

Pat Dalton, Legislative Analyst
(651) 296-7434

Questions can be referred to either Pat Dalton or Jeanne LeFevre (651) 296-5058

The Internet: A Primer

The Internet is a complicated and constantly evolving system of interconnected computer networks and related equipment, software, and transmission links. The Internet has or will have an impact on various public policy issues such as taxation, regulated industries, consumer protection, banking, and crime. Because of this, it is useful to have a basic understanding of what it is and how it works. This information brief contains a simplified overview of the Internet and a glossary of terms that are commonly encountered in discussing the Internet. Terms that are included in the glossary are *italicized* when they are used elsewhere in the paper. This is one in a series of information briefs related to the Internet and taxation.

What is the Internet?

The Internet is a network of computer networks. At the most basic level, the Internet is composed of computers, software, routers, switches, and transmission lines. The computers on the Internet can talk to each other because they all use a suite of agreed upon networking protocols called *TCP/IP*.¹

On a more personal level, the Internet is an interactive means of communicating and gathering information. In addition to finding information on the Internet, users can respond to and add to the information on the system by sending e-mail, posting messages to discussion groups, or setting up their own web page. Even the act of getting information is interactive, requiring the user to input information on the topic(s) of interest and select Internet sites to visit.

¹ The Federal Networking Council, a group of federal agencies with large investments or stakes in Internet activities, defined the Internet as a global information system that (1) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons; (2) is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols; and (3) provides, uses, or makes accessible, either publicly or privately, high-level services layered on the communications and related infrastructure described herein.

The core of the Internet is the high-capacity transmission lines and *routers* that move information between major connection points around the world. The entities that provide these transmission and routing services are called *backbone providers*. Internet users are connected to each other through these *backbones* and through local and *regional networks* that connect into these *backbones*.

Who developed the Internet?

The Internet (or its predecessors) has been around since the 1970s. *APRAnet*, designed for the U.S. Defense Department, was the first network to use the *TCP/IP* protocols and the concept of *packet switching*, which is fundamental to the operation of the Internet. It was supplemented by a more general network or *backbone* called *NSFnet*, a network funded by the National Science Foundation (NSF) to provide service to the government and higher education sectors.

In 1994, because of the increasing commercial use on the Internet, the NSF decided to get out of the *backbone* business. Although the government still maintains its own *backbone* which provides the Internet link to many federal agencies and educational institutions, the vast majority of Internet transmissions are now carried by private, commercial *backbone* providers.

How global is the Internet?

According to the Internet Society (ISOC), as of June 15, 1997, 195 countries were connected to the Internet and 42 countries were not connected.² However, “connected” means different things in different countries. In the United States, Canada, and much of Western Europe, local access to the Internet exists in most areas. Certain lesser developed areas, particularly some countries in Africa, South America, and parts of Asia, may only have Internet access in the country’s capital city. Access is often limited by existing telecommunications systems. In some countries the government, universities, and big businesses may have access to the Internet but private individuals do not. By far the most Internet activity currently occurs in the United States.

Who “owns” the Internet?

Nobody and everybody owns the Internet. Different parts of the Internet are owned and controlled by different entities and are only part of the Internet because the owners agree to use the same standards and protocols. The physical components of the Internet run the gamut from the high-speed transmission lines controlled by telecommunications companies to government computer networks to privately owned computers hooked up to the Internet via a dial-up modem.

² For information on the specific countries that are connected and the type of connection available, please refer to the list on the ISOC web site at: <http://www.isoc.org>.

Who controls or runs the Internet?

Unlike standard commercial industries, the Internet is controlled and operated on a complex, informal, and cooperative basis. The model for this cooperative control comes from the academic and research environment, which is where the Internet was first developed. Entities cooperate in this unstructured system because it is in their best interest to do so.

The closest the Internet gets to a controlling body is the Internet Society (ISOC) which is a voluntary group open to all interested persons. The purpose of the ISOC is to promote global information exchange through Internet technology. This group is led by a group of “invited volunteers” called the Internet Architecture Board (IAB) that is responsible for adopting standards and allocating resources. Another subset of the ISOC is the Internet Engineering Task Force (IETF) which discusses and finds solutions to operational and technical problems of the Internet. Through working groups, the IETF suggests new procedures and standards for adoption by the IAB. Anyone can suggest problems for the IETF to study.

Although the ISOC and IAB adopt standards and methods for how to make the Internet work, they have no real authority to enforce the standards. Policing is done on an informal basis. The Internet is run essentially by peer pressure. If an entity refuses to play by the accepted rules and standards then nobody else will “play with them.” In effect, the other groups that make up the Internet refuse to recognize miscreants as part of the system.³

With the recent growth of the commercial Internet, there is agreement for the necessity of establishing a more formal structure for the governing and standards mechanisms so that things like address assignment, domain name assignment, and official protocol establishment are better defined for fairness and equity.

What are the organizational components that make up the Internet?

Figure 1 is a simplified diagram of the physical structure of the Internet, divided into four levels:

- *Internet backbone providers*
- *regional networks*
- *Internet service providers (ISPs)*
- users

Internet backbone providers are entities operating major networks of high-speed routers and high-speed bulk transmission lines that carry most of the Internet traffic. *Internet backbone providers* may own the transmission lines or lease them from other telecommunications companies. *Backbones* connect with each other at public exchange points called *network access points (NAPs)* which are recognized routing points for the entire system. *NAPs* are distributed all over the world. The United States has 11 official and unofficial *NAPs*. A *backbone provider*

³ Vigilante justice is also alive and well on the Internet. One example is the *flaming* that occurs when a user is perceived as violating accepted norms or standards regarding message content or engaging in inappropriate commercial activities such as *spamming*.

may have multiple lines connecting two *NAPs* to ensure service if one line “crashes.” In addition, *backbone providers* have private interconnections or exchange points with each other to provide transmission shortcuts and to avoid potential bottle necks around the *NAPs*.

Prior to 1994, the major *backbone* of the Internet was *NSFnet* which was sponsored by the National Science Foundation. However, the government funding of this *backbone* legally limited the amount of commercial activity on the Internet. In 1994, the government decided to get out of the *backbone* business, and most of the Internet transmissions are now carried by private, commercial *backbone providers* such as MCI and Sprint and UUNET.

A *backbone provider* earns money by charging *regional networks* and *ISPs* for connections to and use of their transmission lines. However, *backbone providers*, through agreements known as “peering,” carry transmissions from other *backbones* without charge. It is these interconnections between *backbones* that make the Internet a global “network of networks.”

Regional networks provide transmission connections and system routers within a more limited geographic region. Often a *regional network* will be directly connected to a *NAP*, but because they are not recognized as peers by the *backbone providers* they must buy⁴ access to the global portion of the Internet from a *backbone provider*. Sometimes a *regional network* will connect to several *backbone providers*. *Regional networks* provide the transmission lines that connection most of the smaller cities in a region to the Internet.

ISPs connect intranets and individual users to the Internet. An *ISP* controls the hardware (*routers* and modems) and transmissions lines that allow users in a local area to connect to the Internet. Most *ISPs* pay a *regional network*⁵ or an Internet *backbone* to connect to the Internet. Places where a *regional network* or *ISP* allows others to connect into its system is called a *point of presence (POP)*. *ISPs* can vary from small two- or three-person operations to very large businesses with a national presence.


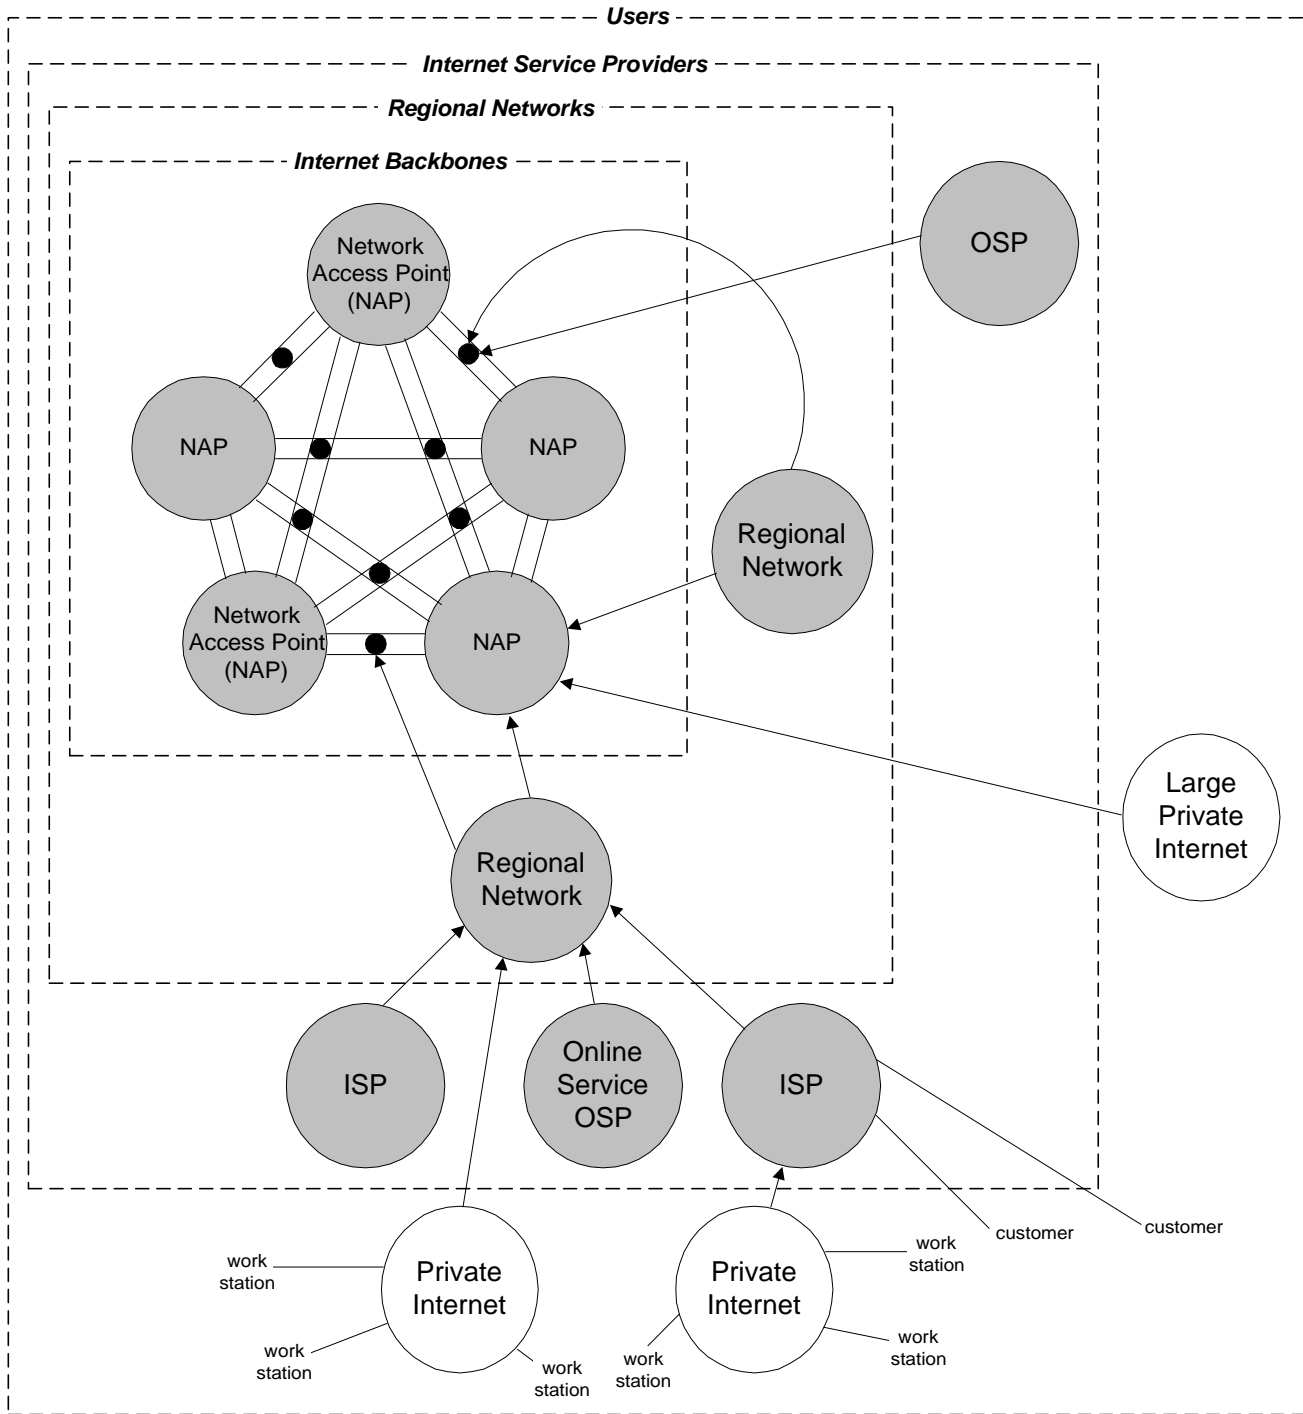
A subset of *ISPs* are the *online service providers (OSPs)*. *OSPs* provide enhanced services in addition to a connection to the Internet. These enhanced services may include proprietary data bases, *web sites*, interactive games, and *chat rooms*. An example of an *OSP* is America Online (AOL) or Microsoft Network.

Users consist of businesses, nonprofit agencies, governments, educational institutions, and individuals who use the Internet to communicate with other users. They usually pay an *ISP* for access to the Internet. Private individuals and some businesses generally connect to an *ISP* by dialing in over their telephone line. Many business, government, and educational entities have local area networks (LANs) or *intranets*. Often they have dedicated transmission lines that continuously connect them to their *ISP*. Larger *intranets* may connect directly to *regional networks* or *national backbones*.

⁴ The term “buy” may be a little simplistic. Some business entities that are *backbones* also own some *regional networks*. *Backbone providers* may also have arrangements with *regional networks* to use their systems as transmission “shortcuts.” Generally though, *regional networks* pay a *backbone* for access to the Internet core.

⁵ The levels in the schematic on the following page have been simplified. In the real world a business entity may be both a *regional network* and an *ISP*, or even a subsidiary of a *backbone provider*.

Figure 1: Physical Structure of the Internet



How is something transmitted over the Internet?

A message sent over the Internet is divided into individual parts called *packets*. Each *packet*, which may be up to 1,500 characters long, contains information on where it is sent from and to, how many *packets* are in the message, and which *packet* this one is (i.e., first, second, last) in the message. The *packets* are sent separately and reassembled at the other end. If some of the *packets* do not make it to their destination, the destination computer sends a message back to the originating computer requesting it to resend the missing *packets*.

Routers on the Internet handle each *packet* separately. This process is known as *packet switching*. At each *router* location, a *packet's* Internet address is examined and the *router* sends it on the next leg of the journey depending on its address and how crowded the transmission routes are. If a certain route is crowded or a portion of a route goes down, the *packets* for one message may travel several different routes.

How does the system know where to send a transmission?

The Internet standards have two methods for addressing transmissions: the *Internet protocol (IP) address* and the *domain name*. The *IP address* is a unique numeric address assigned to each piece of hardware (i.e., computer) permanently attached to the Internet. *Domain names* are alphanumeric names that can be linked to an *IP address* via a cross reference table. The *IP address* is what the system uses to actually figure out where to send a transmission; the *domain name* is an alias for the *IP address* because many users find numbers hard to remember.

An *IP address* consists of four numbers between zero and 255 which are separated by periods. An *IP address* looks like this:

156.99.41.200

Since the Internet is a network of networks, the leftmost number tells the Internet what network the recipient is on, followed by sub-networks, and finally computer. This is similar to the country code, area code, local exchange, and number for a telephone.

Most *routers* connected to the Internet will have routing tables that are fairly specific for *IP addresses* in their own network but may only contain routing information for the higher levels of other *IP addresses*. For example, the router map for a *router* may show that all *IP addresses* starting with a "206.145" are connected through a certain *backbone* and therefore route a *packet* with one of those addresses to the nearest/quickest connection to that *backbone* for distribution.

When people talk about an Internet address they are usually referring to the *domain name* rather than the *IP address*. A member of the Minnesota House of Representatives, if asked for their Internet address, would give the following answer:

jdoe@house.leg.state.mn.us

The series of “words” on the right hand side of the “@” which are separated by periods is the *domain name*. *Domain name* hierarchy is read from right to left. All *domain names* have as the rightmost word one of the top domain levels; either a country designation (in this case “us”) or an organizational designation (e.g., “com” for commercial, “edu” for higher education, “gov” for federal government). Each “word” as you read toward the left represents another level or subdomain within the top domain.

A company called Network Solutions Inc. currently oversees the registering of *domain names* in the United States, although different groups are actually responsible for maintaining the registries for the various domains and sub-domains. Other countries have their own registries. Computers called domain name servers (DNSs) are connected to the Internet at different levels throughout the system and provide the tables that link *domain names* with *IP addresses*.

***Domain names* are transportable but *IP addresses* are not.** *IP addresses* are permanently attached to a computer at a given location. However, a company that has registered a *domain name* may switch to a different *ISP* and although its *IP address* would change, its *domain name* would not. Network Solutions Inc. and the subsidiary registries supervised by them are responsible for updating the cross-reference tables used by the DNSs.

Is the Internet free?

It is a myth that the Internet is free or at least heavily subsidized by the federal government; this myth is probably due to government involvement in early Internet development. The myth is reinforced by the fact that the cost of Internet transmissions is relatively low compared to conventional telephony. However, participants in the Internet pay directly or indirectly for their own use of the system.⁶ In the case of many users, the cost of Internet use is subsidized by advertising sold by their *ISP*, *browser*, or *search engine*. This is similar to the financing of “free” network television by the sale of commercials.

Why is communication over the Internet so inexpensive?

Packet switching allows the Internet to transmit messages quickly and cheaply. With regular telephone service a dedicated connection is established between the two telephones; only those two telephones can use that piece of the transmission system until the connection is broken. Any pause in the communication, including the time spent receiving a message before replying, results in wasted transmission capacity. *Packet switching* allows *packets* from many different transmissions to flow across transmission lines in an unending stream. This dramatically increases the amount of information that can be transmitted over a transmission line, as compared to when the line is used for dedicated connections. Instead of one communication occurring over this part of the transmission system, *packet switching* allows 100 communications

⁶ Actually, an argument can be made that the Internet is subsidized to a limited degree by the traditional telecommunication industry that share the same local transmission lines. This is because the Internet is exempt from the universal service fund charges and federal access charges that are imposed on regular telephony. The Federal Communications Commission is re-examining these exemptions.

to occur simultaneously over the same piece of the transmission system. *Bandwidth* describes the amount of information that can be transmitted across any line, and it is limited only by *router* speeds and type of transmission line.⁷

Is the *World Wide Web (WWW)* just another name for the Internet?

The *World Wide Web (WWW)* or *Web* is an application system that runs on the Internet. It is the *hypertext*-based information service on the Internet. The *Web* was developed in 1989 by a researcher in Switzerland and is extremely useful because of three characteristics:

- the use of *hypertext* to link different documents at many different sites;
- the integration of graphics, video, and sound with text documents; and
- a *uniform resource locator (URL)* protocol to access information created through Internet protocols other than the *World Wide Web*.

The development of the World Wide Web with its multimedia *web sites* and interactive links, along with easy to use *web browsers* such as Netscape Navigator, made the Internet increasingly attractive for commercial use. Its development was a major impetus to the explosive growth of the Internet in the 1990s.

How is economic activity affected by the development of the Internet?

The Internet is becoming an increasingly important factor in the national economy, spurring the development of new types of businesses, changing how business is conducted, and potentially making some current products and services obsolete.

Many different types of companies are involved in providing goods and services that are necessary to operating the Internet. These goods and services can be divided into three basic types: physical infrastructure, access provision, and applications. The physical infrastructure is provided by the telephone and cable companies that provide the transmission lines and the manufacturers of hardware such as *routers*, switches, and host computers. Access is provided by the *ISPs* and the *Internet backbone providers* that take the physical infrastructure and the established computer protocols and use them to interconnect the networks that make up the Internet in a way that allows communication between the various users. Software companies and electronic commerce development companies provide the applications that make the Internet a useful tool for business and consumers.

⁷ There are two main types of transmission lines, copper twisted pair line or fiber optic. Hooking into the Internet using a standard modem and plain old telephone service (POTS) allows information to be transmitted at a speed between 33 and 56 kilobits per second (kbps). The “fastest” copper line (called a DS-3 line) carries 45 megabits per second (mbps) or more than 1,000 times as much. An OC-12 level fiber optic line can carry 622 mbps.

A major area of growth is in specialized computer software companies that develop software to meet needs generated by Internet activity. Examples of the types of software include:

- graphical *browsers* such as Netscape Navigator which provide a user-friendly way for people to move around the Internet;
- *firewall* software which acts as a gatekeeper to limit unauthorized access to information on a computer or *intranet* that is hooked to the Internet; and
- *encryption* software which allows secure transactions over the Internet.

A growing Internet-related industry is the provision of electronic “value-added” information. This includes any information on the Internet that has been summarized or organized to be easier for consumers to find and use. Some of this value-added information is in proprietary data bases that the user must pay a fee to access, while some information is provided “free” but paid for by advertising included with the information. An important subset of this industry is the design and maintenance of *search engines*. *Search engines*, such as “Yahoo!,” are sites on the Internet that act as “portals” to allow users to easily find Internet sites that contain information on specific topics.

The business of advertising is undergoing major changes due to the Internet. In addition to the development of electronic advertising on *browsers* and value-added information sites, company web pages are a new venue for advertising. Companies now exist that design and maintain web pages for other entities and/or rent computer space for web pages on the company’s server.

One of the most talked about issues relating to the Internet is the growth of *electronic commerce*. At its simplest, this is an electronic version of mail-order catalogues. However, there is also the development of *virtual equivalents*, or electronic goods transmitted over the Internet that replace tangible goods or products. Examples of *virtual equivalents* include downloaded software, and electronic versions of books, music, and videos. *Electronic commerce* also includes the provision of services through the Internet. Services that can be provided over the Internet include banking, investment, insurance, and education.

Finally, telephony, cable television, and the Internet may cease to exist as separate industries as they start to provide the same services and share technologies. Currently the Internet uses the same transmission system as the telephone, but the cable television transmission system is equally adaptable to delivering Internet services. The technology now exists to provide voice telephone service via the Internet and telephone companies are beginning to move to *packet switching* for some of their service provision. In the future, television may be delivered via the Internet as well as via cable and a satellite. As it becomes possible to deliver the same services in multiple ways, the government needs to re-examine how it currently treats these industries differently regarding regulation and taxation.

Glossary of Internet Terms

The following glossary only includes Internet terms used in this information brief. A source for definitions of other computer and Internet-related terms is the free online dictionary of computing located at the following address: <http://wombat.doc.ic.ac.uk/foldoc/>.

ARPAnet (Advanced Research Projects Agency Network): A predecessor of the Internet and the first wide area network to use *TCP/IP* for data transmission. In the early 1970s, the U.S. Department of Defense funded this project to develop a means of communication in the case of a war or other emergency that knocked out a portion of the normal communications system.

Backbone: A set of high-speed *TCP/IP routers* and interconnecting transmission lines that form a major trunk system for carrying multiple transmissions. Local and *regional networks* connect to backbones which then provide the long distance connections to other networks that are also attached to the same backbone or another backbone. Backbones are connected to each other at *NAPs* and private exchange points.

Bandwidth: The capacity of a transmission line in a given amount of time. Bandwidth is generally measured in bits per second (bps). A T-1 dedicated copper telephone line that might be used to hook a business to the Internet can carry 1.544 megabits per second. An OC-12 fiber optic cable can carry 622 megabits per second or more than 400 times as much.

Browser (web browser): A software program that allows a user to view and download information from web sites. The most well-known examples are Netscape Navigator and Microsoft Internet Explorer.

Chat rooms: An electronic place where multiple users with similar interests can hold simultaneous “real-time” conversations by inputting text via their computer. Chat rooms are usually dedicated to conversations on a specific topic or area of interest. Chat rooms are electronic meeting or gathering places.

Domain name: The hierarchic alphanumeric name or names that correspond to an *IP address* (site) on the Internet. Domain names are managed by the Domain Name System (DNS). All domain names may contain several sub-domain names but all end in a root domain which is one of the following:

- ▶ .com (commercial sites)
- ▶ .edu (higher education sites)
- ▶ .gov (federal government sites)
- ▶ .mil (military sites)
- ▶ .net (network resources)
- ▶ .org (other noncommercial organizations)
- ▶ .us, .uk, or any other two-letter country code (a site registered in that country)

Although domain names can share parts, similar to people having the same first name, each registered domain name is unique. For example “wwf.org” and “wwf.com” are each unique

domain names (note the different root domains). The first will connect you to the web site of the World Wildlife Fund whereas the second will connect you to the World Wrestling Federation.

Electronic commerce: Commerce conducted through electronic transmissions. This phrase is sometimes used carelessly and means different things to different people. Often it is used to refer only to the sale of goods and services over the Internet. In other cases, electronic commerce also includes the sale of access to the Internet, and the sale of advertising on the Internet. In its broadest sense it includes all commerce facilitated by electronic means, not limited solely to the Internet. Withdrawing money from an automatic teller machines (ATM) would be considered electronic commerce in its broader sense.

Encryption: A process of coding information so that (a) only the intended recipient can read it, and/or (b) it can be identified as having come from a specific sender. Various encryption software and schemes exist but this is still a developing area. “Unbreakable” encryption is an important issue for the online growth of certain businesses, such as banking. On the other hand, law enforcement entities are concerned that “unbreakable” encryption might make it impossible to identify certain organized criminal activities (i.e., drug dealing, money laundering) conducted via the Internet.

Flaming: The sending of vitriolic messages to a user, usually in response to an activity deemed inappropriate by the sender. Persons who send these messages are called “flamers.” Activities that can lead to flaming are sending *spam* or engaging in an activity that causes a logjam on a portion of a system. On a more personal note, flaming is also done in response to a posting to a bulletin board or *chat room* that another user disagrees with or finds offensive.

Firewall: A computer and associated software that limits the amount of access outside users have to data and information contained on a private network or computer. The object of a firewall is to stop unauthorized electronic intruders (computer hackers) from getting confidential information or destroying computer files and systems.

Home page: Any *HTML* document on the *World Wide Web*. Because of the availability of *HTML* editors it is relatively easy for anyone to create an *HTML* document. This ease of creation has led to an explosion of home pages which range from the professional-looking electronic gateway to a company to the idiosyncratic expression of a person’s individuality. Home page is also sometimes used to refer to the top or main page of a *web site*.

HTML (hypertext markup language): The *hypertext* language used to create and link documents on the *World Wide Web*.

Hypertext: Text that contains links to other documents or other places in the same document. Selecting a link will automatically display the second document or move you to the indicated place in the original document.

Internet backbone provider: A company with high-speed *TCP/IP routers* located in a number of different major locations (including *NAPs*) that are linked together with high-speed data transmission lines that it either owns or leases. These transmission routes are called *backbones*.

A company is only a backbone provider if it is recognized as such by the other companies that are backbone providers.

Intranet: A company or group mini-version of the Internet. An intranet is an integrated information system using Internet technology that is deployed on either a local area network or a wide area network of computers. An intranet could link an entire company or government agency, or just one department, or a multi-corporate group. The computers linked to the intranet can share resources and information with each other in the same fashion as the Internet. However, there is usually a *firewall* at the connection of the intranet to the Internet that limits access by outside users to certain types of files and data on the intranet.

Internet service provider (ISP): A company that connects users to the Internet. The connection can be either a direct line or a dial-up connection via a modem. ISPs vary in size from small local operations to entities that serve a national market. *Online service providers (OSPs)* are a specialized type of ISP.

IP (Internet protocol) address: The numerical address assigned to a computer or other piece of electronic equipment that is hooked in to the Internet. Each address consists of four numbers separated by periods. Each number is between zero and 255 (the numbers of combinations possible with eight bits (on-off switches) in the binary code used as the basis for all computer languages). The numbers are hierarchical from left to right, similar to a telephone number which starts, with a country code, then the area code, and then the local exchange.

Network access point (NAP): The network access points are the major public interconnection points for the Internet *backbone* providers. There are four official NSF-endorsed NAPs, one each in San Francisco, Chicago, New York (actually located in Pennsauken, New Jersey), and Washington, D.C.. Three historical NAPs which are leftover from precursors to the commercial Internet are at the University of Maryland, NASA Ames Research Center in California, and Santa Clara, California. Four “unofficial” NAPs are operated by Metropolitan Fiber Systems Inc. in San Jose, Dallas, Chicago, and Washington, D.C.

NSFnet: NSFnet was a *backbone* developed in the 1980s by the National Science Foundation to connect universities to its five supercomputer centers. The growth in commercial Internet use in the early 1990s caused the National Science Foundation to get out of the *backbone* business. The NSFnet was the core of the U.S. Internet in the pre-commercial era between 1987 and 1995.

Online service provider (OSP): A subset of *ISPs* which provide a variety of value-added services in addition to Internet access. These services can include items such as proprietary data bases, *chat rooms*, interactive games, and special e-mail between subscribers.

Packet: A bundle of data. Each transmission on the Internet is broken into pieces called packets that are generally no longer than 1,500 characters. Each packet includes information on the sending address, the destination address, the number of total packets in the transmission, and where this packet fits in the transmission (i.e., first, second, last). Each packet is sent independently through the Internet and the transmission is reassembled at its destination.

Packet switching: The method used to send information over the Internet. *Packets* or pieces of information from lots of different transmissions flow constantly across a transmission line. Whenever that transmission line intersects or connects with other transmission lines, a *router* computer reads the *IP address* on each *packet* and sends it on the next leg of its journey to its destination. The route chosen will be based on both the final destination and the congestion on the various connected transmission lines at that instant. Different *packets* from the same transmission sometimes take different routes to their destination.

Point of Presence (POP): The physical point where a user can connect to an *ISP's* hardware and software and hence into the Internet. The connection can be made either through a dedicated line or a telephone call (dial-up connection).

Regional network: A network of computers and transmission lines that links a limited geographic area. Although the region served by the network may be quite large and provide much of the transmission service in an area, it is not recognized as an *Internet backbone provider* by the other *Internet backbone providers*.

Routers: Computer hardware and software that directs the flow of information *packets* over the Internet. Routers exist at the intersection or connection points of the transmission lines used by the Internet.

Search engine: A program that lets you do keyword searches for information on the Internet. The user inputs the criteria for the search and the search engine gives a list of Internet sites that meet the criteria. Different search engines will often return different lists of sites even if the search criteria used are identical.

Spamming: Indiscriminate mass transmission of unsolicited information, usually advertising or a sales pitch, over the Internet. It is the electronic equivalent to bulk mailing of advertising to an entire city.

TCP/IP (transfer control protocol/internet protocol): This is the system of communication rules (protocol) used for exchanging data between computers on the Internet. TCP/IP are the rules regarding how to send data and how to handle errors in the transmission. It was developed for the U.S. Defense Department's *ARPAnet* in 1970.

Universal resource locator (URL): The *World Wide Web* system of addresses that allows linking and moving between various documents on the Internet. A URL takes the following form:

protocol://host/pathname/filename/

The first part (protocol) indicates the type of Internet document it is, the second part (host) is the *domain name* for the computer on which the document resides, and the last two parts (pathname and filename), which are optional, connect you to the location of the document on the host computer. The URL for the House Research Department *home page* is <http://www.leg.state.mn.us/hrd/hrd.htm/>.

Virtual equivalents: An intangible, electronic replacement for an existing good. The most commonly used example is a virtual book. A book is a document made up of text and illustrations that is printed on paper and bound together. A virtual book also consists of text and illustrations but it has no physical substance; it exists merely as an electronic stream of information flowing across a computer screen.

World Wide Web (WWW): The principal graphical portion of the Internet. The World Wide Web is based on a protocol conceived of by Tim Berners-Lee of the European Particle Physics Laboratory (CERN). The web uses *hypertext markup language (html)* to construct and link Web documents with other documents on the Internet. What makes the Web particularly useful is that links are not limited to other *hypertext* documents on the same computer; it can link to document located on any computer host on the Internet anywhere in the world.