the mainframe environment and then with traditional client/server setups, the concept of groupware became popular. Today, the Web allows collaboration and groupwide scheduling functions. Employees can fire up a whiteboard function like Microsoft Net Meeting and have a conference online. Scheduling is a snap on the intranet with functions like Schedule + by Microsoft. If you wanted to schedule a meeting, you just put in the attendees and let the system handle the coordination. When it finds a slot, you ask it to notify the attendees via e-mail. Each can respond that he will or will not make the conference.

21.5 Intranets: A Distributed Strategy

Intranet applications are a distributed computing strategy. In this strategy, the server and content are located closer to the owners of the content.

Intranet servers may be located within a group or department-level organization, shortening or eliminating delays for posting content. This computing strategy is one of the key elements of keeping information current. It allows the owner of the information to both develop and maintain the content. When the user points his browser to the company’s home page on the intranet, he will see links to pertinent information but will have no indication of where the data is stored. With hypertext links, he can extract a page of information, and it will appear as if it were coming from the hard disk on his own machine.

The Costs of an Intranet

You see that Microsoft has bundled the basic and necessary components for an intranet implementation with its Windows NT Server 4.0 operating system. Your first thought is that establishing an intranet will be a snap. You buy the operating system, and you fire up the Internet Information Server, which also come with File Transfer Protocol (FTP) and Gopher clients. You can get all the copies you need of Internet Explore from the Microsoft Web site for your clients. This is going to be cheap, you think. Well, relatively inexpensive might be a better way to put that. Certainly, it will not be free.

How We Got Here

Compared to the mainframe client/server solution, the Web server/Web browser solution is significantly lower in cost. First, let’s take a look at our corporate computing roots. Below figure shows a traditional mainframe-computing environment that has existed for decades and is currently the most predominant architecture of large businesses.

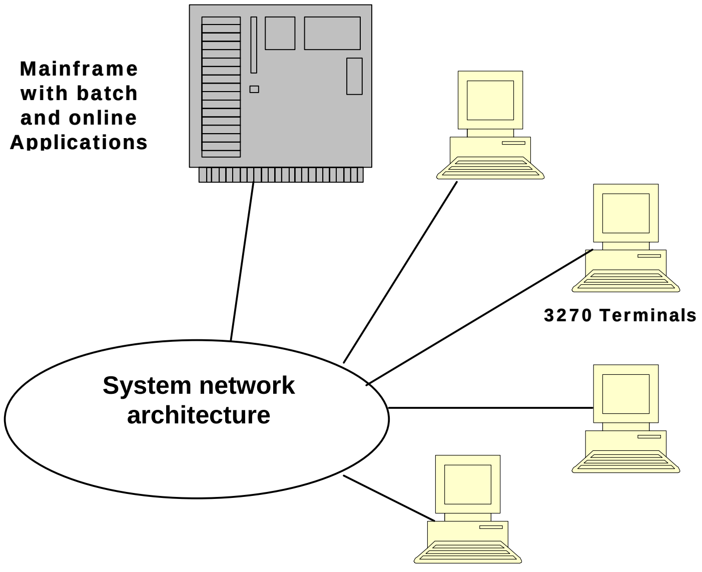
Mainframe with batch and online Applications
3270 Terminals
System network architecture

The centralized mainframe is the classic model for the large-business computing environment. The cost of this architecture caused the centralization of applications; application support in the mainframe is a costly venture. Once the infrastructure is cast in stone, a dramatic shift in culture is required to break down the model.

When the first IBM personal computer was developed, the client/server model was born. Network infrastructure began to replace the dedicated mainframe to 3270 terminal environments, and personal computers came into play. First the personal computers played the role of multi-application 3270 terminals. This was a significant advancement in technology because a single machine could do different tasks. Word processing, spreadsheets, and 3270 terminal applications became possible, although not all at the same time.

The development of more powerful and smaller machines brought about the local area network (LAN) and centralized machines called servers. This has led to moving applications from large mainframe computers to dedicated server computers located throughout an organization. Depending on the organization's size, it may operate with one or several servers, and each can share information with the others. Below Figure shows the client/server environment.

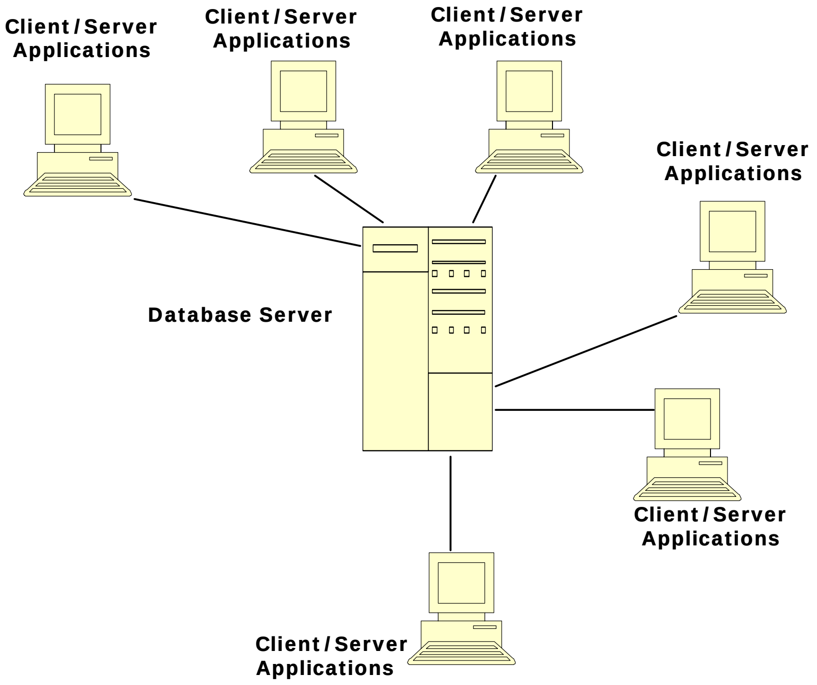
Client / Server Applications
Client / Server Applications
Client / Server Applications
Client / Server Applications
Database Server
Client / Server Applications
Client / Server Applications

The interesting fact about the shift from the mainframe to client/server computing is the shift of cost from the capital-intensive mainframe purchase to the labor intensive client/sever support. As client/server installations increase in complexity, the cost to maintain them also increases.

The current Internet model depends on the GUI and the infrastructure of the network. Remember, an intranet is not just a network; it is a philosophy for information flow. This appears to be a return to the centralized computing environment of the past. The real advantage is the reduction in support labor for Web-based functionality. Web browsers are simplistic and use a relatively simple concept for moving information. The use of the Internet TCP/IP provides the roadbed for our packet traffic.

The Infrastructure Cost

When you consider the cost of the basic business network, you will spend about $500 per user on an office suite like Microsoft Office. There will be labor involved to set up each workstation on your network, so add in about $150 per seat for this. A reasonably swift, leased client computer will run about @150 per month.

Somebody has to run it all, so a network administrator will cost about $50,000 to $70,000 per year, depending on location. The users will clamor for applications support, so add another body at $60,000 per year.

A 100-user LAN would cost about $98,000 up front for the physical infrastructure, plus $15,000 per month for the leased machines and $11,000 per month for the support labor. It sort of adds up, doesn't it? In reality, though, this is not a huge expense for a 100-employee organization.

The Intranet Add-on Expense

Now we're ready to consider the intranet cost on top of this supposedly already-in-place infrastructure.

An intranet feeds and nourishes itself on content. It is a network-based publishing system. It can just be text-based documentation, but it can also be video and audio publishing. All of the publishing methods mean one thing: lots of labor. You already have tons of documents that would fit nicely on an Internet server. They may be able to reside just where they are and you can simply point to them with Web page hot links. In any case, it takes that magic labor to hook it all together.

You will need to convert documents to HTML if you want them to retain formatting and graphical content. If not, you can use a tool like the Gopher server in Windows NT 4.0 and display your documents as plain text. This will cut down on a significant amount of labor, although it will not be as pretty.

You will want your intranet to have a common look and feel across the various departments and organizations in your company. The coordination of this is no trivial task, and that means more labor.

After you have your wonderful new connectivity in operation, you will find that it needs periodic loving care. Your documents will need to be indexed periodically so your search tools can locate information when it is needed. If you cannot reach the information you need when you need it, you will not have accomplished anything with the intranet implementation. Again, indexing files implies more labor.

When you add the intranet to your network, you will be employing applications that can possibly bring the need to upgrade your network infrastructure. This means the users may demand and run applications that are network intensive. Applications like video consume much more bandwidth than do simple Web HTML screens. If your Web pages are graphics intensive, you will also see increased traffic. If there is too much traffic at the same time, everybody's Web pages will begin to refresh their screens slowly. It can get so bad that you may not be able to access the Web server at all. Your users will quickly let you know that this is unacceptable.

You will want to monitor your network traffic and predict when your company's growth will create the need to upgrade your network infrastructure. You may not need to increase bandwidth across the board, but instead improve areas where there is more traffic because of the business need. Speaking of business need, you might consider developing a charge-back system so each department or organization in your company carries its proportional share of the infrastructure cost. Organizations that use graphics-intensive applications might be expected to cough up more financing than organizations that reference text-only Web pages.

21.6 The Web Master

The Webmaster could be though of as the director of communications on an intranet. This is the person or persons directly responsible for the content of your intranet pages and the one who makes the links function again when they break.
Intranets are to distribute information and, with some creativity, they can start processes and automate manual tasks. Your intranet will be used by your entire organization, and that implies a uniform look and feel. By that I mean a consistent method of maneuvering from screen to screen and link to link. If you use one method on one screen and another on the next screen, your user population will be confused and your phone will surely ring. Look-and-feel standards need to be designed and supported by your Webmaster and by your management. All of this needs to be done without excluding your users. It is a good idea to employ focus groups to define the look and feel. That way the employees can feel some connection with the GUI they will use every day.

What you want to get out of your intranet will dictate the levels of Webmaster support you put into it. If you just want to be with the current technology and have an internal Web page with some good-to-know company information, perhaps one of your network administrators can handle the job.

On the other end of the spectrum is the need to support dozens of Web servers. You obviously will need several Webmasters in this situation. This is a specialized technology individual. He will need familiarity with TCP/IP and other Internet protocols like FTP, SMTP, and HTTP. Prior experience with Web servers is recommended, and you should view some of a prospective Webmaster's work on the Internet to gauge his quality of delivery. The personality of a Webmaster is different than that of the network administrator. The Webmaster is usually a people person who is required to work closely with the users, whereas the network administrator is traditionally concerned with machines and network performance.

21.7 Short Summary

Before developing an Internet or intranet Web site, it is worthwhile to explore the differences between the two, as well as the issues related to developing Internet and intranet Web sites. The next two sections, “Internet Web-Site Development” and “Intranet Web-Site Development,” discuss Internet and intranet Web development as they relate to the following topics:

- ❖ Bandwidth availability
- ❖ Server latency
- ❖ Web-browser compatibility
- ❖ Platform compatibility
- ❖ Scope of network
- ❖ Security
- ❖ Scope of audience

21.8 Brain Storm

1. What is web master?
2. Explain the distributed strategy of intranets.
3. How to develop web site explain.
4. Explain the application of intranets.

Lecture 22

Microsoft Internet Tools

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Microsoft Internet information server.
- Describe the Microsoft Internet assistants.
- Describe about Intranet search engine.
- Describe about Alta vista and excite.
- Describe the Intranet toolbox.
- Discuss about platform independent pages.

Coverage Plan

Lecture 22

22.1 Snap Shot

Microsoft® Internet Information Server is a Web server that enables you to publish information on a corporate intranet or on the Internet. Internet Information Server transmits information by using the Hypertext Transfer Protocol (HTTP). Internet Information Server can also be configured to provide File Transfer Protocol (FTP) and gopher services. The FTP service enables users to transfer files to and from your Web site. The gopher service uses a menu-driven protocol for locating documents. The gopher protocol has been largely superseded by the HTTP protocol.

22.2 Internet Information Server

The Web is fundamentally a system of requests and responses. Web browsers request information by sending a URL to a Web server. The Web server responds by returning a Hypertext Markup Language (HTML) page.

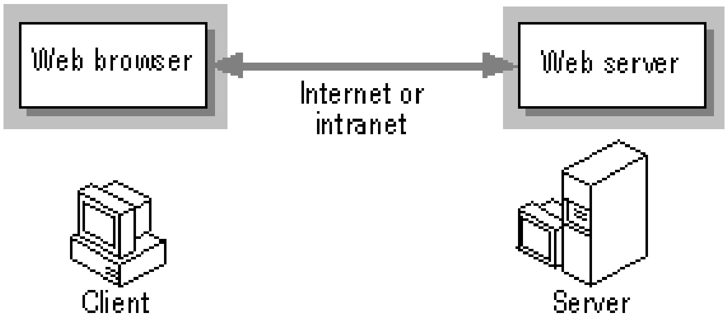

The HTML page can be a static page that has already been formatted and stored in the Web site, a page that the server dynamically creates in response to information provided by the user, or a page that lists the available files and folders on the Web site.

Internet Information Server (IIS) is the heart of your intranet Web-based services and is included in the NT 4.0 release. In addition, there are other server products in this release that are necessary for a complete Web-based intranet: the HTTP server, the Gopher server, and the FTP server. You should install all of the servers when you install IIS.

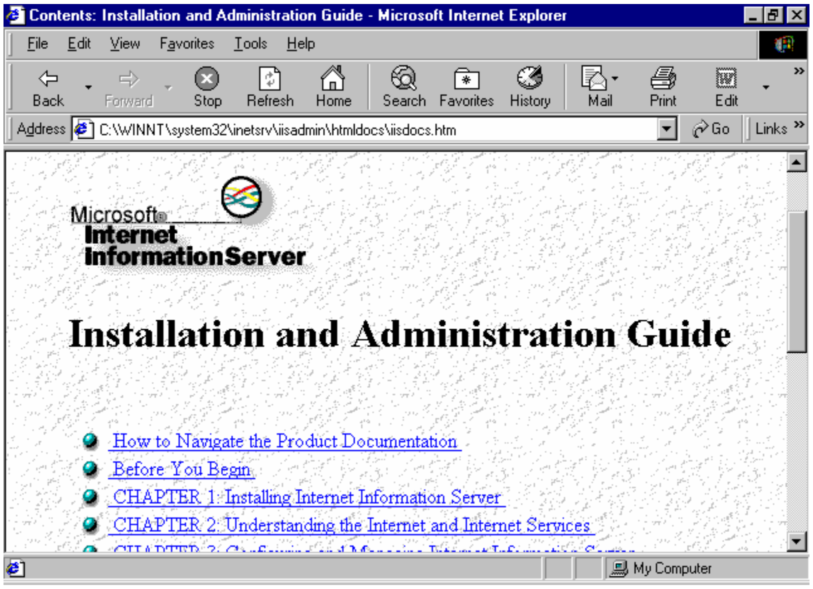

Address: C:\WINNT\system32\inetsrv\iisadmin\htmldocs\iisdocs.htm

Installation and Administration Guide

- How to Navigate the Product Documentation
- Before You Begin
- CHAPTER 1: Installing Internet Information Server
- CHAPTER 2: Understanding the Internet and Internet Services
- CHAPTER 3: Configuring and Managing Internet Information Server

When you have installed IIS, you can stop any or all of the servers if you like. However, if you think you should stop the gopher and FTP servers to conserve memory, you are treating the symptom and not the cause – get more memory. A lack of sufficient memory will starve your intranet to the point where it will lose credibility.

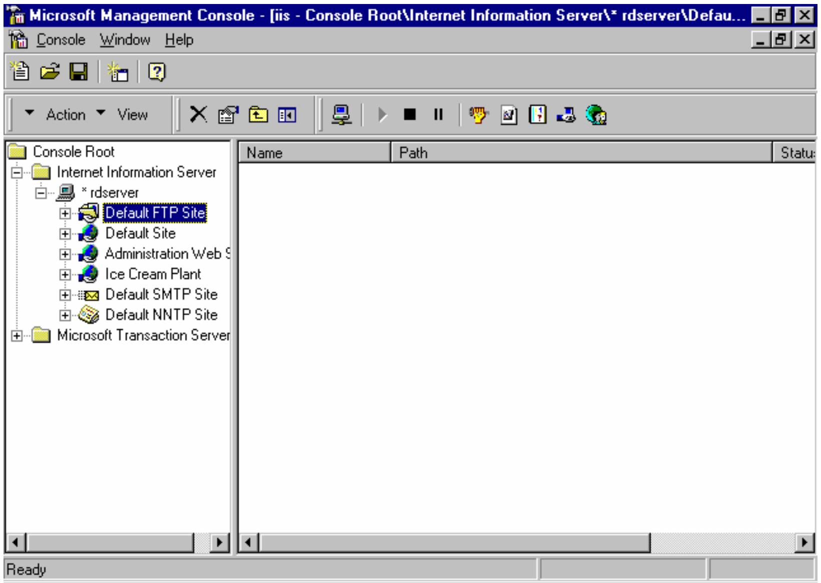

IIS include interface products as well. You will find the Common Gateway Interface (CGI) and the Internet Server Application Programming Interface (ISAPI), which will allow you to scripting languages on the server.

If you intend to use the dbWeb product or you have an Oracle, Microsoft Access, Microsoft SQL Server, Sybace, or other database that uses the Open Database Connectivity standard, you will need to install the ODBC drivers. ODBC 2.5 is also required to use the FrontPage product with IIS.

IIS gives you encryption and access security using the Windows NT access control lists. It provides support for ActiveX and Java programming when you get ready to try your hand at these advanced languages for developing Web pages.

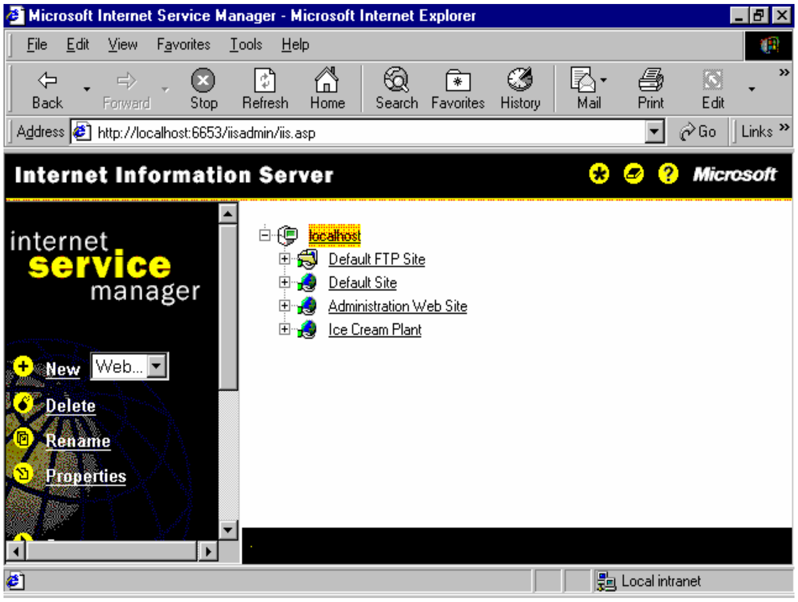

Microsoft Internet Service Manager - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Forward Stop Refresh Home Search Favorites History Mail Print Edit
Address
Go
Links
Internet Information Server
Microsoft
internet service manager
New
Web...
Delete
Rename
Properties
Local intranet

Address: http://localhost:6653/iisadmin/iis.asp

- localhost
  - Default FTP Site
  - Default Site
  - Administration Web Site
  - Ice Cream Plant

You will eventually have your intranet server running securely in a cool little room somewhere at your site. Then you can sit at your computer in your office and manage the IIS with the remote Internet Service Manager that is included. When you are ready for this application, you will want to load your Windows 95-based machine with the Windows 95-compatible Internet Service Manager.

Like NT Server, IIS has a set of administration tools. On the Start \ Programs menus you will find the Internet Information Server, which contains the following:

❖ Internet Information Server Setup

Opening this menu item allows you to add/delete IIS functions or reinstalls IIS from the CD-ROM

- Internet Service Manager

This tool shows the status of the servers currently loaded: Gopher, FTP, and WWW. It gives you control over the servers, allowing you to start, stop, or pause each individually.

- Internet Service Manager (HTML)

This tool provides the same functionality as the previous one, except it uses HTML and Explorer 3.0 to display the status of the servers. It is a good example of the coming transition from networked services and functions to HTML based GUIs.

- Key Manager

This tool allows you to establish encryption keys for secure sockets layer for HTTP.

After you have installed IIS, you can test it using Internet Explorer to view HTML files in the wwwroot folder. First, make sure that your Web server has HTML files in the wwwroot folder. Start Internet Explorer on the same computer that is running IIS. Type in the URL for the home directory of your new Web server. If your server is registered in DNS as www.mycompany.com and you want to view the file Mypage.htm in the root of the home directory, in the Address box you would type http://www.company.com/ Mypage.htm and then press the Enter key. Your home page should then appear on the screen. You can do this step first to verify that IIS is configured and running. The next step should be to access IIS from another computer on the same domain.

22.3 Microsoft Internet Assistants

If your intranet is compact enough to be manageable or your boss is too cheap to spring for Microsoft FrontPage you can use the Microsoft Internet assistance series. The Internet assistants are conversions tools that take you existing Microsoft word documents, excel spreadsheets, PowerPoint's slides, and schedule + calendars and import them into web pages as if they were HTML documents. The really good thing about Internet assistants is that they are freely available on the Microsoft web site.

Word Internet assistant

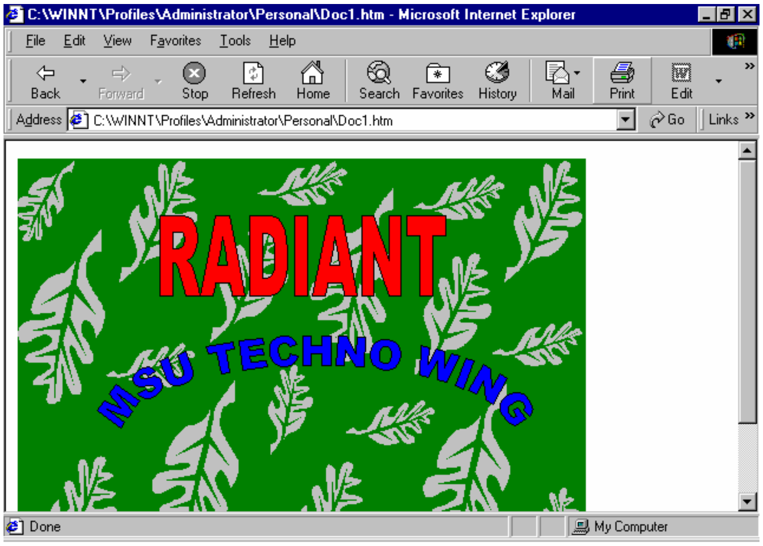
C:\WINNT\Profiles\Administrator\Personal\Doc1.htm - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back
Forward
Stop
Refresh
Home
Search
Favorites
History
Mail
Print
Edit
Address C:\WINNT\Profiles\Administrator\Personal\Doc1.htm
Go
Links
RADIANT
MSU TECHNO WING
Done
My Computer

Internet assistant for Microsoft word is free extension to Microsoft word that lets you create web compatible documents for your intranet within Microsoft word itself.

If you use Microsoft word, you are ready to create intranet documents. Internet assistant converts your word documents to HTML for you. Just save your word documents as HTML and Internet assistant automatically applies the correct HTML commands.

Internet assistant for Microsoft word provides an easy to use interface for inserting hyperlinks, Internet forms, and the formatting that people use most on the Internet.

Internet assistant lets you insert hyperlinks into word and HTML documents. You can test the hyperlinks and retrieve information from other web pages without leaving Microsoft word.

Internet assistant for Excel

Lets you convert any range of cells that exists in an excel document into HTML format for including in a web page or as a page by themselves. The file is html.xla in excel format and is copied into your msoffice/excel/library directory. When you start excel you will find the internet assistant wizard under the tools menu.

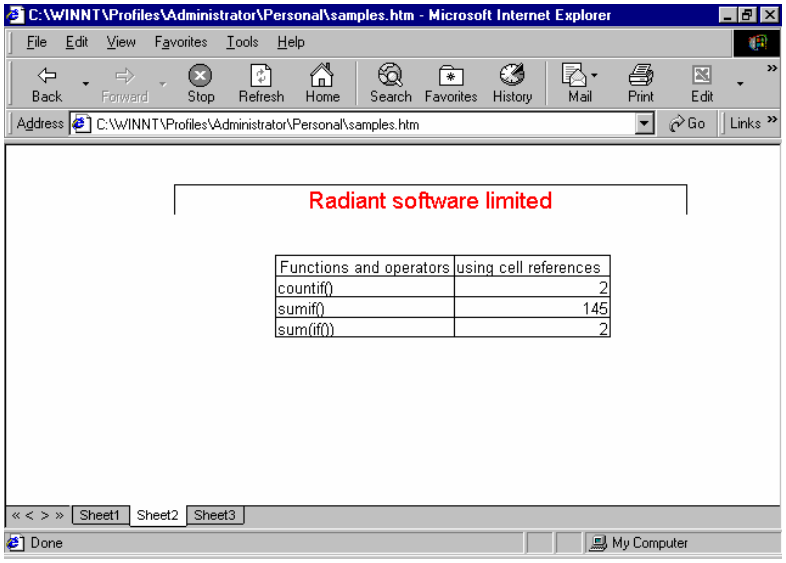

Address C:\WINNT\Profiles\Administrator\Personal\samples.htm

Radiant software limited

Functions and operators	using cell references
countif()	2
sumif()	145
sum(if())	2

Power point Internet assistant

Internet assistant for power point coverts a power point slide to HTML format. You will need to copy the PowerPoint IA into a directory an expand it by double clicking the filename. The click the .exe.file, and it will install itself into PowerPoint.

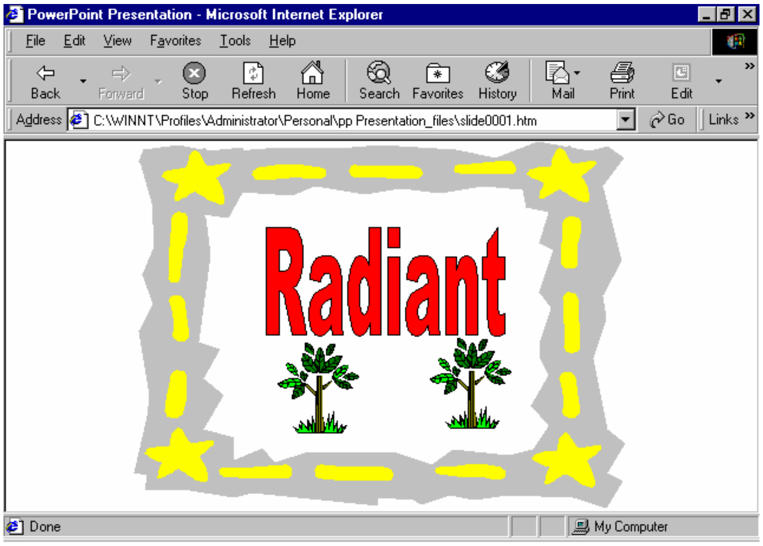
PowerPoint Presentation - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Forward Stop Refresh Home Search Favorites History Mail Print Edit
Address C:\WINNT\Profiles\Administrator\Personal\pp Presentation_files\slide0001.htm
Go Links
Radiant
Done
My Computer

Access Internet assistant

The Internet assistant for Microsoft access is for windows 95 but it will work with NT. It lets you create HTML web pages from a Microsoft access structured database. The Internet assistant can work with data stored natively in an access database, or with any other ODBC compliant database or other source. Like the other IA products, you can download access IA from Microsoft.

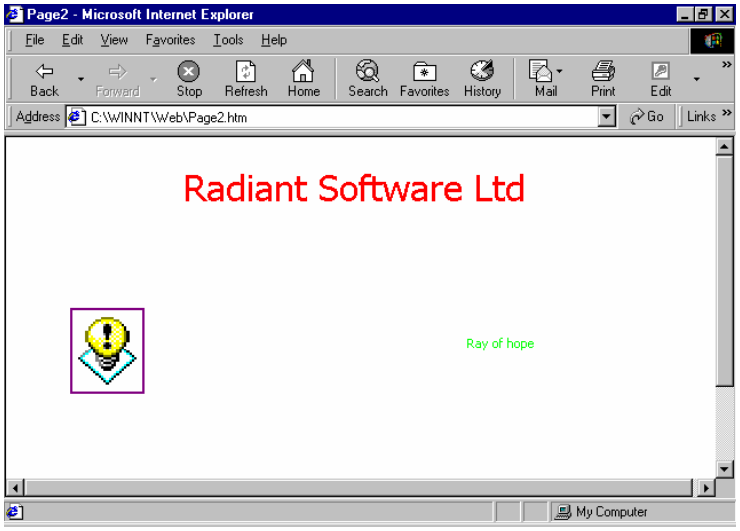
Page2 - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Forward Stop Refresh Home Search Favorites History Mail Print Edit
Address C:\WINNT\Web\Page2.htm
Go Links
Radiant Software Ltd
Ray of hope
My Computer

When you have downloaded the file page2.exe, double click the filename and it will install itself into your Microsoft access. Open your Microsoft access with a database like the northwind sample database that is included with Microsoft access. Now choose add ins from the tools menu and you will see the internet assistant. Select IA on the access menu and you will start the access IA wizard. At the opening screen of the wizard. click next and you will see all of the northwind tables (if you selected the northwind database) Put a check in the table you want to convert to HTML and click next. You can then accept the default HTML template or choose the option to select one that is to your liking. Select the one you want and click the finish button. Use your Internet explorer to view your handiwork, and there you are access data on a web page.

22.4 Intranet Search Engines

One of the big features in Microsoft products like word is the answer wizard . imagine how great it would be to have a question about your company's operations and have a great big answer wizard online. You could type in things like what is the specification for the double-ended wizbang model 3. This is not far far-fetched at all a with an intranet in place.

The intranet with search engines is a winning team. It takes a lot of back end work to set up the files so they can be found, unless you use an automated tool. When that is done, your users will be able to find that needle in the haystack that has been missing for so long.
Your company president probably likes to send out little missiles of information in hopes that the employees will read it. Instead of this ancient method, the real information that the newsletters often simply reference can be put online, and it will be there when the employees want to read it.

This section covers three search engines that exist on the Internet and shows some example of how they work. These same search engines can work for your internal intranet. I will show you how to obtain train version of these engines so you can race them yourself.

22.5 Alta Vista

Alta vista is a search service provided by digital corporation. The web site is a public service and is accessible through the world wide web from any standard web browser. Alta vista provides a simple search interface for entering a few words of a query and an advanced search that lets you use Boolean arguments. It produces a list of all the web pages that mention the content of your inquiry in order of priority. Like most of the search services, each reference in the list is hyper linked to the actual web page.

The AltaVista search public service database contains about 17 billion words indexed from over 32 million web pages. This is a database of practically everything on the Internet. Of course your intranet will not have that many indices. The AltaVista software consists of two products the data collector and the indexer. The data collector follows the web's rules, known as the standard for robot exclusion. SRE is a specification that lets a webmaster omit pages from robot probing by listing them in an exclusion file. When the AltaVista robot probes a site for data, it checks the exclusion file and skips the pages that it contains.

Alta vista probes web serves and attempts to use less than one percent of the server resources by applying delays to subsequent probes.

Collecting data is just half of the job. Once collected, the data needs to be sorted and indexed. AltaVista's indexer produces links to every word on every page brought back by the data collector, which is how you can type in a word and get back a priority listing of URLs.

AltaVista ranks indexed words and their respective URLs. The ranking systems scores each found document according to the number of in included words that were specified in the query. Documents having a higher score are presented first, and those having a lower score last. You will notice if you use the AltaVista search that the lower you go in the list of referenced URLs, the more you begin to deviate from your intended search topic.

AltaVista is available in four versions. They range in capability from indexing your PC to indexing your entire corporation.

22.6 Excite

The excite search function for the Internet is located at http://www.excite.com you can use excite to search for web documents Usenet newsgroups and the excites net review abstracts. In each of these you can perform keyword searches or concept searches. Keyword searches return documents containing the keyword concept searches return documents that relate to the meaning of the keywords.

Excite contains about 50,000 abstracts from web documents. You can search them using multiple keywords, which are treated as Boolean AND/OR logic at the same time. The search tool assigns a confidence level to each match. Documents with more instances of the keywords receive a higher confidence rating.

Excite includes netreviews, which are brief descriptions of web sites and their contents.

When you open your browser on this start page, you will be prompted step-by-step to build your very first search index. The instillation is easy and very well described. If you have an interest in search engines, this is a good example to enhance your knowledge of how searching is accomplished. The installation builds the index database and CGI query strings so you can review the entire process.

22.7 The Intranet Toolbox

Building Web pages using HTML is not a complex effort like writing a few routines in msu.HTML is human readable, and you can see instant results by opening your HTML development file with your Web browser.

In this section, were going to look at Microsoft FrontPage. Although you can write a text file of HTML commands and save it as an HTML file, you will soon find that you need a way to keep up with the far-reaching tentacles of Web links in your organization. FrontPage is designed to do just that for you.

Microsoft FrontPage

Microsoft FrontPage is a tool for building and maintaining Web pages whether they are used on an intranet or the Internet. It provides point-and-click features that follow the familiar editing and navigation modes found in all Microsoft Windows products. FrontPage includes templates and wizards that you can use to lead you through the complete design and development of Web pages that are professional in appearance, function, and form.

In the FrontPage lingo, webs are collections of HTML pages. You could have a main web like your home page for your company, an engineering web, an HR web, an accounting web, and dozens of others.

There are three Web tools in FrontPage: The FrontPage Explorer, the FrontPage Editor, and the To Do List. These tools give you all the functionality you need to build your pages. You can use FrontPage to create and maintain webs on your hard disk, on your local area network, or on a remote machine on the Internet.

You can use FrontPage at a location that does not have Internet Information Server. FrontPage is not dependent on IIS because it has its own personal Web server included. You can even use it to let others browse your own personal Web pages on your networked computer.

When you install FrontPage, you will need to also install the FrontPage Server Extensions, which are programs and scripts that support FrontPage and extend the functionality of the Personal Web Server.

The FrontPage Explorer

The FrontPage Explorer is a tool for viewing the hierarchy of the Web pages and their linkages to each other. The graphical representations of your web allow you to visualize the web while you build and maintain it.

You can build your web with the FrontPage wizards or templates. When you use a wizard or a template, your Web page is based on the features you choose from the fixed set that FrontPage include. The available functions on the FrontPage Explorer screen include

- Edit a page
  Lets you open a page with the FrontPage editor and edit the page content.

- Create a link to a page
  You can graphically link another page to a hyperlink on the page you are building.

- Control access
  You can assign permissions as to who can access your web.

- Associate your own editors
  You can associate your own editors to launch when you click on a file type.

- Verify links
  The FrontPage Explorer shows you graphically which links are valid.

- Copy a web

  You can copy a web from one server to another, even remote servers, and all relationships will remain intact.
- Import existing webs

  You can import Web pages even if they were not created with FrontPage.
- Rename or move any file

  You can rename or move any file within a web. Link relationships to that file are automatically updated to reflect the new filename and path.
- Import text and image files into a web
- Allow multiple authors to update the same web simultaneously from remote locations.

FrontPage Explorer gives you three views, which are available by clicking one for the view icons in the toolbar. The Outline view shows all the pages in a particular web and their linkages. The Linkview centers the Web page that you select and then shows all its links to all the other pages regardless of where they are located. Links are shown with arrows to designate which end of the link is the referenced end. The third view is the Summary view. This lists each page or file that shows up in the web you are exploring and lets you list its page or file properties.

22.8 Platform independent pages

FrontPage includes the personal web server to let you develop and test your web on your own computer. Because FrontPage is platform independent, you can develop a web on a personal computer and then copy it to a different hardware platform and it will retain the programming, a access control information and clickable images.

The front page editor

The front page editor lets you create, maintain, and test pages in your web. You do not need to know HTML commands or coding; you enter and edit text just as you would type a memo in Microsoft word. You can insert images or clip art, create tables, make links insert bookmarks and hotspots, and build frames and forms using the editor. The front-page editor screen with a typical HTML web page loaded for editing.

As you are having a grand time using these powerful functions, the FrontPage editor creates the HTML code used by web browsers. You will find that building a web page loaded for editing.

As you are having a grand time using these powerful functions, the FrontPage editor creates the HTML code used by web browsers. You will find that building a web page with the front page editor is very similar to using Microsoft word functions, when you are done, you can save the file as an HTML file or as an HTML template.

Using the FrontPage editor you can

- **Create a page using a wizard or template**
  You can use the editor stored templates, copy another HTML file, edit it, and save you work as a web page or as a template for other pages. If you use the frames wizard, you create a frame set that is stored within the web and has pointers to the files pointed to within the frame.

- **Insert text in different HTML styles**
  The editor lets you format text size, color, position, and other attributes

- **Insert images into a page**
  You can select from 12 graphics file types and insert them into your page.

- **Create hypertext links**
  You can insert hypertext links to anywhere in the web.   You can then follow the link and display the target page. If the link is to a multimedia file, you can run the file from its associated viewer.

- **Make a clickable image**
  The editor will let you by adding hotspots to your images. These let you click on a graphic an link to another page or another location on the same page.

- ♠ **Open pages from the World Wide Web.**

  You can open web pages and save them to your web or to your file system if permission allows you the right to copy. When you save a page, the editor can import all the images on the page into your web or file system so you will have a local copy.

- ♠ **Add tables**

  Tables let you organize your data or structure the layout of a page.

- ♠ Insert forms

  Forms let you input things like user responses and make you web interactive letting you collect feedback and data from your users.

- ♠ **Use web Bots**

  FrontPage includes a collection of web Bots. An example of a webBot is the save results Bot, which gathers information from any form in your web and stores it in various formats.

- ♠ **Convert RTF and ASCII**

  RTF is a word processing format that preserves the attributes of a document such that other types of word processing can duplicate the formatting. The editor lets you convert world, processing documents to HTML. All images in an RTF document are converted do GIF image files.

22.9 Short Summary

Microsoft Internet information server supports World Wide Web, gopher, and ftp services. The creative possibilities of what you can offer on an Internet Information Server Web site are endless. Some familiar uses are to:

- ❖ Publish a home page on the Internet for your business featuring a newsletter, sales information, or employment opportunities.
- ❖ Publish a catalog and take orders from customers.
- ❖ Publish interactive programs.
- ❖ Provide your remote sales force easy access to your sales database.
- ❖ Use an order-tracking database.

22.10 Brain Storm

1. Explain about Internet information server.
2. Explain the Internet assistants.
3. Explain some of the search engines
4. Explain some of the platform independent pages.

Lecture 23

Internal Network Considerations

Objectives

After completing this lesson, you should be able to do the following

- Discuss about wires and fibers.
- Describe the remote access considerations.
- Describe about network management.
- Describe about resource and device management.
- Describe the cross platform management.
- Discuss about web based network management.

Coverage Plan

Lecture 23

23.1 Snap Shot

This lecture discusses the impact that the intranet way of doing business can have on your existing network and what you might have to consider about the physical network itself. The key to the intranet is its platform independence. You can find a web browser for almost any kind of reasonably modern computing platform that you can imagine. The browser and TCP/IP have the potential of placing every bit of information in your company at your user's fingertips. This combination is an omni-present resource, one that is everywhere at the same time.

23.2 The Wires and Fibers

You can bet that the greatest expense in building a network is the labor to pull the wires. The ongoing expense of supporting the users comes next because, again, this is labor. The wires, connectors, bridges, routers, hubs, and more are just one-time expenses that should last through much iteration of machine replacements; you should expect to use your cable infrastructure for 10 years or more. The key to this extended life is the drive for vendors to produce more speed out of existing cable implementations. Semiconductor technology is consistently producing more speed with higher sensitivity. This translates into pushing higher data rates through the same pair of copper wires.

You want your network to be easy to maintain. It is the labor wasted searching for connections that are not labeled and wires that are not documented that will eventually by your downfall, not the router replacements or other upgrades. Use wiring closets and keeps them neat and clean. The wiring closet is just that – usually a small, closet-like where network cables terminate and routing equipment is mounted.

Put more money in your cable than you think you will need for the near term. In other words, do not scrimp on quality, and use the next higher grade cable than your predictions require. One day you will be faced with upgrading the speed on a network segment, and you will smile at your great foresight when you discover that the cable can already handle the upgrade. The little bit of money you save today on a less-capable cable will be only a tiny portion of the labor involved in pulling in a replacement later.

When you pull a cable into a room, pull it all the way to the farthest corner. You will again smile to yourself when you have to reroute the cable and find that it is actually long enough to reach. You will also avoid splicing the cable, which will only result in future problems down the road. Plan to put a data jack on every wall in every room when you wire a new room location. It's more work, but while you are up in the ceiling or under the floor, do it all once and be through with it.

Design for Reliability

You will be wise to plan on two LAN segments for every workgroup. This is not for capacity purposes. It is for redundancy purposes. Connecting adjacent users to alternate segments will have only half of the workgroup out of service if a segment fails. This is a rare occurrence, but if you follow this as a philosophy in all your cable implementations, you will have a good strong environment.

Spend even more money between your wiring closets and your servers. Go with glass fiber, even if it is running at 10Mbps today. When the day comes that you need 100Mbps or greater speed, it will be there and ready to go. It is also highly reliable and noise free. Today, the cost of fiber is coming down to the reasonable level. Remember, the money you spend today will be an embedded, fixed cost. The labor rate and escalation of labor to replace cable in the future will be far greater.

23.3 Remote-Access Considerations

Your intranet will provide a great service to your users while they are on the job. Some of them, though, will do their jobs on the highway. Your mobile users are, no doubt, out there to make your company revenue. You will want to support their information needs from your intranet through remote dial-in. They should be able to call in with a laptop computer and have all the services that they normally have at their work from home via remote connection. This section discusses the unique considerations for remote connections and their support.

The remote user will want to do one of two things: have remote access to his data or perform a function like Telnet into a machine for remote control. With remote access, you are essentially extending a network node outside of your secure intranet perimeter. Telnet is a remote connection to your perimeter.

Who Are Your Users?

Remote access is a labor-intensive task for the support staff. It is true because connecting computers together with modems is better than it has ever been, yet it is still a process that breaks frequently. Think about why your users need remote access and start from there with your plan to put it in place and give it the support it will require.
Find out how large your remote user population is going to be. Who are these people? Are they all in one sales organization or are they from several departments? Are they likely to be PC savvy, or will they need hand-holding in the middle of the night with their connection problems? Will they need national or international calling capabilities, or are they on the road inside the city limits and using wireless connections?

You will probably have a mix of home computers and company-owned laptop configurations. You will want the home user to set up his machine in a secure manner; this means that you have to write the setup instructions, and for several different types of machines. You may want to handle the laptop setup yourself and profile the machine to be user proof if possible.

Remote Applications

You will want to make yourself a list of all of the applications that your mobile users need to have. That is need, not want. Users having their laptops loaded down with illegal shareware and games are essentially a misuse of company property. Make it clear to them that laptops brought in for configuring will be reformatted to clean up the disk and then loaded with the corporate-specific applications. This will give you the opportunity to have the machines' remote-access services set up and working correctly at least once.

If your company is working to promote telecommuting, you can expect a large increase in the number of remote-access users. This will seriously impact your remote-connection hardware and supporting LAN structure. You may have more than simple dial-in modem access. Users, especially higher-level department heads, could demand ISDN service from their homes to the office. You will probably not be in a position to refuse these requests, but try to establish a charge-back mechanism to make each functional group pay for its own connections and your support requirements.

Remote-Access Support

You must understand how much your support of the remote-access user is going to impact your overall operation. You might consider checking with a peer to another company that has come through the experience of supporting a remote-access installation. Ask him what the impacts have been on his operations. Be sure to understand the overall picture of how the remote access was implemented, and for what reasons, to be sure you can gauge the relationship to your own environment.

You should tailor your support to be the best fit for the most users. You will be fortunate to have all your customers be of one class, like salespersons. It is more likely that you will have a full spectrum of users, from the CEO to yourself, as remote-access users. If your user group is very large and diverse in job function and computer applications, you might need to recommend that the support come from within each user group. This way, applications that are specific to each group of users can be supported by persons who are more familiar with what the application does and for what it is used.

Response Requirements

Some application problems might be able to wait until the next business day when you can deal with them during normal hours. You may want to set the policy for emergency response with organizational managers who can authorize next-day service. Individual users might not want to wait until the next day, and you might need the support of their management to say no to instant service.

It is a good idea to question the need for middle-of-the-night response to an application failure. You may have hundreds of sales people using remote access, yet their orders require a one-week delay before processing so there is no need for instant service.

23.4 Network Management

What do you do when there are more systems to administer than you can possibly handle? This is the great opportunity to employ a network-management software tool. Today’s business calls for fewer people and more automated, time-saving uses of technology.

Network-management applications are becoming a necessity as networks spread their tentacles into the far-reaching and dark corners of your intranet world. Probably the first problem with network-management applications is the cost. Network management is like writing software; it is one of those things you must do that no one sees or appreciates.

The key to network management is letting the software point out problems to eliminate the fires that flare up every hour of every day and a lot of nights. If you can spend less time getting it working again, your staff has more time to shape up the network. This results in less downtime and more free time. It becomes an iterative function. The better you maintain your LAN, the less time you will spend repairing your LAN.

So do you want several of your highly paid network engineers crawling around pulling on cables and looking for the problem, or do you want one engineer using a workstation to pinpoint the problem quickly? It is easy to justify the managed network solution based on manpower saved.

What Is Network Management?

Network management consists of using software tools to isolate problems in the operation of the network. it generally provides the following:

- ❖ Fault management. The usual fault-management opportunity is the inability of a user to log in to a machine to which he is supposed to have legitimate access.
- ❖ Performance management. This relates to monitoring the network for gradual degradation of the components or gradual loading of the network from increasing user access requirements.
- ❖ Configuration management. This includes monitoring the configuration of routers, server permissions, and the registry and making evaluations of the network configuration and expansion as a result of the performance management data analysis.
- ❖ Security management. Security management is specific to watching over the firewall configurations and remote-access security applications.
- ❖ Resource management. This is a longer-term support service that looks after the upgrading of operating system and peripheral device software upgrade and patches.

23.5 Resource and Device Management

Application locking is one service that resource and device management can provide. This function monitors the applications running on a network and compares the number of open applications against the license limits. When limits are reached, users cannot start another copy of the application in question. This prevents inadvertent license abuse.

Running audits of the installed software on each user machine is another application for resource-management services. Users may bring in their own, unlicensed applications and place your facility in legal jeopardy. Workgroup products allow you to find the user's machine on the network regardless of where it may be currently located in the addressing scheme. This capability also allows you to download new applications from servers to update machines. Using profiles and scripts, you can schedule machine updates in the middle of the night or on weekends with no one in attendance.

Resource and Device Monitoring

It is far better to monitor the network and take action to prevent a total failure than to react once the network has failed. With monitoring applications, the network manager can automatically poll each device, and compare its configuration and status against a known good state. If a device drifts out of a bandwidth of acceptable tolerance, the monitoring application alerts the network manager with an alarm. Alarm thresholds can be set for limits on the amount of free disk or main memory space that remains available. The amount of traffic flowing through router ports and firewalls can be sampled and displayed in graphs, making it easy to get a visual picture of the bandwidth usage.

Each device in a network is critical to effective operation of the network as a whole. Being able to automatically alert the network manager in advance of total failures keeps the user confidence in the network at a high level. Users rarely see minor degradation in network performance, but they will always note and rarely forget total network failure.

Alarms from network-management applications can be handled in various ways. There is always a screen indication, like a flashing symbol or line of text, and there can be at terminal beep or other audible alarm. The more sophisticated type of alarm is e-mail or a message to an electronic paging service or both. This provides a system that could eliminate the need for 24-hour, onsite attendance. Your staff could rotate having one or more individuals on call with pagers and remote terminal equipment that allows them to review system problems from home. Many apparent network problems can be avoided or eliminated by parameter modification and remote booting of systems.

23.6 Cross-Platform Management

It is rare to see an organization that uses all devices from one vendor. Usually vendor and device selection are cost driven. Over time, the installed base of computers, modems, multiplexers, routers, and bridges, to name a few, are from various vendors. This requires that the network-management application be flexible and work across many different platforms. There are several protocols available for the network-management application.

The worldwide implemented standard protocol for network management is the *Simple Network Management Protocol (*SNMP). This protocol was designed to provide remote management of network routers for the Internet in the early 1980s. It has been the recognized standard network protocol since 1990 and is now universally implemented in networking equipment. This protocol provides the mechanism for performing the network-management functions I have described so far.

SNMP consists of the management console, the agents, and the management information base (MIB). Commands entered at the management console query the agents, and the results are stored in the MIB. Windows NT 4.0 includes support for SNMP within the Internet Information Server product.

Other conventional methods for network management include

- ❖ RMON (remote monitoring)– This protocol is an extension to the original MIB that includes over 200 objects in nine groups. It performs analysis of the network, statistics, historical data, threshold detection, event reporting, and packet capture.

- ❖ DME (distributed management environment) – This is the current Open Software Foundation design. DME manages bridges, routers, hubs, and similar equipment but also includes support for server-based operating systems and applications.

- ❖ SNMP 2(Simple Network Management Protocol 2) – This is the current update to SNMP. It offers better security features and improvements to device MIBs communicating with

SNMP consoles. SNMP2 also has support for manager-to-manager communication, allowing event notifications to be sent between management consoles.

23.7 Web-based Network Management

Windows NT Server 4.0 includes several Web-based network-management tools. While they do not provide the depth that a protocol like SNMP does, the fact that they exist is a sign of the current direction that the network industry is headed. One advantage of Web-based management is the elimination of familiarization with a dedicated tool like the console function of SNMP. The Web-based tool uses the browser and TCP/IP or HTTP to interact with the embedded management software in routers and similar equipment.
A side benefit of the Web-based tools is a lower cost of implement. Vendors do not have to design special console software, and the time to develop for the Web environment is much shorter than for dedicated protocols and special screen software. The hardware for the manager's console can be a simple 486 PC with an operating system like Windows 95 and a browser like Internet Explorer. The bottom line is that network equipment fitted with Web-based network-management software is considerably less expensive than its SNMP equivalent. Although the SNMP protocol comes with the network equipment, the customer usually has to pay for the network-management console and a vendor-specific network-management software package to run on the console.

The capability of Web-based products to deliver remote management can help cut management costs. Web tools can help reduce the need for network technicians to be deployed to remote sites. Web management packages also work with easy-to understand browsers. Because users do not have to learn the complex SNMP consoles, Web-browsing the network-management realm can help cut the cost of staff training.

23.8 Short Summary

This chapter covers a range of subjects, all pointed at some of the more glaring aspects of supporting an internal network. The network is the point of beginning for the intranet. An intranet does not force a network reconstruction project, though. An intranet is portfolio of applications that will definitely increase the use of the physical network that is to transport the data packets. The interest and enthusiasm of the users will force the upgrading of a 10Mbps LAN to 100 Mbps. Soon the network cables will be so critical your business will screech to a halt when a major problem occurs. You will want to be ready, but you can do much more in the prevention mode to significantly decrease the odds of a failure ever occurring.

23.9 Brain Storm

1. What is network management?
2. Explain about resource and device monitoring.
3. Explain about remote applications.
4. Explain about cross platform management.
5. What is web based network management?
6. Differentiate cross platform and web based network management.

Lecture 24

On-Line Services

Objectives

After completing this lesson, you should be able to do the following

- Discuss about online services.
- Describe about the technology trends.
- Describe about profiles of major online service provider.
- Describe the America online and prodigy.
- Describe about the Microsoft network.

Coverage Plan

Lecture 24

24.1 Snap Shot

On-line services have their roots in computer time-sharing. From the mid-1960s to the late 1970s, computers were so expensive that the only economical way to use them was to conduct all processing at a central site. Users in remote locations, using terminals with no computer power of their own, were connected to a central computer via telephone lines. Independent computer time-sharing companies were created to serve firms that did not buy in-house systems.

This infrastructure was the foundation of the on-line industry. Publishers realized that loading information onto time-sharing services provided them with instant potential users. Early on-line services, such as LEXIS, were based on the time-sharing model. The proliferation of personal computers in the 1980s provided another base of potential users for collection of information. Consumer on-line pioneers such as CompuServe began allowing personal computer owners with modems to access on-line services.

24.2 Definition of Online Services

On-line services provide access to information, entertainment, communications, and/or transaction services via telecommunications. The telephone network is the typical distribution system for on-line services. Cable TV networks, satellite, wireless networks, and the unused portion of FM radio or television signals may also be used. A PC and modem are the most common devices used for accessing on-line services. However, on-line services can also be accessed via dumb terminals, screen phones, video-game machines, and handheld wireless devices. Television cable "boxes" may also be used to interact with online services, although this technology is mostly in test mode.

Companies that operate on-line services are called on-line service vendors. They are also sometimes called system operators or hosts. Content providers offer the services that are available through on-line service vendors. Content providers are segmented into two groups: (1) companies that create information, called database producers, information providers, information service providers, or publishers; and (2) marketers that provide advertisements, transaction services, promotional services, or product information services.
Many companies are both on-line service vendors and content providers; that is, they supply the host computer system and some of the information that is distributed through the system.

One key distinction between on-line services and the Internet is that on-line services companies either provide content or have a close association with a content provider, while the Internet relies on distributed, usually nonaffiliated content providers, including individual corporate entities.

24.3 Technology Trends

The graphical capabilities of on-line services have been limited by available distribution technologies. The narrowband public-switched telecommunications network now in place to deliver multimedia elements such as sounds, pictures, animation, and video. The Internet has increased modem speed requirements. To access the graphic-intensive services of the World Wide Web, for example, minimum 14.4Kbps modem speeds are required. Increased processing speeds and faster modems have allowed on-line services to offer more multimedia elements, such as graphics.

Cable TV companies, telephone companies, hardware and software providers, and on-line services companies are all experimenting with even faster speeds than the 28.8-Kbps modem for on-line services delivery. Cable modems and ISDN promise faster speeds and real-time multimedia via on-line services, but these are still in the early phase.

Cable modems will enable users to download sound files, video clips, or large text files in seconds rather than in minutes or hours. By using a combination of coaxial cable lines and high-speed telephone lines, cable modems can distribute information at 3 Mbps, compared with 28.8 Kbps for the best modems currently in the market. Cable modems will be one of the primary drivers that will bring the on-line services industry to the next level. Manufacturers include Motorola, Zenith, Intel, General Instrument, and Digital Equipment. On-line services such as Prodigy, CompuServe, and America Online are all experimenting with on-line delivery through cable modems.

Until higher bandwidth can be delivered through telephone networks or cable modems become the norm, on-line services are using CD-ROMs to deliver graphics and sound presentations to customers. CompuServe was the first on-line service to offer a CD-ROM product to its on-line members. The CD-ROM allows users to seamlessly enter the on-line service from the CD-ROM.

24.4 Profiles of Major On-Line Service Providers

America Online (AOL) is the largest for-profit on-line service in the United States. The company has the highest growth rate of any paid-access on-line service, including the Internet. The service is currently available for MS-DOS, Microsoft Windows, Apple, and Macintosh operating systems, and for Personal Digital Assistants (PDAs).

Prodigy is the third-largest subscription on-line service in the United States, with approximately 1.3 million subscribers. Prodigy's basic services include e-mail, bulletin boards, newswires, entertainment resources, and multimedia services. Prodigy software currently is offered for the MS-DOS, Macintosh, and Windows operating systems. Prodigy currently is available at modem speeds of up of 14.4 Kbps, and 28.8 Kbps is expected to be coming soon.

24.5 America Online

AOL currently maintains its operations through three primary business units: AOL Enterprises, AOL Services, and AOL Technologies. AOL Enterprises represents the company outside of North America, and it is now negotiating to bring AOL service to Japan and Europe. AOL provides all basic services, including Internet access. The AOL technologies unit is responsible for developing new technologies for use on AOL both domestically and internationally.

AOL has been aggressive in building its international operations. AOL established AOL International in 1994 to develop potential international services. The first major program developed is AOL Away From Home, which provides U.S. subscribers with a way to access the service from more than 140 cities in 42 nations worldwide.

In addition to introducing its own service to the international market, AOL is also establishing a new European Online service with partner Bertelsmann AG, a German media conglomerate. Under an agreement struck by the two companies, AOL and Bertelsmann will jointly own 90 percent of the service. The remaining 10 percent will be offered to European investors. The new service will offer the same basic service as AOL, but with a European feel. Bertelsmann's music, video, and publishing interests will all be featured on the service. The two companies plan to launch the service in France, Great Britain, and Germany. AOL and

Bertelsmann expect to expand the service through-out Europe after it has been established in its three core markets.

Service offerings/applications:

AOL offers numerous basic services covering a variety of topics. It also offers electronic messaging services that range from e-mail to Instant Messages (Ims) that can be sent while on-line. AOL also sponsors special events on-line, including guest appearances by politicians, movie and television celebrities, and musicians, among others. AOL's newest offering is the limited Internet connectivity, with Internet e-mail support and access to Internet news sites, mailing lists, WAIS, and Veronica.

Electronic messaging:

On AOL, all screen names are assigned an Internet address, which enables AOL users to receive mail from and transmit mail to the Internet as well as to other services with Internet links, including CIS and Prodigy. E-mail can be sent while on-line or in the course of a "Flash Session", which allows users to write and read mail off-line, and therefore save on-line time.

In addition to e-mail communication, AOL members can send messages to other users on-line in real time using the "Instant Message" option. With this option, a message is sent to another member regardless of the AOL service he or she is using at that moment. Instant Message cannot be delivered if the addressee is not on-line at the time the message is sent. If the message fails to reach the desired party, the sender is informed immediately.

General services: AOL's general services are divided into several "departments". Each department includes a download and upload section, a chat room, and message boards that allow users to post messages that can be read at any time by other members. In addition, each department now spotlights popular sites on the World Wide Web that reflect the department's primary interests. For example, the Sports department now offers advice on where to find the top sports site on the World Wide Web. The following is a review of the general services offered.

Clubs & Interests: The Clubs & Interests department features activities and chat rooms dedicated to special interests, hobbies and clubs, careers, and sports.

Computing: More than 250 hardware and software manufacturers maintain a presence in this area in order to answer consumer questions and respond to comments concerning their products. Articles concerning the computer industry and written by computer user group members are also available on-line for downloading.

Education: The education department enables subscribers to reach outside research resources, including the Library of Congress. Electronic University Network, National Geographic, and the Smithsonian,.

Internet Connection: The Internet Connection provides full access to the Internet, including an Internet e-mail gateway and access to newsgroups, FTP (file retrieval), and mailing lists. AOL subscribers can also download the company's World Wide Web browser directly from the service. The browser can be used to view WWW pages from throughout the world. The browsers feature compression technology that displays graphics faster than most browsers.

Marketplace: The Marketplace department features on-line shopping through several outside sources, such as Computer Express, 800 Flowers, etc.

People Connection: The People Connection includes "The Lobby" chat session where AOL members can mingle and discuss whatever topics come to mind. There is a limit of 20 people per lobby, but a new lobby is automatically created as an existing one overflows.

Personal Finance: The Personal Finance department features financial information, including stock quotes, financial news, Investor's Business Daily, and various other services.

Reference: The Reference department offers access to every reference outlet on AOL, including Compton's Encyclopedia, Software File Search, a Directory of Services, a Member Directory, the Bible, and CNN Newsroom news.

Sports : The Sports department provide up-to-date sports information culled from Headline News and USA Today Sport. Other sports services include the Odds, Standings, and Scoreboards section, which features general sports information on teams and individual players.

Special events: In addition to its departments, AOL offers members the chance to "meet" special guests on-line in the "Center Stage Auditorium." Questions are sent by pressing a special icon that calls up a line to the AOL moderator, who then passes the questions on the guest.

Technological developments:

AOL's rapid growth has resulted in a number of service problems, including a dramatic half-a-day outage. The most common subscriber complaints include busy signals during peak hours and being dropped from the network due to heavy calling volume. The growth has also taken a toll on the company's primary mail system. The substantial mail volume has led to mail delays of hours and sometimes days. Additionally, in the past, mail sent via the Internet has been occasionally returned unsent, or "bounced," because the AOL mail system could not handle the load. These and other problems led the company to announce a series of changes in 1994 and 1995. The solutions that are now being implemented include the following:

- Acquisition of Advanced Network & Services (ANS), BookLink Technologies, and Navisoft. AOL acquired ANS, the company that built the backbone of the Internet, in early 1995 for $20 million in cash and $15 million in stock. AOL purchased BookLink Technologies and NaviSoft, both Internet service companies, in 1995 for more than $35 million in stock. The three companies are now refining AOL's Internet access software as part of the company's Internet Service unit.

- Acquisition of WAIS Incorporated and Medior Incorporated. This move is designed to improve AOL's ability to help information sources get on the Web. AOL has confirmed deals with several software companies to allow "one-button" access to AOL services.

- Additional support staff: AOL has expanded its Technical Service and Member Services & Billing staffs in Vienna, Virginia, and has also established a new-member support site in Tucson, Arizona.

- Web Shark: AOL plans to offer InterCon's Web Shark World Wide Web browser to its Macintosh-based customers in the near future. AOL released its Windows-based WWW browser in May. AOL officials say the company will be the first of the major services to offer a WWW browser for the Macintosh.

Multimedia: AOL is dedicated to introducing the latest technologies, including multimedia, as they become available. The company already offers "multimedia press kits" for downloading, and is participating in interactive trials being conducted by Intel, General Instrument, Viacom Cable, and Comsat Cable. The trials are designed to test the delivery of on-line services via cable directly to PCs. AOL's interest in multimedia led the company to join forces with LANcity, an equipment manufacturer. The companies plan to jointly support the distribution of PC-based multimedia on-line services via cable television lines. LANcity will provide cable access to AOL. AOL's interest in multimedia also led to its involvement in 2Market. 2Marekt's home-shopping services are delivered via CD-ROM as well as via AOL and Apple Computer's on-line service, eWorld. The CD-ROM and on-line services both deliver detailed information about goods for sale, and they feature color photographs of the items

Internet Software: AOL's Internet service unit will offer two new software products, NaviServer and NaviPress, as part of an integrated World Wide Web publishing system. The software will enable content creators and providers to easily establish and maintain Web applications. While NaviServer and NaviPress can server as stand-alone products, they are part of a client/server application development system. NaviServer, the server product, is designed to offer a powerful architecture to support applications by servicing requests and managing content. NaviPress, the client product, offers a point-and-click interface for creating, editing, and linking content.

High-speed, dial-up network: In 1995, AOL signed a deal with Bolt Beranek and Newman (BBN), provider of Internet technology and services, in which BBN will build, maintain, and operate a portion of AOL's nationwide, high-speed, dial-up network. AOL established its own nationwide network, AOLNet, to provide members with improved high-speed access. AOLNet was formed through ANS (Advanced Network & Services), the corporate Internet access provider and builder of the Internet's backbone, which AOL acquired earlier. AOLNet is a high-speed network that enables users to access AOL at 14.4 Kbps, 28.8 Kbps, and ISDN.

24.6 Prodigy

Prodigy has announced plans for an improvement in both the GUI and the process by which services are developed, and a plan involving direct links between the three leading on-line services. The direct links would allow users to access information from any of the three on-line services and may lead to development of enhanced services.

To combat subscriber attrition and to attract new subscribers, Prodigy as introduced several service changes in recent years. In 1993, the company offered a GUI option; this has increased Prodigy's attractiveness to graphics-based computer users. Prodigy has also added full-color photographs to certain sections of its service. And in 1994 Prodigy debuted a new version of its software that offers access to multimedia services.

Prodigy's multimedia package requires a Windows interface, a sound card, and speakers. The service primarily adds sounds to the various portions of the Prodigy service, including a two-minute spoken news update that accompanies the photos and newswire stories on-line.

Service offerings: In 1995, Prodigy, along with Berlitz Publishing company began offering e-mail in foreign languages. This service also provides users access to language-related publishing, translations, and languages training products.

All Prodigy users are assigned an e-mail address. Prodigy's standard e-mail service includes a Mail Manager service that, at the push of a button, goes on-line, sends any mail to be transmitted, and retrieves unread mail from the subscriber's mailbox. The Mail Manager then logs off. This grants customers the freedom to leave the room and attend to other tasks as mail is being exchanged. A similar service is now available on AOL.

In addition to these communication services, Prodigy offers several specialty services designed to meet the interests of its customers. Current topics of interest include the following:

Investments: Prodigy offers stock advice in the strategic investor section; current advisers include Graham & Dodd and CANSLIM. Prodigy also offers access to the nation's largest discount on-line stockbrokers; PC Financial Network. Therefore, users can purchase or sell stock any time of the day or night via computer. Also, Prodigy offers an additional feature, "Quote Track," which automatically calculates the loss or gain of any number of stocks in a user's portfolio while on-line.

Children's Interests: Prodigy offers a variety of educational resources, including an on-line encyclopedia. In addition, Prodigy offers real-time adventures, polls, and interactive games for children. Other features include NOVA experiments and science discussions, and an interactive version of the popular children's television show Sesame Street.

News and weather: Prodigy maintains a newsroom that constantly scans the top national and international news resources and edits stories for transmission on Prodigy.

Travel: Prodigy offers access to Sabre travel service, which can be used to make travel arrangements, including booking flights, on-line. Sabre also offers free overnight delivery of tickets booked through the agency.

The Cable Guide: Prodigy offers an on-screen, interactive version of The Cable Guide's television listings, where customers can scan program listings for the following seven days. Listings can be called up through several categories, including day, time, network, and type of program. Prodigy is now developing enhancements for the on-line Cable Guide, including the possible addition of pictures, sound bites, and video clips from shows listed in the guide.

America's Talking: Prodigy offers on-line subscribers a way to converse with the hosts of several new programs presented live each day on "America's Talking," a cable network owned by NBC.

24.7 Microsoft Network

Technological developments:

ISDN bundle: Microsoft and partner Pacific Bell and CompuUSA are negotiating to bundle and distribute a software ad hardware package that would enable subscribers to access the MSN via ISDN lines. As planned, the bundle would include Windows 95, an ISDN adapter, and ISDN telephone service. The bundle will be sold in CompuUSA stores throughout California.

NBC partnership: Microsoft and partner NBC are joining forces to develop entertainment resources, including CD-ROM, interactive games, and on-line offerings. NBC has agreed to withdraw its sites from other on-line services, thus offering MSN subscribers exclusive on-line access to the network.

Internet strategy: Microsoft has positioned the Microsoft Network as an on-line service offering seamless integration with the Internet. It has also portrayed its development tool, Blackbird, as a robust multimedia tool that will allow developers to implement full motion and sound. Microsoft is counting on Blackbird to deliver content providers and electronic merchants richer and more compelling multimedia features than the Web's HTML programming language or programming languages of existing consumer on-line services.

According to Microsoft, Blackbird is a development platform with broad applications, including on-line, CD-ROM, broadcast, and interactive TV. The tool includes a strong search engine, called The Find, which allows users to create personal profiles. Other Blackbird features include drag-and-drop design, information retrieval, and open extendibility, which allows third-party developers, such as Adobe or Macromedia, to integrate Blackbird into their tool sets.

By mid-1996, any user with a Windows / Mac-based client was expected to be able to access MSN through the Internet. Internet access providers will be able to resell MSN, allowing customers to access MSN through the Internet. Microsoft's Web browser – Microsoft Internet Explorer - is designed specifically for the Microsoft Windows 95 operating system. The program includes built-in RealAudio technology, providing real-time audio capabilities. The Internet Explorer supports Windows 95 shortcuts to the Internet and supports full drag-and-drop text and graphics.

24.8 Short Summary

Microsoft has provided point-and-click access to the MSN's client software built into the Windows 95 desktop. The service is available via a local telephone number in 47 nations worldwide, with support for 20 different languages. An Apple Macintosh version of the software is also in development and was expected to be available by press time. Microsoft does not plan to import the MSN client for IBM's OS/2 platform. Microsoft's expansion plans also include improved modem-access speeds. The service currently supports 14.4 bps (and lower). Microsoft expected to introduce 28.8-bps and ISDN access by press time. Additionally, Microsoft and partner Tele Communications Incorporated (TCI) plan to begin testing cable-based delivery of the service in 1996.

24.9 Brain Storm

1. Define online services
2. Explain about technology trends.
3. Explain about some of online service providers.
4. What is Microsoft network?

Lecture 25

Third-Party Tools

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the shareware, freeware and crippleware.
- Describe the other Internet servers – freeware EMWAC servers.
- Describe about distributed management.
- Describe about desktop management interface.

Coverage Plan

Lecture 25

25.1 Snap Shot

This lecture covers a range of supporting products. First, there are other Web servers you can run under Windows NT. There are some services, such as the Telnet daemon, that are not included in IIS. You can find in this lecture a brief introduction to a few good Web servers, Telnet servers, and other products. You might find that IIS does not satisfy your requirements and want to evaluate another complete Web server. This lecture also discuses the Desktop Management Interface and how it is evolving in the industry. This will lead into a discussion of the Tivoli desktop management product included on the CD.

25.2 Shareware, Freeware, and Crippleware

Once you have connected to the Internet you will find hundreds of new sites from which you can download products to try before you buy. Included on the CDs with this resource library are a lot of shareware, freeware, and crippleware programs. Remember that if you decide to continue to use a product that is labeled shareware after its expiration date, you are legally obligated to pay for it. Products on the CDs that are labeled freeware are just that – free software. One example of this is the Executive Software's Diskeeper Lite product. Crippleware is my term for the commercial products that are my term for the commercial products that are placed in the public domain but that are limited in their functions or have an expiration date on their use.

The main reason software is marketed in this manner is that most software companies now realize that buying software is like visiting a doctor. You've got to really get down to the details to see whether the treatment will work. With software you have to get down to the details to determine whether it's right for your business or whether you'll have to use a software kludge to make it work.

So in evaluating the products included on the CDs and those described in this lecture, remember that people out there depend on your honesty. They make their living writing some very good code, let you use it, and hope you'll pay for it if it does what you want. In the case of crippled software, you will not have a choice when it expires; you will have to find something else, reload the application form scratch for another trial, or let your honesty be your guide.

25.3 Other Internet Servers

Web services offered by non-Microsoft Web servers vary from product to product.

Freeware – The EMWAC Servers

The European Microsoft Windows NT Academic Centre (EMWAC) has been offering a freeware product for quite a while. There is also a professional version of this server that you can get for a fee. However, if you just want to look at other servers to test features and decide what is right for you, you might as well start with one of the first.

The product supports only the basic HTTP protocol. It is run under Windows NT as a service and can be configured by a Control Panel applet that is added during the installation.

25.4 Distributed Management

It's becoming true for more companies today: Your business is dependent on your network. Companies are using networks of distributed server and client systems to support many mission-critical applications. And when a business relies on its client/server network, it's important to ensure that every system and every server keep operating at peak performance.

The challenge is that, without effective systems management, the risks and hidden costs of network computing will increase. It is estimated that the total five-year cost of ownership of a networked PC is as much as six times the purchase price, and the cost each year typically meets or exceeds the initial hardware investment.

The DMI Framework

The Desktop Management Task Force (DMTF) was formed in 1992 to develop and deliver the enabling technology for building a new generation of PC systems and products. The DMTF's goal is to provide a common management framework for PCs and encourage vendors to quickly bring manageable products to market.

The DMTF has delivered the industry-standard Desktop Management Interface (DMI) to help solve desktop system and server interoperability and management problems in a cross-platform environment.

With DMI 2.0, which was completed in April 1996, static as well as dynamic information can be sent across the network to a remote station, providing a remotely accessible framework for management. DMI implementation in operating systems and management applications provides users with asset management, configuration management real time monitoring, and control capabilities.

Motivation

Within a computer system, there is a gap between management software and the system's components that require management. Mangers must understand how to manipulate information on a constantly growing number of products, which is undesirable from a cost standpoint. The DMI acts as a layer of abstraction between these two worlds.

The DMI has been designed to be

- Independent of a specific computer or operating system
- Independent of a specific management protocol
- Easy for vendors to adopt
- Usable locally no network required
- Unable remotely with DCE/RPC, ONC/RPC or TI/RPC
- Mappable to existing management protocols (for example, CMIP, SNMP)

The DMI procedural interfaces are specifically designed to be remotely accessible through the use of remote procedure calls (RPCs). The RPCs supported by the DMI include

- DCE /RPC
- ONC / RPC
- TI / RPC

Basic Terminology

In this section, system means a computer system. Components are physical or logical entities on a system, such as hardware, software, or firmware. Components may come with the system or may be added to it. The code that carries out management actions for a particular component is known as the component instrumentation.

A management application is a program that initiates management requests. It uses the DMI to perform management operations. The management application may be a program such as an application with a graphical user interface (GUI). It may be a network management protocol agent that translates requests from a standard network management protocol (such as SNMP or CMIP) to the DMI and back again.

25.5 DMI Service Provider

A DMI service provider is analogous to the DMI service layer of previous DMI specifications.

Elements of the DMI

The DMI has four elements:

- ❖ A format for describing management information
- ❖ A service provider entity
- ❖ Two sets of APIs – one set for service providers and management applications to interact, and the other for service providers and components to interact.
- ❖ A set of services for facilitating remote communication.

Component descriptions are defined in a language called the Management Information Format, or MIF. Each component has an MIF file to describe its manageable characteristics. When a component is initially installed into the system, the MIF is added to the MIF database.

DMI service providers expose a set of entry points that are callable by component instrumentation. These are collectively termed the Service Provider API for Components. Likewise, component instrumentation code expose a set of entry points that re callable by the DMI service provider. These are collectively termed the Component Provider API.

Component providers to describe access to management information and to enable a component to be managed use the component interface, or CI. The CI and the MIF shield vendors from the complexity of encoding styles and management registration information. They do not need to learn the details of the popular and emerging management protocols.

The DMI service provider also exposes a set of entry points callable by management applications. These are collectively termed the Service Provider API for management applications. Likewise, management applications expose a set of entry points callable by the DMI service providers. These are collectively termed the Management Provider API.

The management interface, or MI, is used by applications that wish to manage components. The MI shields management application vendors from the different mechanisms used to obtain management information from elements within a computer system.

The procedural MI is a remoteable interface designed to be used with one of the supported RPCs. Remoteable means that the interface is available from a remote location through a dial-up connection or over the network.

The DMI service provider is an active, resident piece of code running on a computer system that mediates between the MI and CI and performs services on behalf of each.

The DMI Service Provider

The DMI service provider coordinates and arbitrates requests from management applications to the specified component instrumentation. The DMI service provider handles the runtime management of the MI and CI, which includes component installation, registration at both levels, request serialization and synchronization, and general flow control and housekeeping.

The interfaces have been designed so that commands at the MI level are either satisfied at the DMI service provider or passed directly to the CI.

Below Figure shows a possible DMI service provider block diagram.

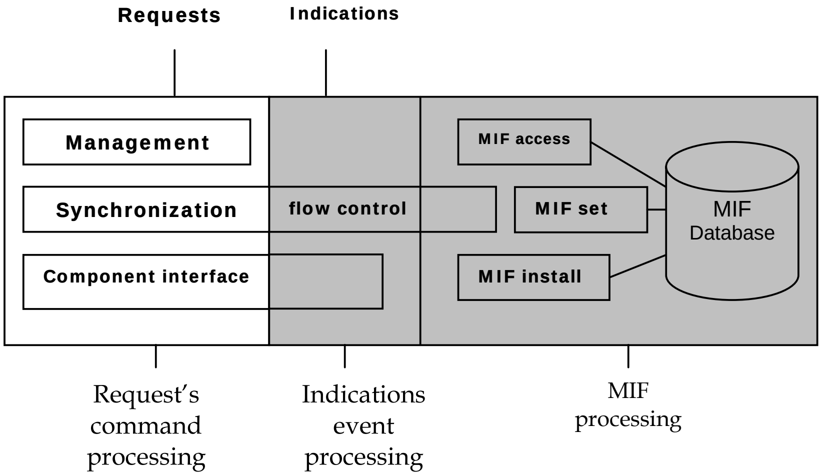

The DMI service provider coordinates the dynamic installation and removal of component instrumentation and management applications. It ensures that at least group 1 (the component ID group) in each is installed.

The DMI service provider coordinates the registration of entities wishing to initiate management activities. The DMI service provider is responsible for all runtime accesses to the MIF data. Implementations of the DMI service provider may choose to store MIF files in an internal format (a MIF database) for performance and ease of access.

The DMI service provider is responsible for launching the component instrumentation, if necessary. The DMI service provider enforces command serialization to a component instrumentation and ensures that commands are allowed to run to completion. Multiple requests for particular component instrumentation must be queued.

The DMI service provider supports event/indication subscription and filtering.

The DMI service provider forwards indicates based on subscription and filters to each registered management application, and must time-stamp incoming indications before forwarding them.

The DMI service provider sends indications to all registered management applications that have subscribed for indications when components are installed or removed form the MIF database.

The DMI service provider appears to management applications as a component with ID1. As a component, it must support the standard component ID group. Additionally, the DMI service provider must support the subscription indication and filter standard groups. Also, like any component, it may define additional groups beyond the component ID groups.

The DMI service provider supports all of the NLS mechanism contained in the DMI version 2.0 specification, including Unicode and multiple NLS installations of schema for each component.

25.6 Short Summary

In this section, system means a computer system. Components are physical or logical entities on a system, such as hardware, software, or firmware. Components may come with the system or may be added to it. The code that carries out management actions for a particular component is known as the component instrumentation.

A management application is a program that initiates management requests. It uses the DMI to perform management operations. The management application may be a program such as an application with a graphical user interface (GUI). It may be a network management protocol agent that translates requests from a standard network management protocol (such as SNMP or CMIP) to the DMI and back again.

25.7 Brain Storm

1. What is shareware?
2. What is freeware?
3. Explain about EMWAC servers.
4. Explain the elements of DMI service provider.
5. Explain about distributed management.

Lecture 26

Broadband Communications for the Internet and Intranet

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the Broadband communication for the Internet and intranet.
- Describe the services and requirements of Broadband communication.
- Describe network architecture supporting broadband communication.
- Describe about integrated services digital network.
- Describe about wireless.

Coverage Plan

Lecture 26

26.1 Snap Shot

This chapter discusses the technologies available for this required upgrade, as well as the issues surrounding the need to support the expected growth in requirements and grade of service. The focus of this chapter is on increasing actual bandwidth by utilizing existing and evolving communication services such as Integrated Services Digital Network (ISDN) – for access – and Asynchronous Transfer Mode (ATM) – for access and for backbone requirements.

26.2 Broadband Communications for the Internet and Intranets

The popularity of the Internet is growing, along with the number of users and applications. As more and more people use the Internet, and as the Web applications available on it increase in complexity, the underlying end-to-end communications technologies will need to be upgraded. Otherwise, the Internet may eventually become an ineffective means of global communication. It is already apparent when downloading a large file or a high-resolution graphic that speed is not one of the Internet's greatest assets at this time. By upgrading to higher speeds and/or broadband communications technologies end to end, the Internet will be able to support a multitude of production-level users and complex multimedia applications. Higher-speed facilities are needed for both the access portion of the network and internally to the Internet itself, namely, in the set of backbones that comprise it. The need for higher speeds is also evident for intranet applications now evolving, as discussed throughout this text.

This chapter discusses the technologies available for this required upgrade, as well as the issues surrounding the need to support the expected growth in requirements and grade of service. The focus of this chapter is on increasing actual bandwidth by utilizing existing and evolving communication services such as Integrated Services Digital Network (ISDN) – for access – and Asynchronous Transfer Mode (ATM) – for access and for backbone requirements. Although the typical corporate planner has little influence on what is inside the Internet in terms of communication facilities, the planner has a lot of control on the access facilities. The planner does have some control by selecting an ISP that has paired with an NSP that has a high-capacity backbone. However, for a specific Internet session, the actual end-to-end network performance also depends on the ISP that the target server has chosen – hence the overall performance is beyond the direct control of the planner. In addition, there is a short discussion of broadband initiatives as they relate to the Internet and to other networks.

26.3 Services and Requirements Driving the Need for Broadband

New services desired on the Internet

Lists some of the services that are currently of interest, or may become of interest, to users over time. Some of these services are realized outside of the Internet, but many do utilize the Internet, at least in some form (e.g., to obtain educational materials, publications, or technical specifications).

A number of the services listed in following Table may find a reasonably high level of market penetration. These include interactive television, video-on-demand, home shopping, videoconferencing, remote monitoring and security, government service delivery, and Internet surfing. Many people want these services in the context of the Internet.

Typical services driving broadband requirements

+ Government
  - Electronic form lodgment
  - Document and data transfer
  - Videoconferencing
  - Electronic access to government information
  - Electronic voting and publish consultation
  - Video libraries

+ Health services
  - X-ray and CAT scan transfer
  - Video consultation
  - Medical records transfer
  - Remote monitoring of outpatients

+ Education
  - Remote interactive teaching
  - Access to local and overseas libraries
  - International classrooms
+ Business
  - Videoconferencing, including "virtual meetings"
  - Electronic commerce
  - Remote interactive training
  - Computer-aided design and manufacturing
  - Multimedia communications

+ Home services
  - Video-on-demand
  - Home shopping and banking
  - Interactive multimedia
  - On-line information services
  - Video mail and videophones
  - Security services and utility metering
  - Telecommuting.

Many of these services utilize hypermedia and GUIs to make it easier for users to access information and to navigate between elements of the distributed repository. Hypermedia / multimedia is emerging as an effective way to communicate, but traditional telephone lines and modems cannot effectively deliver multimedia to homes. To overcome this hurdle many on-line services and service providers are migrating to huger-speed data links such as ISDN, ATM, and Synchronous Optical Network (SONET). As consumers desire more of these bandwidth-intensive services, there is going to be a need to install high-speed communication lines and services in their offices and homes, and to upgrade the speed of the Internet backbones.

People spend much of their lives producing or consuming information – through work, conversations, watching television, or reading newspapers. Almost every home in America has a telephone connection. People receive radio and television services, and in some places, satellite services. A rapidly increasing part of the population has mobile telephones; around

40 percent of homes have persona computers; and a vast majority has videocassette recorders. There is a growing household use of telephone-line modems, and the on-line services industry in the United States is also growing at a significant rate. All of this sets the foundation for increased demand for on-line services in general, and Internet services in particular.

On a going-forward basis, networks must be able to do the following: support broadband services such as high-quality video; allow people to create and distribute their own material; support and interconnect both fixed and mobile systems through a combination of cable, satellite, and local wireless services; enable multiple parties to be connected to an interactive multiparty, multimedia, multiconnection high-speed service; support interactive and switched two-way transmission, as the telephone network does now; and provide global reach. Evolving broadband communications technologies, such as ATM, will play an important role in achieving these goals. These requirements must be supported not only by networks that can be considered as an overlay to the Internet, but by the Internet itself.

Apart from the Internet itself, a seamless high-speed network carrying voice, data, and video services will be used as a pipeline to brig an expanded universe of information and entertainment into the home and the workplace. But with the Internet, thousands of movies, mail-order catalogs, newspapers and magazines, educational course, airline schedules, and other information databases will be available with a few clicks of a PC mouse or perhaps a TV remote control. Two-way video-conferencing may become integral to the family, to social life, and to business in the not-too-distant future. The information and communication infrastructure of the future, based on fiber optics, will provide the main conduit for global entertainment, commerce, information, and communication in the next century.

Consumers are a significant, if not primary, element in the evolutionary process. They will determine which services thrive, which services enhance lifestyles, and which afford greater financial opportunities. We might have a fair idea of what the technologies will look like, but it is the services they support and the content they carry that will ultimately determine a technology’s future. The penetration and affordability of consumer equipment, a population with the necessary skills and training to use the services, a consumer base willing to adopt new means of interaction, and competitive service prices will also be influential in a technology’s success.

The possibilities for the Internet seem almost endless. For example, voice connection via the Internet are now possible. This allows people to participate in telephone conversations anywhere in the world for the price of an Internet connection. Applications like these are what make the internet so attractive to people as more and more uses of the Internet become available, more and more users will be drawn to it. Competition among companies will make WWW pages more complex and will put further loads on the Internet.

A ubiquitous system of broadband backbone networks will allow the United States to be competitive with other countries in terms of economic development and commerce: to create jobs; to ensure equity and quality in education and health care; to share limited resources; to reduce government costs and provide efficiencies; and to provide a platform for the delivery of Information Age services to consumers. Until recently, however, little visible progress has been made toward its realization. Telephone companies, newspaper publishers, cable television operators, and other potential players have been involved in lengthy congressional and court battles. Meanwhile, the pioneers of the computer-mediated communication networks collectively referred to as cyberspace, are not willing to wait. Employing whatever tools they can find, they are constantly pushing the techno-cultural envelope. The Internet has become the de facto information superhighway.

Many expect a replacement of the existing broadcasting model of distributive (one-way) information services by a system in which consumers create and have access to information and exercise choice and control in their communications. This is called multicasting or point casting. In the communications arena, people will be able to send or receive large amounts of information – video, audio, text, graphics, or data – anywhere at any time. But the major benefit will be the speed and ease of access to materials for production of publications. These examples are forerunners of new applications in medial travel, government, advertising, and general publishing, which have the potential to earn income either from the provision of information-content services or export of technology.

26.4 Network Architectures Supporting Broadband

The goal of this discussion is to sensitize the planner to issues related to the need for broadband communications over the Internet and to offer a set of choices for the access subnet work (which is generally under the planner’s control).

Approaches to broadband

In the discussion that follows, you may wish to differentiate between the physical connectivity service (eg. ISDN, ADSL, HFC, or FTTC) and a value-added service such a ATM and frame relay (both of which can, theoretically, run over any physical medium). Note, however, that not only do you need higher speed at the physical layer, but also at the network layer. This is the idea behind new versions of IP.

JPEG (Joint Photograph Experts Group) is a compression format that reduces an image's digital representation to smaller-size files, based on desired resolution. The image is divided into a series of blocks, which JPEG then processes and compresses, resulting in compression ratios of 10 to 1 or 20 to 1. JPEG supports both a lossy and lossless compression scheme. Also, a number of PC-based videoconferencing systems that appeared in the mid-1990s used a version of JPEG, called Motion JPEG, to provide reasonable quality video.

MPEG (Motion Pictures Experts Group) is a standard for compressing and storing video on compact discs. To achieve its ratio of 50 to 1 (or even 200 to 1), MPEG uses a compression scheme similar to JPEG. MPEG goes a step further and also eliminates any redundancies between frames.

Fractal compression, an emerging technology, relies on mathematical modeling to "disintegrate" an image into repetitive shapes. Fractal compression requires only a small amount of information about each shape and can shrink the image file or video frame into a very small space, achieving ratios of up to 2500 to 1 for images and up to 100 to 1 for video playback at 30 frames per second.

The rest of this section examines available Internet access/backbone technologies.

ADSL

Asymmetrical Digital Subscriber Line (ADSL) is a technology that allows for high-speed transfer of data streams over the existing copper plant. It supports between 1.5 and 6.2 Mbps, depending on the technology used. However, the higher rate is achievable only in one direction; hence the name asymmetric. An American National Standards Institute ADSL standards project initiated in 1992 released a first issue of the T1.143 standard for ADSL that specified downstream transmission of up to 6.2 Mbps and a return path of 224 Kbps over

12,000 feet of 24-guage copper plant. It also cited a downstream rate of 1.5 Mbps being achievable over 18,000 feet. A 6.2 –Mbps transmission rate allows the user to download files 200 times faster than a conventional modem using the same copper wire, and it does not affect normal telephone service. ADSL essentially allows the user to achieve T1 and greater speeds over low-cost unrepeated facilities.

HDSL

High-bit-rate digital subscriber line (HDSL) is a full-duplex, symmetrical technology that allows T1-like speeds over ordinary copper wire. This technology is similar to ADSL, except for the fact that it provides a symmetrical data flow with the downstream and upstream channels being equal. Advances in HDSL signal processing equipment have reduced the infrastructure required to implement the technology. New transceivers allow a single twisted pair to carry full T1 or E1 payload. The only drawback of the single pair configuration is that it will be limited to shorter distances than the dual-pair system.

Due to the symmetry and full-duplex nature of HDSL, it will most likely be used for access by small businesses or home workers for applications such as telecommuting, data services, LAN interconnect, and frame relay rather than for entertainment. This will have limited use in residential markets, which are aimed more toward Internet access and entertainment and which tend to be asymmetrical.

BDSL

Broadband digital subscriber line (BDSL) was developed to overcome the bandwidth limits of ADSL by providing 25 Mbps over the existing copper plant, using a fiber-to-the-serving-area approach. Broadband signals are delivered over fiber to remote nodes for distribution to 100 or so homes. The remote nodes receive broadband signals from fiber feeders at rates of OC-12 or OC-24. These nodes are located within 3000 feet of the subscriber. The technology that is being developed for fiber-to-the-curb (FTTC) systems was adapted to provide a low-cost switching and feeder technology. Since the length of copper used was reduced by getting closer to the home, the signal-processing requirements have been considerably simplified. This reduction in signal processing reduces the amount of transistors per transceiver by 50 percent, even though the transmission rates are substantially increased. This reduces the cost of the technology that must be used at each of the nodes, which can serve several hundred subscribers. This technology can be utilized to provide broadband services, including broadband Internet access, in Multiple Dwelling Units (MDUs), by placing the electronics in the basement or common space, and then using ordinary house wire in the raisers.

VDSL

Very high bit rate digital subscriber line (VDSL) can be used in FTTC applications to provide a downstream link of 51.84 Mbps and up to 2 Mbps for a return channel over 200 meters or more of existing copper. This technology can also be used to implement an in-building multidrop wiring bus. This would allow all the devices in the home, such as set-top boxes and smart controllers, to communicate on an ATM bus. The ATM Forum has approved a 25.6-Mbps interface for low-speed ATM User-to-Network Interface (UNI) access.

26.5 Integrated Services Digital Network

Integrated Services Digital Network (ISDN) is a relatively new design of services that are rapidly becoming available, mostly from the switched public telephone systems. The beauty of ISDN is that it uses the subscriber’s existing telephone cables that connect to the switched public network. It also provides a multitude of services. Services that can be used with ISDN include your telephone; video conferencing, LAN-to-LAN connectivity, and really anything that uses packet data to get information from one point to another and needs a relatively speedy method and modest cost. ISDN is a benefit to the subscriber, especially in today’s world of downloading intense graphics and multimedia content for Web pages. All of the services mentioned here are also available on a single telephone line connection. ISDN is rapidly gaining popularity and as it does, prices are just as rapidly coming down to the more reasonable range.

ISDN can provide as many as eight separate connections for devices like telephones, computers, and fax machines. It can support two phone conversations and one data connections at the same time. You can call or receive calls and fax at the same time, and be downloading a multimedia Web page simultaneously – and the Web page can be downloading at the rate of 64Mbps, many times the rate of today’s 28.8Kbps modems. But there is another feature of ISDN that makes it even better. If you establish your ISDN channel as a single telephone grade line and one data line, the speed of the data line becomes 128Mbps. This will move your Web page so fast it will look like it is stored on your own hard disk.

In the modern world, telecommunicating is becoming more and more a reality. The connectivity provided by products such as Microsoft NT Server and Windows 95 enables the office worker to replicate his desk at home. In today's society, there is more emphasis on flexible working hours and employees acting more as independent contractors. The big mainframe in the secret data center with difficult-to-use interfaces like text-based screens is fading fast and ISDN is helping to make all this possible.

ISDN requires a powered box that acts like a multiplexer on subscriber's end of the phone line. This is how the multiple services are provided on a single phone line connection to the switched public network, that is, your local phone company. This box is known as the terminal adapter. There is also a variety of termination devices for the network connection and line termination equipment for the telephone line. On the telephone company's end, there is exchange termination equipment to funnel the multiple services into the phone circuit that ends at the subscriber's phone number. There are two types of ISDN terminal equipment. Special ISDN terminals are terminal equipment type 1 (TE1). Older non-ISDN terminals that are identified as DTE and predate the current ISDN standards are known as terminal equipment type2 (TE2). TE1 is connected to ISDN with a twisted-pair cable having two pairs. TE2 connects to ISDN with an ISDN-compatible terminal adapter. The terminal adapter can be physically mounted inside the TE2 (non-ISDN terminal equipment) or it can be physically separate and standalone.

On the phone line side of the terminal adapter is a device known as the network termination. Network termination equipment is also provided in two types, network termination type 1 (NT1) and network termination type 2 (NT2). Both of these network termination types serve the function of connecting the phone company's two-wire phone line that comes to your location to the four-wire connection that ISDN requires. The point where the phone company ceases to be responsible and where you become responsible is called the demarcation point (also known as the demarc). This is the phone company's way of being able to say "Our stuff is working; it must be your problem."

When you begin to establish a connection with the major carriers, you will quickly learn that there is customer premise equipment (CPE). In North America, the network termination type 1 is a CPE device. That means the network termination is on your property. In the rest of the world, the network termination is part of the service provided by the phone company. The customer premise device concept is an outcome of the deregulation of the telephone industry. The network termination type 2 is a more complex piece of equipment which is needed for digital telephone switches.

Services provided by ISDN consist of a basic rate interface (BRI) and a primary rate interface (PRI).

The basic rate interface consists of two B channels and one D channel; this is known as (2B+D). The B channels are known as bearer channels for the transport of the user's data. The D channel is a data channel the phone company requires for control of the channel. The D channel does not carry anything that belongs to the user's data. BRI B channels operate at 64Kbps for the transport of data packets. The BRI D channel operates at 16Kbps. An ISDN connection can be established in 1 to 4 seconds as opposed to the 10 to 40 seconds that calling a phone number to establish a modem connection requires. It is the separate control path, the BRI D channel that does the connection and allows such high-speed interfacing.

The primary rate interface consists of 23 B channels and 1 D channel in North America and Japan. This is known as 23B+D, and it offers a total of 1.54Mbps of throughput for all 23 channels combined. In the PRI, the D channel runs at 64Kbps for channel control. In the rest of the world, the ISDN primary rate interface provides 30 B channels plus 1 D channel 30B+D, and it has a total speed of 2.048Mbps.

Until recently, equipment was not built to run at the 64Kbps rate of the B channels in ISDN. To match equipment such as serial ports on terminals or personal computers, you need a terminal adapter. The terminal adapter does just that, it adapts terminals. It performs a function known as rate adaptation to make the bit rates match. The terminal adapters in use today operate to two specifications, which are similar to the specifications you see on today's dial-up modems. In North America, the common standard is called V.120. In the European countries, an older standard V.110 is used. Both standards support synchronous and asynchronous data transmission.

Another quick of ISDN and the terminal adapter is a function called switched 56 dial-up. Switched 5 dial-up is a 56Kbps digital service that telephone companies use. You might be expecting to get the full 64Kbps throughput on your B channels and discover that the best you can attain is only 56Kbps. This type of rate adaptation happens when you connect through some of the smaller telephone carriers such as cooperatives in rural areas.

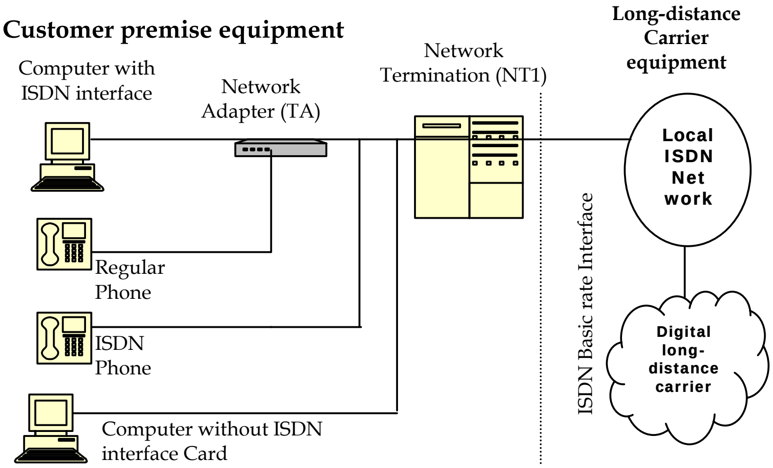
Customer premise equipment
Computer with ISDN interface
Network Adapter (TA)
Network Termination (NT1)
Long-distance Carrier equipment
Local ISDN Net work
Regular Phone
ISDN Phone
Computer without ISDN interface Card
ISDN Basic rate Interface
Digital long-distance carrier

You can dedicate your ISDN connection with the phone company to pure data and no voice. If you do, your telephone service will have to remain on its own telephone company connection separate from the ISDN connection. The two ISDN B channels are then combined into one whopping 128Kbps channel. This is called bandwidth-on-demand. The great thing about bandwidth-on-demand is that you can switch the capacity of the channel interface to combine or split the two B channels. You could use the connection with the phone company at the high rate for a full-color, full-motion digital video conference, and then switch it back to a split network to provide additional phone service or just two separate network interfaces with the outside world.

The availability of ISDN is still an issue. Depending on the phone company your site is served by, you might not be able to get ISDN. Even it is available from your phone company, the cables and drops reaching your site might not support the service. The only way to find out is to check with your phone company.

26.6 Wireless

Wireless applications have the advantage of allowing for rapid deployment of services to subscribers at a low cost. It also limits investment in customer-premise equipment (CPE) to those who have requested the service. However, this technology is not highly desirable for Internet access due to the lack of a return channel and/or low throughput.

26.7 Short Summary

Integrated Digital Services Network (ISDN)

ISDN is a technology designed for the public switched telephone network that allows low-cost communication in data, voice, graphics, and video. It is designed to be used over the existing copper local loop that connects the telephone company's central office to the home. If offered nationwide and subject to affordable tariff rates, narrowband ISDN can support digital service to the home and office without requiring the significantly greater time or expense of infrastructure conversion to broadband (ie., ATM).

ISDN's basic rate interface (BRI), the type that interests most consumers, gives two B channels, each with a data rate of 64 Kbps. The B channels are used for user information. In addition, a third channel, the D channel is used for call control running at 16 Kbps.

26.8 Brain Storm

1. Explain about broadband communication for the internet and intranet.
2. Explain the services of broadband communication
3. What is ISDN?
4. Explain the advantages of ISDN.
5. Explain about wireless.

Lecture 27

Virtual Reality Technology

Objectives

After completing this lesson, you should be able to do the following

- Discuss about the broadband carrier services.
- Describe about the frame relay.
- Describe about Asynchronous transfer mode.
- Describe about virtual reality technologies.
- Describe the evolving virtual reality applications.

Coverage Plan

Lecture 27

27.1 Snap Shot

Virtual Reality (VR) is a set of hardware/software-based applications now entering mainstream commercial application. Applications are appearing both on the Internet and in intranets. The technology is particularly well suited to applications that simulate the experience of actually "being at" the remote site, such as enhanced videoconferences, industrial and architectural walkthroughs, interactive travelogues of distant cities, house-hunting, product modeling, marketing, and training. Major commercial on-line service providers are either readying virtual reality on-line offerings, or are planning to do so.

Fortune 500 companies are finding savings, competitive advantages, or greater customer satisfaction – or all three – from virtual reality business applications. The most promising uses are in training and in design. Motorola and Nortel have applied the technology to train their staff and/or clients, and have documented significant savings compared to other methods of achieving the same goal.

27.2 Broadband Carrier Services for Intranets and for The Internet

Broadband carrier services for intranets and for the Internet are frame relay, ATM, Cable TV, economics approaches and direction.

When regulations restricting competition are relaxed, new service options will become available. As the public policy continues to enable more and more competition, the regulatory barriers that have kept cable and telephone companies out of each other's base businesses are surely going to fall.

The digital technology for these systems is still relatively undeveloped. They are expected to begin with analog video, limiting the number of channels to around 60. This system operates in much the same way as existing broadcast services: all the channels are transmitted down every cable. The system is relatively simple because services do not need to be "switched" at the local exchange to be sent to each customer.

27.3 Frame Relay and SMDS

Frame relay and switched multimegabit data service (SMDS) meet the requirements for multimedia application only in a borderline mode. They are both oriented to data traffic and could not support voice and video applications effectively. However , a service such as frame relay may be reasonable Internet access vehicle for a number of years.

27.4 ATM

ATM is a standards-based network technology designed for high-speed transmission of sound, images, and video over a single network. This networking technology is being widely deployed on public WANs, LANs, and, in some measure, at the desktop. Speeds today include 25, 45, 155, and 622 Mbps. ATM includes a switching mechanism that provides a means of switching high-capacity channels that have time-delivery constraints, such as video and voice signals, as well as supporting more bursty information flows with widely variable bit rates. Migration to broadband-based services is really being driven by the business sector. Today's users need more bandwidth and isochronous services over a packet-switched infrastructure in order to support new time-sensitive multimedia and videoconferencing applications. ATM is the solution because of its high speed, quality of service, connection on demand, scalability in deployment, and traffic-management features. It is applicable to both the Internet access and Internet backbone component; also, it can be utilized in corporate intranets.

ATM combines support of high bit rates and simplicity of circuit switching with the flexibility of packet switching. With switched virtual circuit service, users will be charged by the amount of bandwidth they use, and the length of time they use it instead of paying for a dedicated line that they may not be fully utilized.

ATM supports a number of service classes. The classes determine how the information will be converted into ATM format and how the cells will be treated in the network. Class A is for constant-bit-rate (CBR) services such as voice and video. Class B includes variable bit rate (VBR). Class C deals with connection-oriented services, and Class D with connectionless services. Class X is a pure cell relay service that requires no adaptation. In reality, only Class A and Class X have emerged commercially (even there, the support of circuit emulation is limited). Class D was slated to be used to support SMDS internetworking, but the commercial viability of this service is bleak. ATM utilizes a transfer mechanism that is independent of the type of service. However, provisions must be made to delivering at ATM services is simply to focus on the kind of cell-transport contract that is supported: CBR, VBR, unspecified bit rate (UBR), and available bit rate (ABR).

ATM uses the cell as its basic unit throughout the network. An ATM cell is a protocol data unit with a fixed length, a header that contains routing information to be used by the network switches, and an information field. The ATM standard specifies that all data be encapsulated within a 48-byte section; each cell also includes a 5-byte header section containing virtual circuit routing information. These fixed-length cells of 53 bytes require less processing overhead to switch them, so they can be routed at much higher speeds. Due to the fixed format of the two cell fields, switching can be done in hardware, effectively reducing the intra switching processing time. The network uses only the header information, regardless of the payload, except in case of service interworking.

ATM uses virtual channels that have a prespecified rate, based on either administrative-based provisioning, or on a per-call signaling basis (specifically for switched virtual connections). Both a maximum rate is specified (this bit rate is less than or equal to the physical bit rate of the user-to-network interface), and an average rate (called sustainable cell rate, or SCR). In effect, this mechanism supports bandwidth on demand. This much-abused term simply means that the user can send cells at a time-averaged rate of SCR, but can occasionally “burst” up to a time-averaged rate of PCR.

Cells are identified as belonging to a virtual channel by information in the header. This information can change as it flows from switch to switch because this information is only a relative connection pointer. However, ATM is connection-oriented, which means that a connection must be set up before information is transferred. The switch analyzes the header information and switches the cell from the incoming multiplexed stream to the outgoing multiplexed stream. The switch also maintains the integrity of the cell sequence.

Independent of transmission speed, ATM is also unconstrained by the physical media of the network. It can operate over twisted-pair coaxial and fiber and also by means of wireless technologies. Public carriers offering ATM-based services include Teleport Communications Group, AT&T, Sprint, MCI, BellSouth, Pacific Bell, Ameritech, and Wiltel.

ATM provides for an integrated, service-independent network that can transport the services of today as well as the services of the future without requiring a new infrastructure.

Organizations are now planning to use ATM for their enterprise networks. Service providers offering various services are connected to the access subnetwork; a core ATM is used for citywide or region wide aggregation. In ATM, just like in ISDN and frame relay, bi-directional logical channels are used for user-to-network signaling and control.

In the Internet application, ATM can be used as a customer-visible access technology or as a backbone NSP technology to interconnect the provider's routers. The computer systems and the Web servers can be equipped with ATM interfaces compatible with those available on ATM switches. The belief is that the backbone network of the Internet should ultimately be ATM-based. If the access network is also ATM-based, the resulting end-to-end network can support high-speed communications. ATM is currently available at 45-Mbps to OC-3 speeds. In the future, ATM is expected to run at OC-12, or even higher speeds.

Various high-speed user services can use ATM for hardware-level, high-speed switching. This enables ATM to be a switching technology for MultiProtocol applications. An ATM interface adapter card can provide MultiProtocol access by converting traffic coming from different end systems into ATM format.

Interworking units (IWUs) could be used to link non-ATM-access networks to the core ATM network. The IWU terminates the ATM virtual channel by providing an appropriate protocol peer. An example is frame-relay-to-ATM or Ethernet-to-ATM interworking. The processing required by the IWU compared to the hardware-based cell routing will affect the throughput of the network.

The user-to-network interface can be implemented according to the needs of the user and the primary function of the user's network. Interfaces can be developed for physical interfaces of different rates and formats.

There is also a need for standards to specify the interface between ATM and the customer-premise equipment (CPE), as well as the residential broadband access network. Standards are also needed for the internal home bus of the future where PCs and set-top boxes can communicate. ATM must have different-speed interfaces for access lines slower than SONET rates and must also overcome the problems of technology costs and interoperability problems among ATM products.

The business created by broadband/ATM networking by the end of the decade could be approximately $3 million. David Dorman, chairman and CEO of Pacific Bell, predicted the information superhighway “media frenzy” will only intensify in the years ahead. A ubiquitous, nationwide, broadband network, he said, will cost upward of $250 billion but will create $3 trillion ($3000 billion) in commerce by the end of the decade. Dorman predicated that telecommuting would become the first commercial hit of broadband networks, although others disagree.

Companies manufacturing and designing ATM systems are trying to make the transition to higher-speed networks as painless as possible. In an enterprise context, the corporate information data warehouses are probably the best locations to upgrade. This is where a major portion of network traffic occurs as users access the warehouse seeking information. If the throughput is increased, then productivity is increased.

27.5 Virtual Reality Technologies

Over the past five years there has been increased commercial interest in VR. Some think that VR will have as great an impact on society as television. At this juncture, however, VR is still an evolving technology in the early stages of development. The term virtual reality was coined by Jaron Lanier, the founder of the first commercial VR company (VPL Incorporated), in 1989.

VR covers a range of definitions. A definition advanced by Shemen and Judkins, termed the “five i’s of VR”, applies to the characteristics it possesses. The “five i’s” are intensive, interactive, immersive, illustrative, and intuitive.

Intensive. In a VR setting, the user should be concentrating on multiple, vital information, to which the user will respond.

Interactive. In a VR setting, the user and the computer act reciprocally over an appropriate interface.

Immersive. VR should deeply involve or absorb the user.

Illustrative. VR should offer information in a clear, descriptive, and illuminating way.

Intuitive. Information should be easily perceived.

Some see VR as the fusion of three other technologies – telephone, television, and video game - to arrive at a technology that is better than the linear sum of the three. In VR, there is strong emphasis on communicating with virtual worlds (VW) via a more intuitive and more sense inclusive human-computer interface (HCI) than would otherwise be possible. The goal is to make the technology assessable to non-computer-language-literate users. Until the present, the input/output to a computerized system has been restricted to a fairly unintuitive textual user interface (TUI) or to an improved, but still limited, mouse-based GUI. With VR, the interface is the human's hands, legs, eyes, or body; a mouse can also be used.

In a VR environment, what happens in the virtual world depends on the user. The user's inclusion in the virtual world and the ability to influence what happens in it are key advancements of VR compared with earlier technologies. VR facilitates spatial navigation, a method of using graphical symbols to shop on-line and to travel through cyberspace.

It is possible to trace VR concepts back 70 years, to work in vehicle simulation, teleoperation, three-sided screens in "cinerama", and head-mounted CCTV- based teleoperation work at Philco and Argonne National Laboratory. Flight simulations conducted by NASA in the mid-1960s started to give shape to VR as we know it today. Computer-generated synthetic displays for virtual environments started to appear in late 1960. work continued in the 1970s, when the term artificial reality was coined. In 1989, the term VR was coined to refer to all virtual environments. More work continued into the 1990s, with added emphasis on technology and applications. As early as 1990, you could purchase $90 gloves, although these had no tactile feedback. High-end systems used by organizations such as NASA cost millions of dollars.

27.6 Evolving Virtual Reality Applications

Market Status

According to The Perth Institute (Hawley, Pennsylvania), the VR market will grow from more than $100 million in 1994 to about $6 billion by 1999. Businesses applications will be a sizable fraction of that total. Companies such as Silicon Graphics and Hewlett-Packard are expected to be big beneficiaries. Other analysts have advanced a more conservative $1 billion forecast.

There has been a large amount of media hyperbole about the technology. Fortunately, VR is a lot closer to reality now than just a while ago, but widespread introduction is not going to happen overnight. Researchers expect that it will be 5 to 10 years before the technology actually does what the press currently claims it can. Cost remains a factor, but prices are dropping. Computing power is getting cheap. You can buy a VR workstation, software, and peripherals for as little as $75,000. For some applications these prices may well are justifies: for example, if a company needs to redesign the floor space every 12 months or so, then the investment in a VR-based layout tool may well be worthwhile. Low-end VR software also runs on high-performance PCs.

In spite of the limitations, VR has already been commercially successful for computer games and arcades. The main market for VR has so far been in the area of entertainment; however, new applications are evolving. For example, in Japan, people can shop for new kitchens. Applications to architectural designs are being developed. For rapid prototyping, VR is well suited. Training applications are appearing. Military simulations also continue to be a sizable market.

There have been documented applications in the following fields:

+ Banking, financial trading, investments
+ Art, science
+ Engineering
+ Games, amusement, entertainment
+ Building
+ Marketing (including shopping)
+ Military
+ Education
+ Shopping

There is a growing community of integrators that can assist users in deploying the technology.

Examples of applications

Some apply Fubini's law to VR technology, which describes a four-step evolution:

1. People initially use the technology to do what they do now, but faster.

2. They gradually begin to use technology to do new things.
3. New things change lifestyles and work styles.
4. New lifestyles and work styles change society… and eventually change technology, introducing new technology, at which point the cycle starts over.

Initially, VR can be seen as a natural progression of 3-D computer generated models. Applications in development and design, for example, could include a more complete analysis of the design due to viewing it from any angle, as well as the communication of designs to clients using 3-D walkthroughs, which might have voice commands, sounds, and touch. As a future application, clients may walk to an empty building site and put on a pair of glasses that lets them see their proposed designs, as well as the site, and move things around to meet their architectural expectations. Automotive companies in Detroit, for example can access real-time feedback from design consultants in Europe. By linking geographically dispersed CAVEs, designers can simultaneously view a virtual life-size car model and interactively move or restyle body parts so participants at other sites get a firsthand look at proposed alterations.

Some specific examples of application follow (these are often called "VR Demonstration Prototypes")

∗ Furniture companies in Europe, Japan, and the United States have used VR as a marketing and presentation tool. The VPL Incorporated VR system (introduced in 1989) was used commercially in Matsushita's Shinjuku store to sell kitchens.

∗ At the University of North Caroline, VR technology has been used for the architectural design/walkthrough of an addition to a church.

∗ At the University of North Carolina, VR technology has been used to design the University's new computer science building (Sitterson Hall)

∗ Architectural prototypes at the ACG Center, although still expensive, are nonetheless proofs of concept.

∗ Intel and Sensed cosponsored "Designing a Virtual home", a two day demonstration of a prototype architectural application. This two person "VR station" represented the architect working with the client. It included helmets, a joystick, and a pointing wand used to move surfaces and to change their appearance.

∗ Work is under way to make virtual explorations of long-lost monuments and buildings.

∗ EDS has opened the Detroit Virtual Reality Center to generate demand for virtual reality solutions. About 2600 business people from more than 350 companies have seen the demos as of late 1995. The company believes that the opening of the center signifies a transition from primarily entertainment and government use to a broader VR use. The company has a VR theater and a VR CAVE. It develops applications, lets clients use the facilities, consults with companies that want to set up their own operations, and works with partners to improve hardware and software.

∗ Albert Kahn is an architectural firm that does $600 million a year in construction design for factories, hospitals, auto makers, R&D centers, and public institutions. The company has used VR for initial design work. They want to give the customer a sense and feel of different options.

∗ GM R&D Center uses VR for model design. CAD data is fed into VR programs using homegrown software packages and a VR CAVE. The CAVE is a barren room with a car seat in the middle. The walls of the CAVE are screens that, with the help of shutter glasses flickering on and off every 1/60 of a second, provide the illusion of a 3-D environment. The company uses it to market-test new design without having to spend time and money building prototypes out of wood.

∗ Newbridge showcased VR, multimedia, and other broadband applications at Interop 1996.

∗ Nortel joined Superscape to train telephone attendants using VR methods.

∗ Three-dimensional VR tours use real 360° photography to show travel destinations.

∗ VR videoconferencing

∗ The Berlin Multi-media Arts and Communications Center also has used the VPL Incorporated VR system

- ∗ Computer graphics have been developed using VR techniques by the Advanced Computer Graphics Center, Royal Melbourne Institute of Technology (Australia).

27.7 Short Summary

When we want to know about virtual reality we must know the following five 'i's'. They are intensive, interactive, immersive, illustrative and intuitive.

Intensive. In a VR setting, the user should be concentrating on multiple, vital information, to which the user will respond.

Interactive. In a VR setting, the user and the computer act reciprocally over an appropriate interface.

Immersive. VR should deeply involve or absorb the user.

Illustrative. VR should offer information in a clear, descriptive, and illuminating way.

Intuitive. Information should be easily perceived.

27.8 Brain Storm

1. What is virtual reality?
2. What is ATM?
3. Explain about frame relay.
4. Explain the services of Broadband carrier.

MANONMANIAM SUNDARANAR UNIVERSITY

Centre for Information Technology and Engineering

Tirunelveli

Syllabus for MS(IT&EC)

2.3 Internet and Intranet Administration

Lecture – 1

Internet - Linking to the Internet - Internet address – Internet tools-information retrieval tools – communication tools – multimedia information tools – information search tools

Lecture - 2

Intranet – Intranets Vs Groupware – Intranet hardware - Intranet software – Intranet services – Extranet

Lecture - 3

Internet server – web protocols – browser – Futures of Internet & Intranet application

Lecture – 4

The evolution of TCP/ IP –The rise of the Internet – TCP/IP protocol architecture – TCP/IP core protocols – TCP / IP application interfaces

Lecture - 5

IP addressing – Basic address scheme - Address classes - dotted-decimal notation – Networking basics – Host restrictions - Subnets – domain Name system

Lecture - 6

Subnet mask – Subnetting – Variable length subnetting – Supernetting and classless interdomain routing - Public & private addresses

Lecture - 7

TCP / IP operation & applications – Internet protocol- The IP header –Address resolution – TCP source and destination port – FTP – Telnet –HTTP

Lecture - 8

Network infrastructure – component of network Infrastructure –Ethernet – Token ring – FDDI – fast Ethernet – asymmetric digital subscriber line -virtual local area network

Lecture - 9

ISP – Need of an ISP - Basic Connections –Bulletin board systems- not-so-basic connection – need of Not-so-basic connection

Lecture – 10

Services of Internet – Domain name services – web site development and specialized browsers – account management – network operations center - connection type – Security

Lecture – 11

Virtual server – evolution of ISP – quality of service – ISP backbone – value added services

Lecture – 12

Internet services – secure services – web client security issues – web server security issues- remote access - Real time conferences services – administrative services

Lecture – 13

Router technology – Network fundamentals – Internet routing - routing protocols- routing software – Mobile routing

Lecture – 14

Web roots - Web servers – Transport issues - web access – connection modes - publishing web server

Lecture – 15

Internet security – firewalls –securing internet applications- virus checking and scanning – security administration

Lecture – 16

Firewall – firewall technologies –packet filtering – firewall architecture – firewall design

Lecture – 17

Proxy server – proxy services - why proxy? – How proxying works? – Proxy server Terminology – proxying without a proxy server

Lecture – 18

Packet filtering – filtering by address –filtering by service –packet filtering router – rules of filtering

Lecture – 19

Internet application – Tools & utilities – communication –email – internet relay chat – www – URL - browser

Lecture – 20

Intranetting Vs Internetting – Do you need to Intranet? – Increase communication – User friendly - Benefits of Intranet

Lecture – 21

Internet & Intranet web site development – Intranet application – organizational web pages – Distributed strategy- web master

Lecture – 22

IIS - Microsoft Internet assistants – intranet search engine – Alta vista – Excite – Intranet toolbox – Platform independent pages

Lecture – 23

Wires and fibers – remote access considerations – network management – resource and device management – cross platform management –web based network management

Lecture – 24

Definition of online services – technology trends - profiles of major online service provider – America online – prodigy – Microsoft network

Lecture – 25

Shareware, freeware & crippleware – other Internet server – distributed management – DMI service provider

Lecture – 26

Broadband communication for the Internet & Intranet – Services & Requirement – network architecture supporting broadband – ISDN – wireless

Lecture – 27

Broadband carrier services – frame relay – ATM - virtual reality technologies – Evolving Virtual reality applications.

Best of Luck