The Net and the Internet

Table of Contents

Network

In information technology, a network is a series of points or nodes interconnected by communication paths. Networks can interconnect with other networks and contain subnetworks.

The most common topologies or general configurations of networks include the bus, star, and ring topologies. Networks can also be characterized in terms of spatial distance as local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs).

A given network can also be characterized by the type of data transmission technology in use on it (for example, a TCP/IP or SNA network); by whether it carries voice, data, or both kinds of signals; by who can use the network (public or private); by the usual nature of its connections (dial-up or switched, dedicated or nonswitched, or virtual connections); and by the types of physical links (for example, optical fiber, coaxial cable, and copper wire). Large telephone networks and networks using their infrastructure (such as the Internet) have sharing and exchange arrangements with other companies so that larger networks are created.

File sharing

File sharing is the public or private sharing of computer data or space in a network with various levels of access privilege. While files can easily be shared outside a network (for example, simply by handing or mailing someone your file on a diskette), the term file sharing almost always means sharing files in a network, even if in a small local area network. File sharing allows a number of people to use the same file or file by some combination of being able to read or view it, write to or modify it, copy it, or print it. Typically, a file sharing system has one or more administrators. Users may all have the same or may have different levels of access privilege. File sharing can also mean having an allocated amount of personal file storage in a common file system.

File sharing has been a feature of mainframe and multi-user computer systems for many years. With the advent of the Internet, a file transfer system called the File Transfer Protocol (FTP) has become widely-used. FTP can be used to access (read and possibly write to) files shared among a particular set of users with a password to gain access to files shared from an FTP server site. Many FTP sites offer public file sharing or at least the ability to view or copy files by downloading them, using a public password (which happens to be "anonymous"). Most Web site developers use FTP to upload new or revised Web files to a Web server, and indeed the World Wide Web itself can be thought of as large-scale file sharing in which requested pages or files are constantly being downloaded or copied down to the Web user.

More usually, however, file sharing implies a system in which users write to as well as read files or in which users are allotted some amount of space for personal files on a common server, giving access to other users as they see fit. The latter kind of file sharing is common in schools and universities. File sharing can be viewed as part of file systems and their management.

Any multi-user operating system will provide some form of file sharing. Among the best known network file systems is (not surprisingly) the Network File System (NFS). Originally developed by Sun Microsystems for its UNIX-based systems, it lets you read and, assuming you have permission, write to sharable files as though they were on your own personal computer. Files can also be shared in file systems distributed over different points in a network. File sharing is involved in groupware and a number of other types of applications.

LAN (local area network)

A LAN is a network of interconnected workstations sharing the resources of a single processor or server within a relatively small geographic area. Typically, this might be within the area of a small office building. However, FDDI extends a local area network over a much wider area. Usually, the server has applications and data storage that are shared in common by multiple workstation users. A local area network may serve as few as four or five users or, in the case of FDDI, may serve several thousand.

The main LAN technologies are:

- Ethernet
- Token ring
- ARCNET
- FDDI (Fiber Distributed Data Interface)

Typically, a suite of application programs can be kept on the LAN server. Users who need an application frequently can download it once and then run it from their local hard disk. Users can order printing and other services as needed through applications run on the LAN server. A user can share files with others at the LAN server; read and write access is maintained by a LAN administrator.

A LAN server may also be used as a Web server if safeguards are taken to secure internal applications and data from outside access.

Ethernet

Ethernet is the most widely-installed local area network technology. Now specified in a standard, IEEE 802.3, Ethernet was originally developed by Xerox and then developed further by Xerox, DEC, and Intel. An Ethernet LAN typically uses coaxial cable or special grades of twisted pair wires. The most commonly installed Ethernet systems are called 10BASE-T and provide transmission speeds up to 10 Mbps. Devices are connected to the cable and compete for access using a Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol.

Fast Ethernet or 100BASE-T10 provides transmission speeds up to 100 megabits per second and is typically used for LAN backbone systems, supporting workstations with 10BASE-T cards. Gigabit Ethernet provides an even higher level of backbone support at 1000 megabits per second (1 gigabit or 1 billion bits per second).

Token ring network

A token ring network is a type of local area network. In a token ring network, all workstations are connected in a ring or star topology and a token-passing scheme is used to prevent the collision between two workstations who want to send messages at the same time. Here, very briefly, is how it works:

1. Empty information frames are continuously circulated on the ring.
2. When a workstation wants to send a message, it inserts a token in an empty frame (this may consist of simply changing a 0 to a 1 in the token bit part of the frame) and inserts a message and a destination identifier in the frame.
3. The frame is then examined by each successive workstation. If the workstation sees that it is the destination for the message, it copies the message from the frame and changes the token back to 0.
4. When the frame gets back to the originator, it sees that the token has been changed to 0 and that the message has been copied and received. It removes the message from the frame.
5. The frame continues to circulate as an "empty" frame, ready to be taken by a workstation when it has a message to send.

The token scheme can also be used with bus topology LANs.

The standard for the token ring protocol is IEEE-802.5. Perhaps the best known implementation of this protocal is the IBM Token Ring Network, which is widely used today in many corporate LANs. The Fiber Distributed-Data Interface (FDDI) also uses a token ring protocol.

Arcnet (Attached Resource Computer Network)

Arcnet is a widely-installed local area network technology from the Datapoint Corporation, the originator of the local area network. Arcnet uses a token-bus scheme for managing line sharing among the workstations and other devices connected on the LAN. The LAN server continuously circulates empty message frames on a bus (a line in which every message goes through every device on the line and a device uses only those with its address). When a device wants to send a message, it inserts a "token" (this can be as simple as setting a token bit to 1) in an empty frame in which it also inserts the message. When the destination device or LAN server reads the message, it resets the token to 0 so that the frame can be reused by any other device. The scheme is very efficient when traffic increases since all devices are afforded the same opportunity to use the shared network.

Arcnet can use coaxial cable or fiber optic lines. Cable lengths can be up to 2,000 feet per segment with a total network span of 20,000 feet without bandwidth loss. Arcnet's bandwidth or information flow capacity is 2.5 Mbps. Of the four major LAN technologies (which also include Ethernet, token ring, and FDDI), Arcnet is said to be the least expensive to install.

FDDI (Fiber Distributed-Data Interface)

FDDI is a standard for data transmission on fiber optic lines in a local area network that can extend in range up to 200 km (124 miles). The FDDI protocol is based on the token ring protocol. In addition to being large geographically, an FDDI local area network can support thousands of users.

An FDDI network contains two token rings, one for possible backup in case the primary ring fails. The primary ring offers up to 100 Mbps capacity. If the secondary ring is not needed for backup, it can also carry data, extending capacity to 200 Mbps. The single ring can extend the maximum distance; a dual ring can extend 100 km (62 miles).

FDDI is a product of American National Standards Committee X3-T9 and conforms to the open system interconnect (OSI) model of functional layering. It can be used to interconnect LANs using other protocols. FDDI-II is a version of FDDI that adds the capability to add circuit-switched service to the network so that voice signals can also be handled. Work is underway to connect FDDI networks to the developing Synchronous Optical Network SONET which in turn is part of broadband ISDN.

WAN (wide area network)

A WAN (wide area network) is a geographically dispersed telecommunications network and the term distinguishes a broader telecommunication structure from a local area network (LAN). A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared user) networks. An intermediate form of network in terms of geography is a metropolitan area network (MAN).

MAN (metropolitan area network)

A MAN (metropolitan area network) is a network that interconnects users with computer resources in a geographic area or region larger than that covered by even a large local area network (LAN) but smaller than the area covered by a wide area network (WAN). The term is applied to the interconnection of networks in a city into a single larger network (which may then also offer efficient connection to a wide area network). It is also used to mean the interconnection of several local area networks by bridging them with backbone lines. The latter usage is also sometimes referred to as a campus network.

Examples of metropolitan area networks of various sizes can be found in the metropolitan areas of London, England; Lodz, Poland; and Geneva, Switzerland. Large universities also sometimes use the term to describe their networks.

SNA (Systems Network Architecture)

SNA is a proprietary IBM architecture and set of implementing products for network computing within an enterprise. It existed prior to and became part of IBM's Systems Application Architecture (SAA) and it is currently part of IBM's Open Blueprint. With the advent of multi-enterprise network computing, the Internet, and the de facto open network architecture of TCP/IP, IBM is finding ways to combine its own SNA within the enterprise with TCP/IP for applications in the larger network.

SNA itself contains several functional layers and includes an application program interface called the Virtual Telecommunications Access Method (VTAM), a communications protocol for the exchange of control information and data, and a data link layer, Synchronous Data Link Control (SDLC). SNA includes the concepts of nodes that can contain both physical units that provide certain setup functions and logical units, each associated with a particular network transaction.

DSN (Deep Space Network)

The Deep Space Network (DSN) is a sophisticated data communications system used by NASA (the U.S. National Aeronautics and Space Administration) in conjunction with manned and unmanned space missions. The DSN is also used by radio astronomers.

The main terminal of the DSN is located at JPL (Jet Propulsion Laboratory) headquarters in Pasadena, California. There are three primary antennas, spaced equally on a great circle that slants around the world. All three are large paraboloid (dish) antennas that can be used for transmitting and receiving signals over a wide range of radio frequencies. One antenna is located in California, another is in Spain, and another is in Australia. The antennas are located in such a way that all existing operational spacecraft can be monitored and controlled, and communications maintained with them, almost 100 percent of the time. This is true of both earth-orbiting satellites and interplanetary space vehicles.

Signals transmitted and received by DSN equipment include satellite control and telemetry, e-mail (including text, graphics, video, programs, and sound attachments), communications with the Space Shuttles, and radio-frequency emanations from distant celestial objects.

The NASA Deep Space Network - or DSN - is an international network of antennas that supports interplanetary spacecraft missions and radio and radar astronomy observations for the exploration of the solar system and the universe. The network also supports selected Earth-orbiting missions.

The DSN currently consists of three deep-space communications facilities placed approximately 120 degrees apart around the world: at Goldstone, in California's Mojave Desert; near Madrid, Spain; and near Canberra, Australia. This strategic placement permits constant observation of spacecraft as the Earth rotates, and helps to make the DSN the largest and most sensitive scientific telecommunications system in the world. NASA's scientific investigation of the Solar System is being accomplished mainly through the use of unmanned automated spacecraft. The DSN provides the vital two-way communications link that guides and controls these planetary explorers, and brings back the images and new scientific information they collect. All DSN antennas are steerable, high-gain, parabolic reflector antennas.

The antennas and data delivery systems make it possible to:
- Acquire telemetry data from spacecraft.
- Transmit commands to spacecraft.
- Track spacecraft position and velocity.
- Perform very-long-baseline interferometry observations.
- Measure variations in radio waves for radio science experiments.
- Gather science data.
- Monitor and control the performance of the network.

The network is a facility of NASA, and is managed and operated for NASA by the Jet Propulsion Laboratory. The Telecommunications and Mission Operations Directorate (TMOD) manages the program within JPL.

Fiber-Optic

Fiber optic (or "optical fiber") refers to the medium and the technology associated with the transmission of information as light impulses along a glass or plastic wire or fiber. Fiber optic wire carries much more information than conventional copper wire and is far less subject to electromagnetic interference. Most telephone company long-distance lines are now fiber optic.

Transmission on fiber optic wire requires repeating at distance intervals. The glass fiber requires more protection within an outer cable than copper. For these reasons and because the installation of any new wiring is labor-intensive, few communities yet have fiber optic wires or cables from the phone company's branch office to local customers (known as local loop).

Digital

Digital describes electronic technology that generates, stores, and processes data in terms of two states: positive and non-positive. Positive is expressed or represented by the number 1 and non-positive by the number 0. Thus, data transmitted or stored with digital technology is expressed as a string of 0's and 1's. Each of these state digits is referred to as a bit (and a string of bits that a computer can address individually as a group is a byte).

Prior to digital technology, electronic transmission was limited to analog technology, which conveys data as electronic signals of varying frequency or amplitude that are added to carrier waves of a given frequency. Broadcast and phone transmission has conventionally used analog technology.

Digital technology is primarily used with new physical communications media, such as satellite and fiber optic transmission. A modem is used to convert the digital information in your computer to analog signals for your phone line and to convert analog phone signals to digital information for your computer.

Analog

Analog technology refers to electronic transmission accomplished by adding signals of varying frequency or amplitude to carrier waves of a given frequency of alternating electromagnetic current. Broadcast and phone transmission have conventionally used analog technology.

Analog also connotes any fluctuating, evolving, or continually changing process. Analog is usually represented as a series of sine waves. The term originated because the modulation of the carrier wave is analogous to the fluctuations of the voice itself.

A modem is used to convert the digital information in your computer to analog signals for your phone line and to convert analog phone signals to digital information for your computer.

Modem (modulator/demodulator)

A modem modulates outgoing digital signals from a computer or other digital device to analog signals for a conventional copper twisted-pair telephone line and demodulates the incoming analog signal and converts it to a digital signal for the digital device.

In recent years, the 2400 bps modem that could carry e-mail has become obsolete. 14.4 Kbps and 28.8 Kbps modems were temporary landing places on the way to the much higher bandwidth devices and carriers of tomorrow. From early 1998, most new personal computers came with 56 Kbps modems. By comparison, using a digital ISDN adapter instead of a conventional modem, the same telephone wire can now carry up to 128 Kbps. With Digital Subscriber Line (DSL) systems, now being deployed in a number of communities, bandwidth on twisted-pair can be in the megabit range.

V.xx telephone network standards

The V Series Recommendations from the ITU-T are summarized in the table below. They include the most commonly used modem standards and other telephone network standards. Prior to the ITU-T standards, the American Telephone and Telegraph Company and the Bell System offered its own standards (Bell 103 and Bell 212A) at very low transfer rates. Another set of standards, the Microcom Networking Protocol, or MNP Class 1 through Class 10 (there is no Class 8), has gained some currency, but the development of an international set of standards means these will most likely prevail and continue to be extended. (Some modems offer both MNP and ITU-T standards.)

In general, when modems "handshake," they agree on the highest standard transfer rate that both can achieve.

Beginning with V.22bis, ITU-T transfer rates increase in 2400 bps multiples. (bis stands for "second version." terbo stands for "third version.")

Standard	Meaning
V.22	Provides 1200 bits per second at 600 baud (state changes per second)
V.22bis	The first true world standard, it allows 2400 bits per second at 600 baud
V.32	Provides 4800 and 9600 bits per second at 2400 baud
V.32bis	Provides 14,400 bits per second or fallback to 12,000, 9600, 7200, and 4800 bits per second
V.32terbo	Provides 19,200 bits per second or fallback to 12,000, 9600, 7200, and 4800 bits per second; can operate at higher data rates with compression; was not a CCITT/ITU standard
V.34	Provides 28,800 bits per second or fallback to 24,000 and 19,200 bits per second and backwards compatility with V.32 and V.32bis
V.34bis	Provides up to 33,600 bits per second or fallback to 31,200 or V.34 transfer rates
V.35	The trunk interface between a network access device and a packet network at data rates greater than 19.2 Kbps. V.35 may use the bandwidths of several telephone circuits as a group. There are V.35 Gender Changers and Adapters.
V.42	Same transfer rate as V.32, V.32bis, and other standards but with better error correction and therefore more reliable
V.90	Provides up to 56,000 bits per second downstream (but in practice somewhat less). Derived from the x2 technology of 3Com (US Robotics) and Rockwell's K56flex technology.

An industry standard, ISDN uses digitally-encoded methods on phone lines to provide transfer rates up to 128,000 bits per second. Another technology, DSL, provides even faster transfer rates.

x2 (56 Kbps downstream)

x2 is a technology from US Robotics (now 3Com) for the downstream transmission of data over ordinary phone lines at 56 Kbps (thousands of bits per second). The 56 Kbps speed is achieved in the downstream direction only (to your home or business). Upstream speed is at the regular maximum speed of 33.6 Kbps. (The actual achieved downstream speed is reported by users to be about 53 Kbps.) x2 provided input to and has been replaced by the V.90 ITU standard.

56 Kbps technologies exploit the fact that most telephone company offices are interconnected with digital lines. Assuming your Internet connection provider has a digital connection to its telephone company office, the downstream traffic from your local Internet access provider can use a new transmission technique on your regular twisted-pair phone line that bypasses the usual digital-to-analog conversion. A V.90-equipped modem doesn't need to demodulate the downstream data. Instead, it decodes a stream of multi-bit voltage pulses generated as though the line was equipped for digital information. (Upstream data still requires digital-to-analog modulation.)

Unlike ISDN, the V.90 technology does not require any additional installation or extra charges from your local phone company. On the other hand, the maximum transmission speed of ISDN is twice that of V.90 at 128 Kbps. You also have the flexibility of combining digital and voice transmission on the same line.

ISDN (Integrated Services Digital Network)

Integrated Services Digital Network (ISDN) is a set of CCITT/ITU standards for digital transmission over ordinary telephone copper wire as well as over other media. Home and business users who install ISDN adapters (in place of their modems) can see highly-graphic Web pages arriving very quickly (up to 128 Kbps). ISDN requires adapters at both ends of the transmission so your access provider also needs an ISDN adapter. ISDN is generally available from your phone company in most urban areas in the United States and Europe.

There are two levels of service: the Basic Rate Interface (BRI), intended for the home and small enterprise, and the Primary Rate Interface (PRI), for larger users. Both rates include a number of B (bearer) channels and a D (delta) channel. The B channels carry data, voice, and other services. The D channel carries control and signaling information.

The Basic Rate Interface consists of two 64 Kbps B channels and one 16 Kbps D channel. Thus, a Basic Rate user can have up to 128 Kbps service. The Primary Rate consists of 23 B channels and one 64 Kpbs D channel in the United States or 30 B channels and 1 D channel in Europe.

The typical cost for Basic Rate usage in a city like Kingston, New York is about $125 for phone company installation, $300 for the ISDN adapter, and an extra $20 a month for a line that supports ISDN.

Integrated Services Digital Network in concept is the integration of both analog or voice data together with digital data over the same network. Although the ISDN you can install is integrating these on a medium designed for analog transmission, broadband ISDN (BISDN) will extend the integration of both services throughout the rest of the end-to-end path using fiber optic and radio media. Broadband ISDN will encompass frame relay service for high-speed data that can be sent in large bursts, the Fiber Distributed-Data Interface (FDDI), and the Synchronous Opical Network (SONET). BISDN will support transmission from 2 Mbps up to much higher, but as yet unspecified, rates.

Bits Per Second (bps)

In data communications, (abbreviated bps) is a common measure of data speed for computer modems and transmission carriers. As the term implies, the speed in bps is equal to the number of bits transmitted or received each second. The duration d of a data bit, in seconds, is inversely proportional to the digital transmission speed s in bps: d = 1/s

Larger units are sometimes used to denote high data speeds. One kilobit per second (abbreviated kbps) is equal to 1,000 bps. One megabit per second (Mbps) is equal to 1,000,000 bps or 1,000 kbps.

Computer modems for twisted-pair telephone lines usually operate at speeds between 14.4 and 57.6 kbps. The most common speeds are 28.8 and 33.6 kbps. So-called "cable modems," designed for use with TV cable networks, can operate at more than 100 kbps. Fiberoptic modems are the fastest of all; they can send and receive data at many Mbps.

The bandwidth of a signal depends on the speed in bps. With some exceptions, the higher the bps number, the greater is the nominal signal bandwidth. (Speed and bandwidth are, however, not the same thing.) Bandwidth is measured in standard frequency units of kilohertz (khz) or megahertz (MHz).

Data speed is sometimes specified in terms of baud, which is a measure of the number of times a digital signal changes state in one second. Baud, sometimes called the "baud rate," is almost always a lower figure than bps for a given digital signal. The terms are often used interchangeably, even though they do not refer to the same thing. If you hear that a computer modem can function at "33,600 baud" or "33.6 kilobaud," you can be reasonably sure that the term is being misused, and the figures actually indicate bps.

Baud

Baud was the prevalent measure for data transmission speed until replaced by a more accurate term, bps (bits per second). One baud is one electronic state change per second. Since a single state change can involve more than a single bit of data, the bps unit of measurement has replaced it as a better expression of data transmission speed.

The measure was named after a French engineer, Jean-Maurice-Emile Baudot. It was first used to measure the speed of telegraph transmissions.

Dial-up

Dial-up pertains to a telephone connection in a system of many lines shared by many users. A dial-up connection is established and maintained for a limited time duration. The alternative is a dedicated connection, which is continuously in place. Dial-up lines are sometimes called switched lines and dedicated lines are called nonswitched lines. A dedicated line is often a leased line that is rented from a telephone company.

A dial-up connection can be initiated manually or automatically by your computer's modem or other device.

Dedicated line

A dedicated line is a telecommunications path between two points that is available 24 hours a day for use by a designated user (individual or company). It is not shared in common among multiple users as dial-up lines are. A dedicated line can be a physical path owned by the user or rented from a telephone company, in which case it is called a leased line. A synonym is nonswitched line (as opposed to a switched or dial-up line).

Leased line

A leased line is a telephone line that has been leased for private use. In some contexts, it's called a dedicated line. A leased line is usually contrasted with a switched line or dial-up line.

Typically, large companies rent leased lines from the telephone message carriers (such as AT&T) to interconnect different geographic locations in their company. The alternative is to buy and maintain their own private lines or, increasingly perhaps, to use the public switched lines with secure message protocols. (This is called tunneling.)

Webification

Webification (sometimes seen with a lower case w) is the act of converting some information content from its original format into a format capable of being displayed on the World Wide Web. Some conversion examples are:

- A Postscript source file or ASCII text to a Hypertext Markup Language (HTML) file
- A Microsoft Word document to HTML (sometimes referred to as "DOC to HTML"). More recent versions of Microsoft Word include this capability.
- Hard-copy print publication pages into files in the Portable Document Format (PDF) for viewing on the Web with Adobe's downloadable Acrobat viewer
- A Lotus Notes database to HTML files
- An image in a scanned-in or other format to a Web-ready image, either a GIF or a JPEG file
- A speech or interview into a file in the RealAudio format for playing as streaming audio on the Web
- A video tape recording into a streaming video file

Using the File Transfer Protocol (FTP) from the Web browser, text pages (with files in the ASCII TXT format) can also be "Webified" for display by Web users. Many Internet Request for Comment (RFC) documents are available on the Web in the text format. The only Webification these files need is to simply make them available in a directory accessible to the FTP server.

Webification is the name of a Web site firm in Hollywood, Florida, that specializes in internetworking and in Web site design.

The Internet

The Internet, sometimes called simply "the Net," is a worldwide system of computer networks - a network of networks in which users at any one computer can, if they have permission, get information from any other computer (and sometimes talk directly to users at other computers). It was conceived by the Advanced Research Projects Agency (ARPA) of the U.S. government in 1969 and was first known as the ARPANet. The original aim was to create a network that would allow users of a research computer at one university to be able to "talk to" research computers at other universities. A side benefit of ARPANet's design was that, because messages could be routed or rerouted in more than one direction, the network could continue to function even if parts of it were destroyed in the event of a military attack or other disaster.

Today, the Internet is a public, cooperative, and self-sustaining facility accessible to hundreds of millions of people worldwide. Physically, the Internet uses a portion of the total resources of the currently existing public telecommunication networks. Technically, what distinguishes the Internet is its use of a set of protocols called TCP/IP (Transmission Control Protocol/Internet Protocol). Two recent adaptations of Internet technology, the intranet and the extranet, also make use of the TCP/IP protocol.

For many Internet users, electronic mail (e-mail) has practically replaced the Postal Service for short written transactions. Electronic mail is the most widely used application on the Net. You can also carry on live "conversations" with other computer users, using IRC (Internet Relay Chat). More recently, Internet telephony hardware and software allows real-time voice conversations.

The most widely used part of the Internet is the World Wide Web (often abbreviated "WWW" or called "the Web"). Its outstanding feature is hypertext, a method of instant cross-referencing. In most Web sites, certain words or phrases appear in text of a different color than the rest; often this text is also underlined. When you select one of these words or phrases, you will be transferred to the site or page that is relevant to this word or phrase. Sometimes there are buttons, images, or portions of images that are "clickable." If you move the pointer over a spot on a Web site and the pointer changes into a hand, this indicates that you can click and be transferred to another site.

Using the Web, you have access to millions of pages of information. Web "surfing" is done with a Web browser, the most popular of which are Netscape Navigator and Microsoft Internet Explorer. The appearance of a particular Web site may vary slightly depending on the browser you use. Also, later versions of a particular browser are able to render more "bells and whistles" such as animation, virtual reality, sound, and music files, than earlier versions.

Intranet

An intranet is a network of networks that is contained within an enterprise. It may consist of many interlinked local area networks and also use leased lines in the wide area network. Typically, an intranet includes connections through one or more gateway computers to the outside Internet. The main purpose of an intranet is to share company information and computing resources among employees. An intranet can also be used to facilitate working in groups and for teleconferences.

An intranet uses TCP/IP, HTTP, and other Internet protocols and in general looks like a private version of the Internet. With tunneling, companies can send private messages through the public network, using the public network with special encryption/decryption and other security safeguards to connect one part of their intranet to another.

Typically, larger enterprises allow users within their intranet to access the public Internet through firewall servers that have the ability to screen messages in both directions so that company security is maintained. When part of an intranet is made accessible to customers, partners, suppliers, or others outside the company, that part is called an extranet.

Extranet

An extranet is a private network that uses the Internet protocols and the public telecommunication system to securely share part of a business's information or operations with suppliers, vendors, partners, customers, or other businesses. An extranet can be viewed as part of a company's intranet that is extended to users outside the company. It has also been described as a "state of mind" in which the Internet is perceived as a way to do business with other companies as well as to sell products to customers. The same benefits that HTML, HTTP, SMTP, and other Internet technologies have brought to the Internet and to corporate intranets now seem designed to accelerate business between businesses.

An extranet requires security and privacy. These require firewall server management, the issuance and use of digital certificates or similar means of user authentication, encryption of messages, and the use of virtual private networks (VPNs) that tunnel through the public network.

Companies can use an extranet to:

- Exchange large volumes of data using Electronic Data Interchange (EDI)
- Share product catalogs exclusively with wholesalers or those "in the trade"
- Collaborate with other companies on joint development efforts
- Jointly develop and use training programs with other companies
- Provide or access services provided by one company to a group of other companies, such as an online banking application managed by one company on behalf of affiliated banks
- Share news of common interest exclusively with partner companies

Netscape, Oracle, and Sun Microsystems have announced an alliance to ensure that their extranet products can work together by standardizing on JavaScript and the Common Object Request Broker Architecture (CORBA). Microsoft supports the Point-to-Point Tunneling Protocol (PPTP) and is working with American Express and other companies on an Open Buying on the Internet (OBI) standard. The Lotus Corporation is promoting its groupware product, Notes, as well-suited for extranet use.

Hypertext

Hypertext is the organization of information units into connected associations that a user can choose to make. An instance of such an association is called a link or hypertext link. (And the highlighted word "link" in the previous sentence is an example of a hypertext link.)

Hypertext was the main concept that led to the invention of the World Wide Web, which is, after all, nothing more (or less) than an enormous amount of information content connected by an enormous number of hypertext links.

The term was first used by Ted Nelson in describing his Xanadu system.

Markup

Markup refers to the sequence of characters or other symbols that you insert at certain places in a text or word processing file to indicate how the file should look when it is printed or displayed or to describe the document's logical structure. The markup indicators are often called "tags." For example, this particular paragraph is preceded by a:

<p> (or paragraph tag)

so that it will be separated by an empty line from the preceding line.

There is now a standard markup definition for document structure (or really a description of how you can define markup) in the Standard Generalized Markup Language (SGML).

Markup can be inserted by the document creator directly by typing the symbols in, by using an editor and selecting prepackaged markup symbols (to save keystrokes), or by using a more sophisticated editor that lets you create the document as you want it to appear (this is called a WYSIWYG editor).

Link

Using hypertext, a link is a selectable connection from one word, picture, or information object to another. In a multimedia environment such as the World Wide Web, such objects can include sound and motion video sequences. The most common form of link is the highlighted word or picture that can be selected by the user (with a mouse or in some other fashion), resulting in the immediate delivery and view of another file. The highlighted object is referred to as an anchor. The anchor reference and the object referred to constitute a hypertext link.

Although most links do not offer the user a choice of types of link, it would be possible for the user to be provided a choice of link types, such as: a definition of the object, an example of it, a picture of it, a smaller or larger picture of it, and so forth.

Links are what make the World Wide Web a web.

HTML (Hypertext Markup Language)

HTML (Hypertext Markup Language) is the set of "markup" symbols or codes inserted in a file intended for display on a World Wide Web browser. The markup tells the Web browser how to display a Web page's words and images for the user.

HTML is defined in practice both by Netscape and Microsoft as they add changes to their Web browsers and more officially for the industry by the World Wide Web Consortium (W3C). A new version of HTML called HTML 4 has recently been officially recommended by W3C, making this level an effective standard. However, both Netscape and Microsoft browsers currently implement some features differently and provide non-standard extensions. Web developers using the more advanced features of HTML 4 may have to design pages for both browsers and send out the appropriate version to a user. Significant features in HTML 4 are sometimes described in general as dynamic HTML.

Dynamic HTML (DHTML)

Dynamic HTML is a collective term for a combination of new Hypertext Markup Language (HTML) tags and options, style sheets, and programming that will let you create Web pages more animated and more responsive to user interaction than previous versions of HTML. Much of dynamic HTML is specified in HTML 4.0. Simple examples of dynamic HTML pages would include (1) having the color of a text heading change when a user passes a mouse over it or (2) allowing a user to "drag and drop" an image to another place on a Web page. Dynamic HTML can allow Web documents to look and act like desktop applications or multimedia productions.

The features that constitute dynamic HTML are included in Netscape Communications' latest Web browser, Navigator 4.0 (part of Netscape's Communicator suite), and by Microsoft's browser, Internet Explorer 4.0. While HTML 4.0 is supported by both Netscape and Microsoft browsers, some additional capabilities are supported by only one of the browsers. The biggest obstacle to the use of dynamic HTML is that, since many users are still using older browsers, a Web site must create two versions of each site and serve the pages appropriate to each user's browser version.

The Concepts and Features in Dynamic HTML

Both Netscape and Microsoft support:
- An object-oriented view of a Web page and its elements
- Cascading style sheets and the layering of content
- Programming that can address all or most page elements
- Dynamic fonts

An Object-Oriented View of Page Elements

Each page element (division or section, heading, paragraph, image, list, and so forth) is viewed as an "object." (Microsoft calls this the "Dynamic HTML Object Model." Netscape calls it the "HTML Object Model." W3C calls it the "Document Object Model.") For example, each heading on a page can be named, given attributes of text style and color, and addressed by name in a small progam or "script" included on the page. This heading or any other element on the page can be changed as the result of a specified event such a mouse passing over or being clicked or a time elapsing. Or an image can be moved from one place to another by "dragging and dropping" the image object with the mouse. (These event possibilities can be viewed as the reaction capabilities of the element or object.) Any change takes place immediately (since all variations of all elements or objects have been sent as part of the same page from the Web server that sent the page). Thus, variations can be thought of as different properties of the object.

Not only can element variations change text wording or color, but everything contained within a heading object can be replaced with new content that includes different or additional HTML as well as different text. Microsoft calls this the "Text Range technology."

Style Sheets and Layering

A style sheet describes the default style characteristics (including the page layout and font type style and size for text elements such as headings and body text) of a document or a portion of a document. For Web pages, a style sheet also describes the default background color or image, hypertext link colors, and possibly the content of page. Style sheets help ensure consistency across all or a group of pages in a document or a Web site.

Dynamic HTML includes the capability to specify style sheets in a "cascading" fashion (that is, linking to or specifying different style sheets or style statements with predefined levels of precedence within the same or a set of related pages). As the result of user interaction, a new style sheet can be made applicable and result in a change of appearance of the Web page. You can have multiple layers of style sheet within a page, a style sheet within a style sheet within a style sheet. A new style sheet may only vary one element from the style sheet above it.

Layering is the use of alternate style sheets or other approaches to vary the content of a page by providing content layers that can overlay (and replace or superimpose on) existing content sections. Layers can be programmed to appear as part of a timed presentation or as the result of user interaction. In Internet Explorer 4.0, Microsoft implements layers through style sheets. Netscape supports the style sheet approach but also offers a new HTML <LAYER>...</LAYER> tag set (that Microsoft does not support). Both approaches are being considered by the W3C Working Committee and both companies say they will support whatever W3C decides will be the recommended approach.

SHTML

(Server-side include Hypertext Markup Language file)

A Web file with the suffix of ".shtml" (rather than the usual ".htm") indicates a file that includes some information that will be added "on the fly" by the server before it is sent to you. A typical use is to include a "Last modified" date at the bottom of the page.

This HTTP facility is referred to as a server-side include. (Although rarely done, the server administrator can identify some other file name suffix than ".shtml" as a server-side include file.) You can think of a server-side include as a limited form of Common Gateway Interface (CGI) application. In fact, the CGI is not used. The server simply searches the server-side include file for CGI environment variables, and inserts the variable information in the places in the file where the "include" statements have been inserted.)

When creating a Web site, a good idea is to ask your server administrator which environment variables can be used and whether the administrator can arrange to set the server up so that these can be handled. Your server administrator should usually be able to help you insert the necessary include statements in an HTML file.

Daemon

A daemon (pronounced and sometimes spelled like "demon") is a program that runs continuously and exists for the purpose of handling periodic service requests that a computer system expects to receive. The daemon program forwards the requests to other programs (or processes) as appropriate. Each server of pages on the Web has an HTTPD or Hypertext Transport Protocol daemon that continually waits for requests to come in from Web clients and their users.

Page

1) On the World Wide Web, a page is a file notated with the Hypertext Markup Language (HTML). Usually, it contains text and specifications about where image or other multimedia files are to be placed when the page is displayed. You can think of a Web site as a book (albeit a hypertext kind of book rather than a sequentially arranged kind of book) that arrives a page at a time as you request each one. Each page is an individual HTML file with its own Web address (URL). The first page you usually request at a site is known as the home page. (Most home pages have a default name that doesn't have to be specified; you only need to enter the domain name for the site itself.) With frames, multiple pages (HTML files) can be downloaded to a browser and presented on designated sections of the display screen at the same time.

2) In a computer's random access memory (RAM), a page is a group of memory cells that are accessed as part of a single operation. That is, all the bits in the group of cells are changed at the same time. In some kinds of RAM, a page is all the memory cells in the same row of cells. In other kinds of RAM, a page may represent some other group of cells than all those in a row.

3) In computer systems that use virtual memory (also known as virtual storage), a page is a unit of data storage that is brought into real storage (on a personal computer, RAM) from auxiliary storage (on a personal computer, usually the hard disk) when a requested item of data is not already in real storage (RAM).

4) In printing, a page is what is printed on a single piece of paper.

In creating a Web site, frames is the use of multiple, independently controllable sections on a Web presentation. This effect is achieved by building each section as a separate HTML file and having one "master" HTML file identify all of the sections. When a user requests a Web page that uses frames, the address requested is actually that of the "master" file that defines the frames; the result of the request is that multiple HTML files are returned, one for each visual section. Links in one frame can request another file that will appear in another (or the same) frame. A typical use of frames is to have one frame containing a selection menu in one frame and another frame that contains the space where the selected (linked to) files will appear.

Frames, originally created by Netscape as an HTML extension and now part of the HTML 4.0, specification are defined with HTML FRAMESET and FRAME tags. Sites that use frames need to create an alternative scheme of pages for requests from browsers that don't support them and possibly for users that prefer a non-frames version.

Web Page types: ice, jello, and liquid

Ice, jello, and liquid are related terms describing three approaches to controlling content placement on a Web page. Because the browser user can control and change both screen resolution and window size, the Web page designer is challenged to design a page that will achieve its intended effect in spite of user resizing.

An ice page is one in which the primary content has a fixed width in pixels and assumes a left margin alignment. Such a page is designed to display optimally on one particular display resolution setting and window size and either specifies or assumes that size. If the resolution is set to a different setting, the page may have unneeded space on the right size of the display window, but this is a risk that the designer chooses to take.

A jello page is one in which the primary content is centered on the page. It also has a fixed-width column, but it is always centered in any size window and at any display resolution. Its center alignment eliminates the excessive right margin seen in an "ice" page (when viewed on a large monitor). But it can have excessive margins on both sides of the centered content.

A liquid space, according to Glenn Davis of Project Cool, is (to paraphrase slightly):..."a Web page that will reflow to fit no matter what size window you pour it into." Unlike the ice and jello approaches, a liquid page or content area leaves no blank margin on the right or the left. However, unless special care is taken, elements in a liquid page can sometimes have too much or too little space between them.

More sophisticated designers sometimes use a combination of tables and tables-within-tables to define and control layout space and element placement in the various resolutions and windows sizes users may choose to use. To create liquid content spaces or pages, you may be able to determine the window size and dynamically adjust your content layout to meet each user's requirements.

Home page

1) For a Web user, the home page is the first Web page that is displayed after starting a Web browser like Netscape's Navigator or Microsoft's Internet Explorer. The browser is usually preset so that the home page is the first page of the browser manufacturer. However, you can set it to open to any Web site. For example, you can specify that "http://www.yahoo.com" or "http://whatis.com" be your home page. You can also specify that there be no home page (a blank space will be displayed) in which case you choose the first page from your bookmark list or enter a Web address.

2) For a Web site developer, a home page is the first page presented when a user selects a site or presence on the World Wide Web. The usual address for a Web site is the home page address, although you can enter the address (URL) of any page and have that page sent to you.

A Web site is a collection of Web files on a particular subject that includes a beginning file called a home page. For example, most companies, organizations, or individuals that have Web sites have a single address that they give you. This is their home page address. From the home page, you can get to all the other pages on their site. For example, the Web site for IBM has the home page address of http://www.ibm.com. (In this case, the actual file name of the home page file doesn't have to be included because IBM has named this file index.html and told the server that this address really means http://www.ibm.com/index.html.)

Since it sounds like geography is involved, a Web site is rather easily confused with a Web server. A server is a computer that holds the files for one or more sites. On one hand, a very large Web site may reside on a number of servers that may be in different geographic locations. IBM is a good example; its Web site consists of thousands of files spread out over many servers in world-wide locations. But a more typical example is probably the site you are looking at, whatis.com. We reside on a commercial space provider's server with a number of other sites that have nothing to do with Internet glossaries.

A synonym and less frequently used term for Web site is "Web presence." That term seems to better express the idea that a site is not tied to specific geographic location, but is "somewhere in cyberspace." However, "Web site" seems to be used much more frequently.

You can have multiple Web sites that cross-link to files on each others' sites. This simply means that you've identified two starting places or home pages for all the files.

Some publications have begun using the term "Website." and "Web site".

Web site

This definition is also listed under presence, site and Website.

A Web site is a related collection of Web files that includes a beginning file called a home page. A company or an individual tells you how to get to their Web site by giving you the address of their home page. From the home page, you can get to all the other pages on their site. For example, the Web site for IBM has the home page address of http://www.ibm.com. (The home page address actually includes a specific file name like index.html but, as in IBM's case, when a standard default name is set up, users don't have to enter the file name.) IBM's home page address leads to thousands of pages. (But a Web site can also be just a few pages.)

Since site implies a geographic place, a Web site can be confused with a Web server. A server is a computer that holds the files for one or more sites. A very large Web site may be spread over a number of servers in different geographic locations. IBM is a good example; its Web site consists of thousands of files spread out over many servers in world-wide locations. But a more typical example is probably the site you are looking at, whatis.com. We reside on a commercial space provider's server with a number of other sites that have nothing to do with Internet glossaries.

A synonym and less frequently used term for Web site is "Web presence." That term seems to better express the idea that a site is not tied to specific geographic location, but is "somewhere in cyberspace." However, "Web site" seems to be used much more frequently.

You can have multiple Web sites that cross-link to files on each others' sites or even share the same files.

Web server

A Web server is a program that, using the client/server model and the World Wide Web's Hypertext Transfer Protocol (HTTP), serves the files that form Web pages to Web users (whose computers contain HTTP clients that forward their requests). Every computer on the Internet that contains a Web site must have a Web server program. The most popular Web servers are Microsoft's Internet Information Server (IIS), which comes with the Windows NT server; Netscape's FastTrack and Enterprise servers; and Apache, a Web server for UNIX-based operating systems. Other Web servers include Novell's Web Server for users of its NetWare operating system and IBM's family of Lotus Domino servers, primarily for IBM's OS/390 and AS/400 customers.

Web servers often come as part of a larger package of Internet- and intranet-related programs for serving e-mail, downloading requests for FTP files, and building and publishing Web pages. Considerations in choosing a Web server include how well it works with the operating system and other servers, its ability to handle server-side programming, and publishing, search engine, and site building tools that may come with it.

Protocols

TCP/IP (Transmission Control Protocol/Internet Protocol) is the basic communication language or protocol of the Internet. It can also be used as a communications protocol in the private networks called intranets and in extranets. When you are set up with direct access to the Internet, your computer is provided with a copy of the TCP/IP program just as every other computer that you may send messages to or get information from also has a copy of TCP/IP.

TCP/IP is a two-layered program. The higher layer, Transmission Control Protocol, manages the assembling of a message or file into smaller packets that are transmitted over the Internet and received by a TCP layer that reassembles the packets into the original message. The lower layer, Internet Protocol, handles the address part of each packet so that it gets to the right destination. Each gateway computer on the network checks this address to see where to forward the message. Even though some packets from the same message are routed differently than others, they'll be reassembled at the destination.

TCP/IP uses the client/server model of communication in which a computer user (a client) requests and is provided a service (such as sending a Web page) by another computer (a server) in the network. TCP/IP communication is primarily point-to-point, meaning each communication is from one point (or host computer) in the network to another point or host computer. TCP/IP and the higher-level applications that use it are collectively said to be "connectionless" because each client request is considered a new request unrelated to any previous one (unlike ordinary phone conversations that require a dedicated connection for the call duration). Being connectionless frees network paths so that everyone can use them continuously. (Note that the TCP layer itself is not connectionless as far as any one message is concerned. Its connection remains in place until all packets in a message have been received.)

Many Internet users are familiar with the even higher layer application protocols that use TCP/IP to get to the Internet. These include the World Wide Web's Hypertext Transfer Protocol (HTTP), the File Transfer Protocol (FTP), Telnet (Telnet) which lets you logon to remote computers, and the Simple Mail Transfer Protocol (SMTP). These and other protocols are often packaged together with TCP/IP as a "suite."

Personal computer users usually get to the Internet through the Serial Line Internet Protocol (SLIP) or the Point-to-Point Protocol (PPP). These protocols encapsulate the IP packets so that they can be sent over a dial-up phone connection to an access provider's modem.

Protocols related to TCP/IP include the User Datagram Protocol (UDP), which is used instead of TCP for special purposes. Other protocols are used by network host computers for exchanging router information. These include the Internet Control Message Protocol (ICMP), the Interior Gateway Protocol (IGP), the Exterior Gateway Protocol (EGP), and the Border Gateway Protocol (BGP).

HTTP (Hypertext Transfer Protocol)

The Hypertext Transfer Protocol (HTTP) is the set of rules for exchanging files (text, graphic images, sound, video, and other multimedia files) on the World Wide Web. Relative to the TCP/IP suite of protocols (which are the basis for information exchange on the Internet), HTTP is an application protocol.

Essential concepts that are part of HTTP include (as its name implies) the idea that files can contain references to other files whose selection will elicit additional transfer requests. Any Web server machine contains, in addition to the HTML and other files it can serve, an HTTP daemon, a program that is designed to wait for HTTP requests and handle them when they arrive. Your Web browser is an HTTP client, sending requests to server machines. When the browser user enters file requests by either "opening" a Web file (typing in a Uniform Resource Locator or URL) or clicking on a hypertext link, the browser builds an HTTP request and sends it to the Internet Protocol address indicated by the URL. The HTTP daemon in the destination server machine receives the request and, after any necessary processing, the requested file is returned.

HTTPS (Secure Hypertext Transfer Protocol)

HTTPS (Secure Hypertext Transfer Protocol) is a Web protocol developed by Netscape and built into its browser that encrypts and decrypts user page requests as well as the pages that are returned by the Web server. HTTPS is really just the use of Netscape's Secure Socket Layer (SSL) as a sublayer under its regular HTTP application layer. (HTTPS uses port 443 instead of HTTP port 80 in its interactions with the lower layer, TCP/IP.) SSL uses a 40-bit key size for the RC4 stream encryption algorithm, which is considered an adequate degree of encryption for commercial exchange.

Suppose you use a Netscape browser to visit a Web site such as NetPlaza (http://www.netplaza.com) and view their catalog. When you're ready to order, you will be given a Web page order form with a URL that starts with https://. When you click "Send," to send the page back to the catalog retailer, your browser's HTTPS layer will encrypt it. The acknowledgement you receive from the server will also travel in encrypted form, arrive with an https:// URL, and be decrypted for you by your browser's HTTPS sublayer.

HTTPS and SSL support the use of X.509 digital certificates from the server so that, if necessary, a user can authenticate the sender. SSL is an open, nonproprietary protocol that Netscape has proposed as a standard to the World Wide Consortium (W3C). HTTPS is not to be confused with SHTTP, a security-enhanced version of HTTP developed and proposed as a standard by EIT.

HTTPD (Hypertext Transfer Protocol daemon)

On the Web, each server has an HTTPD or Hypertext Transfer Protocol daemon that waits in attendance for requests to come in from the rest of the Web. A daemon is a program that is "an attendant power or spirit" (Webster's). It sits waiting for requests to come in and then forwards them to other processes as appropriate.

PPP (Point-to-Point Protocol)

PPP (Point-to-Point Protocol) is a protocol for communication between two computers using a serial interface, typically a personal computer connected by phone line to a server. For example, your Internet server provider may provide you with a PPP connection so that the provider's server can respond to your requests, pass them on to the Internet, and forward your requested Internet responses back to you. PPP uses the Internet protocol (IP) (and is designed to handle others). It is sometimes considered a member of the TCP/IP suite of protocols. Relative to the Open Systems Interconnection (OSI) reference model, PPP provides layer 2 (data-link layer) service.

PPP is a full-duplex protocol that can be used on various physical media, including twisted pair or fiber optic lines or satellite transmission. It uses a variation of High Speed Data Link Control (HDLC) for packet encapsulation.

PPP is usually preferred over the earlier de facto standard Serial Line Internet Protocol (SLIP) because it can handle synchronous as well as asynchronous communication. PPP can share a line with other users and it has error detection that SLIP lacks. Where a choice is possible, PPP should be preferred.

FTP (File Transfer Protocol)

FTP (File Transfer Protocol), a standard protocol, is the simplest way to exchange files between computers on the Internet. Like the Hypertext Transfer Protocol (HTTP), which transfers displayable Web pages and related files, and the Simple Mail Transfer Protocol (SMTP), which transfers e-mail, FTP is an application protocol that uses the Internet's TCP/IP protocols. FTP is commonly used to transfer Web page files from their creator to the computer that acts as their server for everyone on the Internet. It's also commonly used to download programs and other files to your computer from other servers.

As a user, you can use FTP with a simple command line interface (for example, from the Windows MS-DOS Prompt window) or with a commercial program that offers a graphical user interface. Your Web browser can also make FTP requests to download programs you select from a Web page. Using FTP, you can also update (delete, rename, move, and copy) files at a server. You need to log on to an FTP server. However, publicly available files are easily accessed using anonymous FTP.

FTP is usually provided as part of a suite of programs that come with TCP/IP.

URL (Uniform Resource Locator)

A URL (Uniform Resource Locator) (pronounced YU-AHR-EHL or, in some quarters, UHRL) is the address of a file (resource) accessible on the Internet. The type of resource depends on the Internet application protocol. Using the World Wide Web's protocol, the Hypertext Transfer Protocol (HTTP) , the resource can be an HTML page (like the one you're reading), an image file, a program such as a CGI application or Java applet, or any other file supported by HTTP. The URL contains the name of the protocol required to access the resource, a domain name that identifies a specific computer on the Internet, and a hierarchical description of a file location on the computer.

On the Web (which uses the Hypertext Transfer Protocol), an example of a URL is:

http://www.mhrcc.org/kingston

which describes a Web page to be accessed with an HTTP (Web browser) application that is located on a computer named www.mhrcc.org. The specific file is in the directory named /kingston and is the default page in that directory (which, on this computer, happens to be named index.html).

An HTTP URL can be for any Web page, not just a home page, or any individual file. For example, this URL would bring you the whatis.com logo image:

http://whatis.com/whatisAnim2.gif

A URL for a program such as a forms-handling CGI script written in Perl might look like this:

http://whatis.com/cgi-bin/comments.pl

A URL for a file meant to be downloaded would require that the "ftp" protocol be specified like this one:

ftp://www.somecompany.com/whitepapers/widgets.ps

Cyberspace

Cyberspace is the total interconnectedness of human beings through computers and telecommunication without regard to physical geography.

William Gibson is sometimes credited with inventing or popularizing the term by using it in his novel of 1984, Neuromancer.

Surf

To surf the Internet is to explore cyberspace without a predefined agenda. By cyberspace, we generally mean the World Wide Web. As of early 1997, few people surfed gopherspace and fewer still "FTP-space."

Note that many Web users find themselves surfing when they started out with a specific Web site to visit. The Web is like that because it is a web of interconnections whose relevance and attraction for any given Web user are impossible to predict in advance.

Browser

A browser is an application program that provides a way to look at and interact with all the information on the World Wide Web. The word "browser" seems to have originated prior to the Web as a generic term for user interfaces that let you browse text files online. By the time the first Web browser with a graphical user interface was invented (Mosaic, in 1992), the term seemed to apply to Web content, too. Technically, a Web browser is a client program that uses the Hypertext Transfer Protocol (HTTP) to make requests of Web servers throughout the Internet on behalf of the browser user. A commercial version of the original browser, Mosaic, is in use. Many of the user interface features in Mosaic, however, went into the first widely-used browser, Netscape Navigator. Microsoft followed with its Internet Explorer. Today, these two browsers are highly competitive and the only two browsers that the vast majority of Internet users are aware of. Although the online services, such as America Online, Compuserve, and Prodigy, originally had their own browsers, virtually all now offer the Netscape or Microsoft browser. Lynx is a text-only browser for UNIX shell and VMS users. Another recently offered browser is Opera.

Lynx

Lynx is a keyboard-oriented text-only Web browser that was developed at the University of Kansas primarily for students who used UNIX workstations. It has also been rewritten to run on VMS operating systems for users of VT100 terminals.

If you use the UNIX shell interface and your access provider offers it, Lynx may be interesting for you since it has a succinct key- (not mouse-) driven user interface. Information about Lynx, including where to download it, is available at the official Lynx server page.

World Wide Web

A technical definition of the World Wide Web is: all the resources and users on the Internet that are using the Hypertext Transfer Protocol (HTTP).

A broader definition comes from the organization that Web inventor Tim Berners-Lee helped found, the World Wide Web Consortium (W3C):

 "The World Wide Web is the universe of network-accessible information, an embodiment of human knowledge."

W3C (World Wide Web Consortium)

The World Wide Web Consortium (W3C) describes itself as follows:

"The World Wide Web Consortium exists to realize the full potential of the Web.

 The W3C is an industry consortium which seeks to promote standards for the evolution of the Web and interoperability between WWW products by producing specifications and reference software. Although W3C is funded by industrial members, it is vendor-neutral, and its products are freely available to all.

 The Consortium is international; jointly hosted by the MIT Laboratory for Computer Science in the United States and in Europe by INRIA who provide both local support and performing core development. The W3C was initially established in collaboration with CERN, where the Web originated, and with support from DARPA and the European Commission."

Organizations may apply for membership to the Consortium; individual membership isn't offered. The W3C has taken over what was formerly called the CERN httpd or Web server. You'll find some useful information at the W3C Web site (http://www.w3.org).

E-mail - (electronic mail)

E-mail (electronic mail) is the exchange of computer-stored messages by telecommunication. (Some publications spell it email; we prefer the currently more established spelling of e-mail.) E-mail messages are usually encoded in ASCII text. However, you can also send non-text files, such as graphic images and sound files, as attachments sent in binary streams. E-mail was one of the first uses of the Internet and is still the most popular use. A large percentage of the total traffic over the Internet is e-mail. E-mail can also be exchanged between online service users and in networks other than the Internet, both public and private.

E-mail can be distributed to lists of people as well as to individuals. A shared distribution list can be managed by using an e-mail reflector. Some mailing lists allow you to subscribe by sending a request to the mailing list administrator. A mailing list that is administered automatically is called a list server.

E-mail is one of the protocols included with the Transport Control Protocol/Internet Protocol (TCP/IP) suite of protocols. A popular protocol for sending e-mail is SMTP and a popular protocol for receiving it is POP3. Both Netscape and Microsoft include an e-mail utility with their Web browsers.

Carbon Copy and blind carbon copy

In e-mail, a carbon copy (abbreviated "cc," and sometimes "fcc" for "first carbon copy") is a copy of a note sent to an addressee other than the main addressee. A blind carbon copy is a copy sent to an addressee that is not visible to the main and carbon copy addressees. For example, you may have a work colleague that acts as a back-up when you're on vacation or not at work. You don't necessarily want the people you correspond with to know that you have a back-up. So, to keep your back-up informed, you always send the back-up a blind carbon copy. The fact that a blind carbon copy was sent is not apparent to the main and carbon copy recipients.

The term is borrowed from the days of the mechanical and later the electronic typewriter (circa 1879-1979) when copies of typed sheets of paper were made by inserting a special sheet of inked paper called carbon paper into the typewriter. For two copies, you would insert carbon paper (sometimes just called a "carbon") between the original being typed and each of the two sheets that would become the carbon copies.

IRC (Internet Relay Chat)

For terms frequently used in online keyboard chatting, see Chat/IRC/BBS acronyms.

Internet Relay Chat (IRC) is a system for chatting that involves a set of rules and conventions and client/server software. On the Web, certain sites such as Talk City or IRC networks such as the Undernet provide servers and help you download an IRC client to your PC. Talk City also offers an IRC client applet that it downloads for you as part of their home page so that you can start chatting right away.

You can start a chat group (called a channel) or join an existing one. There is a protocol for discovering existing chat groups and their members. Participants in some chat groups use nicknames that last only for the duration of the session (you can't "own" a nickname). Other channels encourage you to register a nickname that you always use and even offer space for a personal profile, picture, and personal home page link.

Popular ongoing IRC channels are #hottub and #riskybus. A number of channels are set up and conducted in foreign languages.

Chatting

For terms frequently used in online keyboard chatting, see Chat/IRC/BBS acronyms.

On the Internet, chatting is talking to other people who are using the Internet at the same time you are. Usually, this "talking" is the exchange of typed-in messages requiring one site as the repository for the messages (or "chat site") and a group of users who take part from anywhere on the Internet. In some cases, a private chat can be arranged between two parties who meet initially in a group chat. Chats can be ongoing or scheduled for a particular time and duration. Most chats are focused on a particular topic of interest and some involve guest experts or famous people who "talk" to anyone joining the chat. (Transcripts of a chat can be archived for later reference.)

Chats are conducted on online services (especially America Online), by bulletin board services, and by Web sites. Several Web sites, notably Talk City, exist solely for the purpose of conducting chats. Some chat sites such as Worlds Chat allow participants to assume the role or appearance of an avatar in a simulated or virtual reality environment.

Talk City and many other chat sites use a protocol called Internet Relay Chat (IRC). A chat can also be conducted using sound or sound and video, assuming you have the bandwidth access and the appropriate programming.

Internet telephony

This duplicates the definition of telephony.

Telephony is the technology associated with the electronic transmission of voice, fax, or other information between distant parties using systems historically associated with the telephone, a handheld device containing both a speaker or transmitter and a receiver. With the arrival of computers and the transmittal of digital information over telephone systems and the use of radio to transmit telephone signals, the distinction between telephony and telecommunication has become difficult to find. However, we believe that telephony does connote voice or spoken and heard information predominately and it usually assumes a point-to-point (rather than a broadcast) connection. It also tends to assume a temporarily dedicated connection.

Internet telephony is the use of the Internet rather than the traditional telephone company infrastructure and rate structure to exchange spoken or other telephone information. Since access to the Internet is available at local phone connection rates, an international or other long-distance call will be much less expensive than through the traditional call arrangement.

On the Internet, three new services are now or will soon be available:

The ability to make a normal voice phone call (whether or not the person called is immediately available; that is, the phone will ring at the location of the person called)
The ability to send fax transmissions at very low cost (at local call prices) through a gateway point on the Internet in major cities
The ability to leave voice mail at a called number

Some companies that make products that provide or plan to provide these capabilities include: IDT Corporation (Net2Phone), Netspeak, NetXchange, Rockwell International, VocalTec, and Voxspeak. Among uses planned for Internet phone services are phone calls to customer service people while viewing a product catalog online at a Web site.

You can now add telephone capabilities to your computer by adding a telephony board, available for under $300, that combines the functions of modem, sound board, speakerphone, and voicemail system. A telephony board is often integrated into new machines targeted for small business and home office users.

A Telephony API (application program interface) is available from Microsoft and Intel that allows Windows client applications to access voice services on a server and that interconnects PC and phone systems.

Internet2

Internet2 is a collaboration among more than 100 U.S. universities to develop networking and advanced applications for learning and research. Since much teaching, learning, and collaborative research may require real-time multimedia and high-bandwidth interconnection, a major aspect of Internet2 is adding sufficient network infrastructure to support such applications. But Internet2 also intends to investigate and develop new ways to use the Internet and the Internet2 infrastructure for its educational purposes. Although Internet2 is not envisioned as a future replacement for the Internet, its organizers hope to share their developments with other networks, including the Internet. Internet2 will include and further develop the National Science Foundation's very high-speed Backbone Network Service (vBNS) that currently interconnects research supercomputers in the U.S. The involved institutions plan to continue using the existing Internet for "ordinary" services such as e-mail, personal Web access, and newsgroups.

Internet2 collaborators plan to use Quality-of-Service (QoS) tools so that participants can reserve and use bandwidth for special events or in certain time periods. Here are some possibilities envisioned by Internet2:

Distributed learning modules: Conceptually, teachers and students can be share materials in cyberspace with students learning in a self-directed manner under the supervision of an educational system or teacher. Internet2 foresees tools that would make it easy to create what they call LearningWare, using existing object-oriented programming methodologies. Internet2 may also help realize the Instructional Management System (IMS), a standard process for using the Internet in developing and delivering learning packages and tracking outcomes. One can think of the IMS as a more structured way to exploit the potential learning materials on the World Wide Web.

New ways to envision and retrieve information: In the future, today's text-oriented models of information structure could be replaced by interactive pictures of information structure (compare a textual taxonomy or table of contents with illustrations of interlinked and explodable animal forms, for example). With Internet2's high-bandwidth connections, experiments in such information visualization will be possible; new ideas can be tried out. In environments where up-to-date information is valuable, information can be pushed to users at their request.

Virtual environment sharing: Sometimes called tele-immersion, participants in teleconferences could share the perception that everyone was in the same physical place, possibly with virtual (but somewhat real) models of shared work objects such as architectural models or multimedia storyboards. You would be able to see yourself with others in a far-away conference room, talking and perhaps manipulating objects in the room.

Virtual laboratory: A virtual laboratory would allow scientists in a number of different physical locations, each with unique expertise, computing resources, and/or data to collaborate efficiently not simply at a meeting but in an ongoing way. Effectively, such a project would extend and pool resources while engendering orderly communication and progress toward shared goals. For example, a group of astronomers and computer scientists at the supercomputing centers in the U.S. are attempting to share experiments and knowledge about the origin of the universe. Shared visualizations of alternative possibilities could conceivably suggest additional or refined alternatives. Virtual laboratories are also envisioned for the design and manufacturing of complex systems such as airplanes and for studying and forecasting weather patterns.

Java Programming

Although JavaScript, Java applets, and ActiveX controls were present in previous levels of Web pages, dynamic HTML implies an increased amount of programming in Web pages since more elements of a page can be addressed by a program.

Dynamic Fonts

Netscape includes dynamic fonts as part of dynamic HTML. This feature of Netscape's Navigator browser in its Communicator suite lets Web page designers include font files containing specific font styles, sizes, and colors as part of a Web page and to have the fonts downloaded with the page. That is, the font choice no longer is dependent on what the browser provides.

JavaScript

JavaScript is an interpreted programming or script language from Netscape. It is somewhat similar in capability to Microsoft's Visual Basic, Sun's Tcl, the UNIX-derived Perl, and IBM's REXX. In general, script languages are easier and faster to code in than the more structured and compiled languages such as C and C++. Script languages generallly take longer to process than compiled languages, but are very useful for shorter programs.

JavaScript is used in Web site development to do such things as:

 Automatically change a formatted date on a Web page (see our "Today" page)
 Cause a linked-to page to appear in a popup window (see our "Make a WordPop!" page)
 Cause text or a graphic image to change during a mouse rollover

JavaScript uses some of the same ideas found in Java, the compiled object-oriented language derived from C++. JavaScript code can be imbedded in HTML pages and interpreted by the Web browser (or client). JavaScript can also be run at the server as in Microsoft's Active Server Pages (ASPs) before the page is sent to the requestor. Both Microsoft and Netscape browsers support JavaScript, but sometimes in slightly different ways.

Applet

An applet is a little application program. Prior to the World Wide Web, the built-in writing and drawing programs that came with Windows were sometimes called "applets." On the Web, using Java, the object-oriented programming language, an applet is a small program that can be sent along with a Web page to a user. Java applets can perform interactive animations, immediate calculations, or other simple tasks without having to send a user request back to the server.

ActiveX

ActiveX is the name Microsoft has given to a set of "strategic" object-oriented program technologies and tools. The main technology is the Component Object Model (COM). Used in a network with a directory and additional support, COM becomes the Distributed Component Object Model (DCOM). The main thing that you create when writing a program to run in the ActiveX environment is a component, a self-sufficient program that can be run anywhere in your ActiveX network (currently a network consisting of Windows and Macintosh systems). This component is known as an ActiveX control. ActiveX is Microsoft's answer to the Java technology from Sun Microsystems. An ActiveX control is roughly equivalent to a Java applet.

If you have a Windows operating system on your personal computer, you may notice a number of Windows files with the "OCX" file name suffix. OCX stands for "Object Linking and Embedding control." Object Linking and Embedding (OLE) was Microsoft's program technology for supporting compound documents such as the Windows desktop. The Component Object Model now takes in OLE as part of a larger concept. Microsoft now uses the term "ActiveX control" instead of "OCX" for the component object.

One of the main advantages of a component is that it can be re-used by many applications (referred to as component containers). A COM component object (ActiveX control) can be created using one of several languages or development tools, including C++ and Visual Basic, or PowerBuilder, or with scripting tools such as VBScript.

Currently, ActiveX controls run in Windows 95 and NT and in Macintosh. Microsoft plans to support ActiveX controls for UNIX.

An ActiveX control is a component program object that can be re-used by many application programs within a computer or among computers in a network. The technology for creating ActiveX controls is part of Microsoft's overall ActiveX set of technologies, chief of which is the Component Object Model (COM). ActiveX controls can be downloaded as small programs or animations for Web pages, but they can also be used for any commonly-needed task by an application program in the latest Windows and Macintosh environments. In general, ActiveX controls replace the earlier OCXs (Object Linking and Embedding custom controls). An ActiveX control is roughly equivalent in concept and implementation to the Java applet.

An ActiveX control can be created in any programming language that recognizes Microsoft's Component Object Model (COM). An ActiveX control is a component or self-contained program package that can be created and reused by many applications in the same computer or in a distributed network. The distributed support for COM is called the Distributed Component Object Model (DCOM). In implementation, an ActiveX control is a Dynamic Link Library (DLL) module. An ActiveX control runs in what is known as a container, an application program that uses the Component Object Model program interfaces. This reuseable component approach to application development reduces development time and improves program capability and quality. Windows 95 and NT application development programs such as Powerbuilder and Microsoft Access take advantage of ActiveX controls.

Visual Basic and C++ are commonly used to write OCX or ActiveX controls.

Digital Certificate

A digital certificate is an electronic "credit card" that establishes your credentials when doing business or other transactions on the Web. It is issued by a certification authority (CA). It contains your name, a serial number, expiration dates, a copy of the certificate holder's public key (used for encrypting and decrypting messages and digital signatures), and the digital signature of the certificate-issuing authority so that a recipient can verify that the certificate is real. Some digital certificates conform to a standard, X.509. Digital certificates can be kept in registries so that authenticated users can look up other users' public keys.

Encryption and Decryption

Encryption is the conversion of data into a form, called a cipher, that cannot be easily intercepted by unauthorized people. Decryption is the process of converting encrypted data back into its original form, so it can be understood.

The use of encryption/decryption is as old as the art of communication. In wartime, a cipher, often incorrectly called a "code," can be employed to keep the enemy from obtaining the contents of transmissions. (Technically, a code is a means of representing a signal without the intent of keeping it secret; examples are Morse code and ASCII.) Simple ciphers include the substitution of letters for numbers, the rotation of letters in the alphabet, and the "scrambling" of voice signals by inverting the sideband frequencies. More complex ciphers work according to sophisticated computer algorithms that rearrange the data bits in digital signals.

In order to easily recover the contents of an encrypted signal, the correct decryption key is required. The key is an algorithm that "undoes" the work of the encryption algorithm. Alternatively, a computer can be used in an attempt to "break" the cipher. The more complex the encryption algorithm, the more difficult it becomes to eavesdrop on the communications without access to the key. Encryption/decryption is especially important in wireless communications. This is because wireless circuits are easier to "tap" than their hard-wired counterparts. Nevertheless, encryption/decryption is a good idea when carrying out any kind of sensitive transaction, such as a credit-card purchase online, or the discussion of a company secret between different departments in the organization. The stronger the cipher -- that is, the harder it is for unauthorized people to break it -- the better, in general. However, as the strength of encryption/decryption increases, so does the cost.

In recent years, a controversy has arisen over so-called strong encryption. This refers to ciphers that are essentially unbreakable without the decryption keys. While most companies and their customers view it as a means of keeping secrets and minimizing fraud, some governments view strong encryption as a potential vehicle by which terrorists might evade authorities. These governments, including that of the United States, want to set up a key-escrow arrangement. This means everyone who uses a cipher would be required to provide the government with a copy of the key. Decryption keys would be stored in a supposedly secure place, used only by authorities, and used only if backed up by a court order. Opponents of this scheme argue that criminals could hack into the key-escrow database and illegally obtain, steal, or alter the keys. Supporters claim that while this is a possibility, implementing the key escrow scheme would be better than doing nothing to prevent criminals from freely using encryption/decryption.