INFORMATION SYSTEMS SERVICES

Overview of the Internet

An introduction to the Internet, Internet concepts, the programs necessary to use the Internet and how to connect to the Internet.

AUTHOR: Information Systems Services

DATE: March 2001

EDITION: 2.2

OVE 8

50p

UNIVERSITY OF LEEDS

Contents

Format Conventions

In this document the following format conventions are used:

Commands that you must type in are shown in **bold Courier** font.	`WIN31`
Menu items are given in a **Bold, Arial** font.	**Windows Applications**
Keys that you press are enclosed in angle brackets.	**<Enter>**

Feedback

If you notice any mistakes in this document please contact the Information Officer. Email should be sent to the address **info-officer@leeds.ac.uk**

Print Record

This document was printed on 16-Jun-03.

Introduction

There has been a lot of talk in the press about the "Information Superhighway". People talk about "cruising the net", "surfing in Cyberspace" and a number of other largely meaningless and highly jargonised phrases. But for many people, even some computer users, it is still a thing surrounded in mystique and, to an extent, suspicion.

Unfortunately, getting access to this 'Internet' isn't that obvious. There are some difficult concepts to overcome, often some obstructive hardware to grapple with. Consequently many who have managed to conquer their fear and lack of knowledge when they first came to computing now find themselves in exactly the same situation again facing the world of international computing.

This document has been compiled to help such people. In preparing this overview extensive use has been made of a particular document, details of which are on the back page. Material from the Sunday Times "Getting Wired" supplement for 14th May 1995 is also used. The University of Leeds cannot be held liable for any misinformation produced here.

The aim of this document is to explain:

What is the Internet

A brief explanation of physically what it is and how it came about.

What can I do with the Internet

To whet your appetite. Also, because the Internet is best described from a functional point of view ("What do I want to do") rather than from a list of available software and what it does.

How do I get connected to the Internet

Having whet your appetites, this is how you actually get on.

This document will not explain how to use any of the software programs involved when using the Internet, otherwise it would be huge. Many of those programs are documented separately in other ISS documentation and so may be explored separately.

If you want a more detailed introduction to the Internet and what it can do you are recommended to get a full copy of the "EFF's Guide to the Internet" by Adam Gaffin, on which this document is largely based. Details of how to get hold of the document are given in the Appendix.

What is the Internet?

To explain what the Internet is, we'll start from the bottom up.

What is a Network?

Most of the time, PC users work on their own with their PC connected to nothing more exciting than a printer. However, it has always been possible to connect two PCs with a piece of wire so that they can communicate – usually just to transfer files.
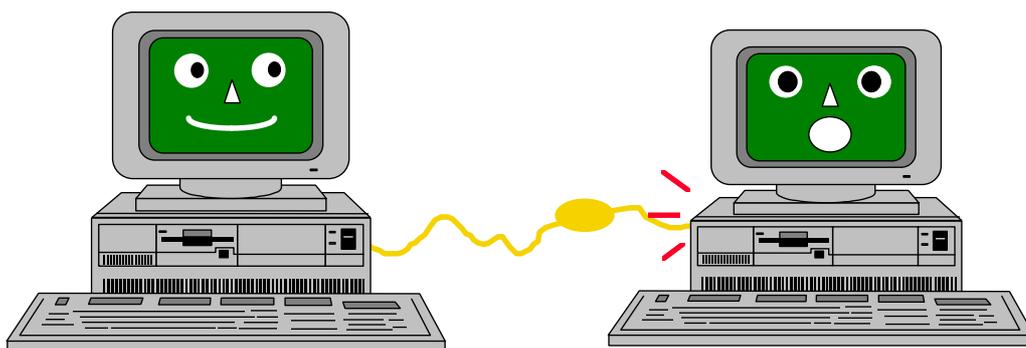
This, in essence, is all a network is – two or more computers connected by a physical piece of wire. The connection between the two computers is then managed and operated by some special software.

Local Area Networks

It makes sense, most often for financial reasons but also for others, to network groups of computers where they share a common workload. All the computers in an administrative office, all the computers to do with a certain ward or discipline. Networking computers means that the people using them can share files easily, send each other messages and share each other's printers. This idea has developed into Local Area Networks (LANs). Nowadays most organisations have a local area network. LANs can be as small as just one shared office or as large as a whole city.
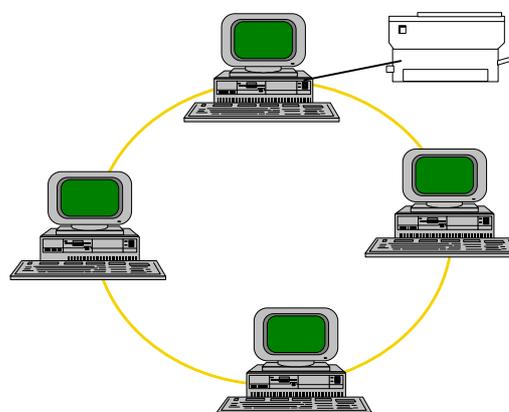
Wide Area Networks

In some cases an organisation is spread over a large area, and you do not have the easy concentration of computing to provide a LAN for. In this instance computers may be connected by a Wide Area Network (WAN). The difference between a WAN and a LAN is partly one of scale (although this is relative) but also relates to the technology. With a LAN you will typically get a fast network that can network PC file servers. With a WAN the network will often be much slower and will usually involve some mainframe computer as the server rather than a PC.

Combinations of Networks
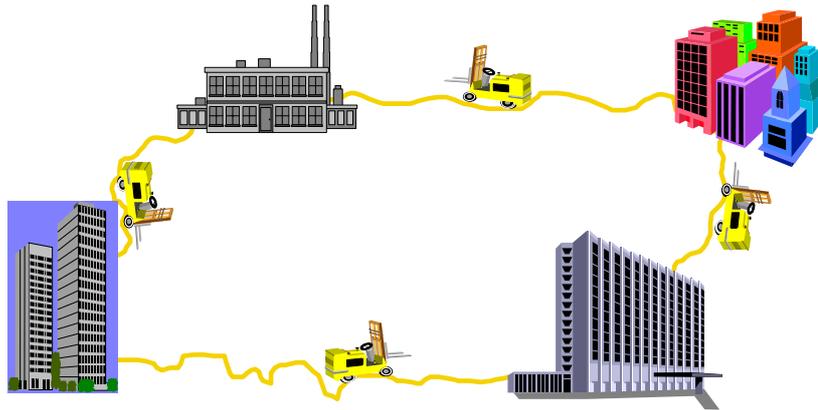
LANs and WANs are not mutually exclusive. In fact they combine very well together. Many organisations now have both, where a WAN has been created by connecting up a series of LANs. A good example of this is the Academic Community. Every University has it's own LAN providing email, printing, file sharing and other facilities. Then every University is connected to the Joint Academic Network (JANET) which is a WAN running the length of the country. Through JANET someone in Leeds can log into a mainframe computer in Southampton.

The Internet

Big fleas have little fleas upon their backs to bit 'em,
The little fleas have smaller fleas, and on ad infinitem.
Great fleas have greater fleas, upon whose backs they ride on,
The greater fleas have greater yet, and greater yet, and so on

If you take the principle of small, locally managed, networks connected up to form a larger network, which in turn is connected up to form a larger network, and so on, you get some idea of what the Internet is. Physically it is just a collection of LANs and WANs that have been connected up at a national and international level.

The Internet is more than this, however. It not only is the physical wire, but also how information is sent from one computer to another. Take the example of making a road journey. If you are sending someone in a car to go somewhere, you can do one of two things:

- give them precise written details of how to get to their destination. As they have no other means of finding their way they will have to stick to your instructions even if they can see roadworks ahead.

- give them a map and let them find their own way. If roadworks, or some other blockage occurs then the person can either read the map and work out a fresh route, or else follow the diversion signs.

Originally information on networks followed the first method – it was given precise instructions on how to reach its destination. However, if there was a hold-up, say a computer was switched off, the information was also held up until the hold-up cleared.

Now information is more intelligent and can work out its own route, so if there is a hold-up it will try and get through some other way. It's all rather clever, and was invented by the American Defence Industry who asked "If someone dropped a very large bomb on our computer network, would it still function?". The intelligent information transfer – a mechanism called the Internetworking Protocol (hence "Internet") – means now that it can.

What do you do with a Network?

Talk to other people

A traveller was very impressed by a hotel in Hong Kong. Upon arrival at the hotel from the airport, he checked in, and went to his room. When the TV was switched on there was a message waiting for him on the TV screen.

This system is available to everyone who uses the Internet. To get on to the Internet you have to have an identification unique to you called your *username*. Everyone who connects to the Internet has a username, and if you have their username and the name of the machine they log into you can send a message to that username and they will receive it when they next log in. This system is called *Electronic Mail* or *email* for short.

The basic concepts behind email parallel those of regular mail. You send mail to people at their particular addresses. In turn, they write to you at your email address. You can subscribe to the electronic equivalent of magazines and newspapers. Sooner or later, you'll probably even get electronic junk mail.

Email has two distinct advantages over regular mail. The most obvious is speed. Instead of several days, your message can reach the other side of the world in hours, minutes or even seconds. The other advantage is that once you master the basics, you'll be able to use email to access databases and file libraries as well as transfer files and programs.

Email also has advantages over the telephone. You send your message when it's convenient for you. Your recipients respond at their convenience. No more frustration whilst you phone someone who's out, then they phone you back whilst you're out. And while a phone call across the country or around the world can quickly result in huge phone bills, email lets you exchange vast amounts of mail for only a few pennies – even if the other person is on the other side of the earth.

Email is your connection to help – your Net lifeline. The Net can sometimes seem a frustrating place! No matter how hard you try, no matter where you look, you just might not be able to find the answer to whatever is causing you problems. But when you know how to use email, help is often just a few keystrokes away: you can ask your system administrator or a friend for help in an email message.

Even if you use the Internet for no other purpose, email makes getting a connection worthwhile all by itself. Very quickly you'll find yourself writing to friends who are on the Internet, contacting colleagues interested in similar professional areas all round the world, discussing politics and asking for advice on how to use some of the software on your computer. It is very powerful and incredibly cheap.

Many-to-Many Discussions

If you want to have an individual discussion with someone about a particular topic, then email is perfect for that. It's a one-to-one communication tool. However, you sometimes want to have a many-to-many discussion, like you would in a meeting. Email can do this but is not very convenient. Instead there is a standard system for many to many discussions called *Usenet*.

Usenet is very like email, in that you compose a message and send it somewhere. However, instead of being sent to a person the message is sent to your local Usenet system. That way the messages are stored centrally, once, for everyone to read, rather than have every message clogging up each individual user's email space. The local system is in constant contact with other Usenet systems around the world, they compare messages and exchange the ones they don't have. That way your message to the local system in Bradford-on-Avon gets passed from computer to computer round the world for everyone to see whenever they access their local Usenet system.

The basic building block of Usenet is the newsgroup, which is a collection of messages with a related theme (on other networks, these would be called conferences, forums, bboards or special-interest groups). There are now many thousands of these newsgroups, in several different languages, covering everything from art to zoology, from science fiction to South Africa.

Some public-access systems, typically the ones that work through menus, try to make it easier by dividing Usenet into several broad categories. Choose one of those and you're given a list of newsgroups in that category. Then select the newsgroup you're interested in and start reading.

Other systems let you compile your own "reading list" so that you only see messages in conferences you want. In both cases, newsgroups are arranged in a particular hierarchy devised in the early 1980s. Newsgroup names start with one of a series of broad topic names. For example, newsgroups beginning with "comp." are about computer-related topics. These broad topics are followed by a series of more focused topics (so that "comp.unix" groups are limited to discussion about UNIX). The main hierarchies are:

- bionet   Research biology
- bit.listserv Conferences originating as Bitnet mailing lists
- biz   Business
- comp   Computers and related subjects
- misc   Discussions that don't fit anywhere else
- news   News about Usenet itself
- rec   Hobbies, games and recreation
- sci   Science other than research biology
- soc   "Social" groups, often ethnically related
- talk   Politics and related topics
- alt   Controversial or unusual topics; not carried by all sites

In addition, many host systems carry newsgroups for a particular city, state or region. For example, ne.housing is a newsgroup where New Englanders can look for apartments. A growing number also carry newsgroups which are aimed at elementary and secondary teachers and students.

Use other people's computer programs

The most powerful thing on the Internet has to be the information that is there. Some of this information you can simply transfer to your own computer and use locally, some you can't. So the next most powerful thing on the Internet is the software that is installed on all the different computers. The only way to access that software is to actually log in to the computer.

The local software you use to log in to other computers is called *telnet*. Telnet is the electronic equivalent of walking up to the computer's front door, knocking, identifying yourself, and walking in. Once you are in, there will be a variety of programs available to you, at the discretion of the computer owners. What you are able to do depends on the computer, but telnet itself will just give you a simple text-based screen interface, known as *vt100 emulation*. Don't expect to start running Microsoft Windows programs directly off a computer in Katmandu!

Telnet is a program that lets you use the power of the Internet to connect to databases, library catalogues, and other information resources around the world. Want to see what the weather's like in Vermont? Check on crop conditions in Azerbaijan? Get more information about somebody whose name you've seen online? Telnet lets you do this, and more. Increasingly, information once available only via telnet is being disseminated by the World Wide Web (see later) – but it can still be useful to learn your way around telnet.

Alas, there's a big "but!" Unlike the phone system, Internet is not yet universal; not everybody can use all of its services. Almost all colleges and universities on the Internet provide telnet access, as do many commercial organisations. Free-Net systems[1] do not give you access to every telnet system and if you are using a public-access UUCP or Usenet site, you will not have access to telnet. The main reason for this is cost.

A good example of the typical use of telnet is in querying large databases. Some of these are open access to anyone who wants to query them (many library catalogueues offer this facility) but many have to be paid for before you get a username and password on to the system.

What places to telnet to, and what services they offer, you have to find out by word of mouth, or through literature. Just start asking!

---

[1] Free-Net systems are where some benevolent groups of people allow you free access to their computer and from there on to the Internet. Whilst clearly being a cheap way to do things for the end user these systems are often quite limited in what they can do. They also only occur in America.

Copy information and programs to your own computer

Hundreds of systems connected to Internet have file libraries, or archives, accessible to the public. Much of this consists of free or low-cost shareware programs for virtually every make of computer. If you want a different communications program for your IBM, or feel like playing a new game on your Amiga, you'll be able to get it from the Net.

But there are also libraries of documents as well. If you want a copy of a recent U.S. Supreme Court decision, you can find it on the Net. Copies of historical documents, from the Magna Carta to the Declaration of Independence are also yours for the asking, along with a translation of a telegram from Lenin ordering the execution of rebellious peasants. You can also find song lyrics, poems, even summaries of every "Lost in Space" episode ever made. You can also find extensive files detailing everything you could ever possibly want to know about the Net itself.

The main way of getting files across the Net is using a mechanism called *file-transfer protocol*, or *ftp*. This is a client server program (see below for details of client-server software) because you have to run the ftp client software on your computer and the computer that you want the files from has to be running the ftp server software so that it can service your request.

Ftp is very quick and allows you to get several files down at once. In its UNIX incarnation it has quite a difficult user interface and so is not so easy to use. There are Windows programs that look rather like the Windows Explorer and allow you to simply drag and drop files from the server computer and your own.

How do you find a file you want, though?

Until a few years ago, this could be quite a pain – there was no master directory to tell you where a given file might be stored on the Net. Who'd want to slog through hundreds of file directories looking for something?

Three students at McGill University in Montreal, asked the same question – then did something about it. They created a database system, called *Archie*, that would periodically call up file libraries and basically find out what they had available. In turn, anybody could dial into the Archie database, type in a file name, and see where on the Net it was available. Archie currently catalogues more than 1,000 file libraries around the world.

Today, there are three ways to ask archie to find a file for you: through telnet, a 'client' Archie program on your own host system or via email. All three methods let you type in a full or partial file name and will tell you where on the Net it's stored. Using Archie is just like filling in a form of your requirements (like the name of the file), like you would any database, and it then returns a list of all the files it has found and the names of the computers where the files are stored.

The Problem of Addresses

Even with tools to help you, telnet and ftp can still be frustrating. There are all those telnet and ftp addresses to remember. Telnet services often have their own unique commands. And, oh, those weird directory and file names!

But now that the Net has become a rich repository of information, people are developing ways to make it far easier to find and retrieve information and files. The principle methods here are Gophers, Wide-Area Information Servers (WAISs) and most recently the World Wide Web.

Gophers

Gophers essentially take a request for information and then scan the Net for it, so you don't have to. Gophers work through menus – instead of typing in some long sequence of characters, you just move a cursor to your choice and hit enter. Gophers even let you select files and programs from ftp sites this way.

Gophers are great for exploring the resources of the Net. Just keep making choices to see what pops up. Play with it; see where it takes you. Some choices will be documents. When you read one of these and either come to the end or hit a lower-case **<q>** to quit reading it, you'll be given the choice of saving a copy to your home directory or emailing it to yourself. Other choices are simple databases that let you enter a word to look for in a particular database. To get back to where you started on a gopher, hit your **<u>** key at a menu prompt, which will move you back 'up' through the gopher menu structure (much like "cd .." in ftp).

In addition to ftp sites, there are hundreds of databases and libraries around the world accessible through gophers. There is not yet a common gopher interface for library catalogues, so be prepared to follow the online directions more closely when you use gopher to connect to one.

Gopher menu entries that end in a / are gateways to another menu of options. Entries that end in a period are text, graphics or program files, which you can retrieve to your home directory (or email to yourself or to somebody else). A line that ends in <?> or <CSO> represents a request you can make to a database for information. The difference is that <?> entries call up one-line interfaces in which you can search for a keyword or words, while <CSO> brings up an electronic form with several fields for you to fill out (you might see this in online "White Pages" directories at colleges).

The point about gophers is that you are dealing with what interests you – the information. Mostly you are not interested *where* the information resides, but simply in getting at it. Gopher hides all the problems of telnetting to obscure sites, changing directory and ftp-ing with arcane commands behind a menu of choices where you choose the information and gopher delivers it. What this results in is a potentially limitless mine of information that can span thousands of computers across many continents but that is very easily accessible.

The World Wide Web (WWW or W3)

Nice as gophers are, there's an even better way to navigate and find information resources on the Net – the World-Wide Web.

Originally developed as a resource for physicists, the Web today is fast becoming the Main Street of cyberspace. You'll find interesting characters wandering around, museums and galleries to visit, schools to teach you new skills, even restaurants (some of which will deliver real food in response to email). You name it, chances are somebody's created a Web server about it. Growing numbers of people even have their own personal Web 'pages' where they let the world know what they're interested in.

The Web has exploded in popularity for two reasons. One is that it is fairly easy to use. As with gophers, you navigate the Web by making selections from your screen – no more cryptic UNIX commands to memorise or addresses to type in. In particular, the Web software allows you to access everything set up under the gopher system, as well as connect to ftp sites and use ftp, login to other computers with telnet and read Usenet news groups.

But what really sets the Web apart is hyperlinks. To understand hyperlinks, think of an encyclopaedia. As you read an article on, say, Africa, your eye is drawn to a picture of an elephant. You want to learn more about the animal, so you get out the "E" volume and look up "elephant" and start reading. The Windows help system works in exactly this way, as do many CD-Rom encyclopaedias.

Hyperlinks are the online equivalent of this browsing process. Tim Berners-Lee, who developed the original Web model, came up with a simple language that lets somebody developing a Web document embed pointers to related resources. When you then call up that document, you'll see some words in a different color or somehow otherwise highlighted. By moving your cursor to one of those words and then hitting enter (or clicking on it with your mouse, depending on your interface) you'll then call up the linked document.

These hyperlinks are easy to create so anybody can put together a Web resource that can become a central clearing house of information on virtually any topic, linking documents that could be physically stored in dozens of locations around the world. The language that Tim Berners-Lee developed is a version of the markup language, SGML, which he called HTML (HyperText Markup Language). HTML is a proper formatting language but it is relatively easy to learn. What is even better, though, is that most of the main Word Processing software companies are now providing HTML converters so that you can create the documents in, say, Word for Windows and then save them as HTML to put into your Web server.

The Web can be browsed using either a text based browser (e.g. Lynx) or a graphical browser (e.g. Netscape). If you can, you are encouraged to use a graphical browser as the Web not only gives you access to text-based information but can also provide pictures, sound and movie clips. To make use of these, however, you need a computer capable of using a graphical browser (4Mb 386 PC or 4Mb Apple Mac) and an appropriate Internet connection (direct, SLIP or PPP connection) and one fast enough to cope with large graphics files. If you are using a 14,400 bps connection you may find using the Web to get graphics files a little slow!

Clients and Servers

One of the concepts that you will have to get your brain around whilst working on the Internet is the idea of *client-server*. Most software on the Internet works on this principle, so it is worth taking some time to explain it.

The client-server model of computing is a common one in networking. It consists essentially of two computers: the client and the server. The client is intelligent, is the one initiating the action and is where the control lies. The server is less intelligent but is there to provide a service on demand. Typically the server is a very large computer that can handle several clients all at once.
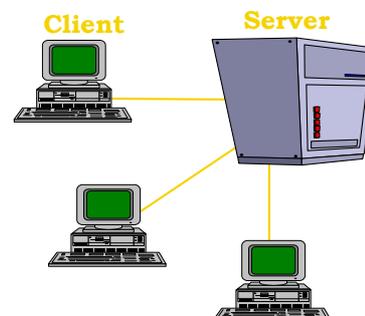
In terms of the Internet, you are almost always the client. You are the one who wants to initiate the transfer of files, the sending of messages, the querying of a database. Every other computer you connect to will be a server. Now in some cases the roles can be reversed, but those are rare and are too advanced for this document.

The thing about client-server is that it is software-specific. For instance, if I want to copy a file from an Internet computer, I would run a program called ftp. For it to work, I must run the *client ftp* on my computer and the computer I wish to copy the file from must be running the *server ftp*. That means that when I ask my ftp software to transfer the file, the client ftp software can then talk to the server ftp software and ask for the file. If I wasn't running the client ftp then I couldn't do the asking (self-evidently), and if the server computer wasn't running server ftp then it wouldn't understand what the client ftp software was saying.

The reason this is important is because most software on the Internet relies on this client-server model. All the following programs require you to run the client part on your PC or Macintosh, but you will only be able to use on computers running the server part.

- ftp
- telnet
- email
- Usenet news
- gopher
- Web

Consequently, you will hear people talking about mail clients (which is the program you use to actually type your mail message), Web servers (the box that actually holds the HTML information and services Web requests) and Web clients, or *browsers*.

How Do I Get On?

There are essentially two types of access to the Internet – direct access or dial-in access. Dial-in access is dealt with in a separate document (FAC 16).

Direct access is available to people whose employers have connected their own computers to the Internet. A good example is our own academic community. Every University has its own local area network. All the separate local area networks are connected together by a Wide Area Network we call JANET. The JANET network is then connected on to the Internet. Consequently anyone who has access to their University network also has access directly out to JANET and the Internet beyond. To connect you must have a computer with a network card and software that allows your computer to 'talk' to your local network.

For staff and students at the University of Leeds, access to the Internet is provided through both our PC-based Novell systems and through various UNIX-based systems. Both systems have a variety of software for using the Internet, including email software, Usenet software and Web browsing software.

On PC Novell systems the best Internet software is available through the Windows desktop. The following applications can be found as desktop icons or under the **Internet Tools** submenu:

- Pegasus mail   Email software
- Netscape   Web browsing software
- News Xpress   Usenet news browsing software

Under the **Network Tools** submenu there is a collection of other tools for using the Network. The important ones here are:

- Telnet   A telnet program
- Ftp   A graphical file transfer program

Alternatively, if you do not want to, or cannot, use Windows there are a few DOS-based programs available[2] once you are logged in to the Novell systems:

- Pegasus mail   DOS version of the email software
- telnet   Telnet program
- ftp   File transfer program

Note that telnet and ftp for both DOS and Windows will only connect you to computers in the University of Leeds unless your PC has special rights to access the Internet directly. This 'right' is called "JIPS Access" and is obtained by applying to the Help Desk. Only secure PCs in offices can be given JIPS access, so cluster PCs can never directly telnet or ftp to a computer outside Leeds. The problem ftp-ing files can often be got round by using Netscape that will allow it, but telnetting from a cluster PC outside Leeds is strictly forbidden and prevented. Users who want to be able to telnet to other computers must either obtain a GPS username or else a computer which can be authorised for JIPS access. It should also be noted that if you have a PC with JIPS access then you can potentially load Internet browsing software on your local copy of Windows and use it from there without having to log in to the Novell network first.

---

[2] DOS-based programs are no longer supported by the ISS

Under UNIX there is both line-based and X Windows-based software available. The line-based software is as follows:

- Lynx   Line-based Web browsing software
- Mail and Elm   Email software, Elm being the much better one
- Telnet   Telnet software
- ftp   File transfer software
- Archie   Archive searching program for finding files to ftp.
- rn   Usenet news reader

The X Windows software is very similar:

- mailtool   X Windows mail
- Netscape   Web browsing software, almost identical to the MS Windows version
- xrn   X Windows news reader
- xgopher   Gopher browsing software
- xarchie   File archive browsing software that will also ftp the file

Internet Jargon Terms

Like any community, the Net has developed its own language. What follows is a glossary of some of the more common phrases you are likely to encounter. But it's only a small subset of *net.speak*. You can find a more complete listing in "The New Hacker's Dictionary," compiled by Eric Raymond (MIT Press). Raymond's work is based on an online reference known as "The Jargon File," which you can find at URL:

`http://www.tuxedo.org/~esr/jargon/html/index.html`

AFAIK   As Far As I Know.

ANSI   Computers use several different methods for deciding how to put information on your screen and how your keyboard interacts with the screen. ANSI is one of these 'terminal emulation' methods. Although most popular on PC-based bulletin-board systems, it can also be found on some Net sites. To use it properly, you will first have to turn it on, or enable it, in your communications software.

ASCII   Has two meanings. ASCII (American Standard Code for Information Interchange) is a universal computer code for English letters and characters. Computers store all information as binary numbers. In ASCII, the letter "A" is stored as 01000001, whether the computer is made by IBM, Apple or Commodore. ASCII also refers to a method, or protocol, for copying files from one computer to another over a network, in which neither computer checks for any errors that might have been caused by static or other problems.

ARPANet   A predecessor of the Internet. Started in 1969 with funds from the Defence Department's Advanced Research Projects Agency.

Backbone   A high-speed network that connects several powerful computers. In the U.S., the backbone of the Internet is often considered the *NSFNet* (see entry), a government-funded link between a handful of supercomputer sites across the country.

Baud   The speed at which modems transfer data. One baud is roughly equal to one bit per second. It takes eight bits to make up one letter or character. Modems rarely transfer data at exactly the same speed as their listed baud rate because of static or computer problems. More expensive modems use systems, such as Microcom Network Protocol (MNP), which can correct for these errors or which 'compress' data to speed up transmission.

BITNet Another, academically-oriented, international computer network, which uses a different set of computer instructions to move data. It is easily accessible to Internet users through email, and provides a large number of conferences and databases. Its name comes from "Because It's Time."

Bookmark A gopher or Web file that lets you quickly connect to page-listed sites.

Bounce What your email does when it cannot get to its recipient – it bounces back to you – unless it goes off into the ether, never to be found again.

Command line On UNIX host systems, this is where you tell the machine what you want it to do, by entering commands.

Communications software A program that tells a modem how to work.

Daemon An otherwise harmless UNIX program that normally works out of sight of the user. On the Internet, you'll most likely encounter it only when your email is not delivered to your recipient – you'll get back your original message plus an ugly message from a 'mailer' daemon.

Distribution A way to limit where your Usenet postings go. Handy for such things as "for sale" messages or discussions of regional politics.

Domain The last part of an Internet address, such as "news.com."

Dot When you want to impress the net veterans you meet at parties, say "dot" instead of "full stop," for example: "My address is john at site dot domain dot com."

Dot file A file on a UNIX public-access system that alters the way you or your messages interact with that system. For example, your `.login` file contains various parameters for such things as the text editor you get when you send a message.   When you do an `ls` command, these files do not appear in the directory listing; type `ls -a` to list them.

Down When a public-access site runs into technical trouble, and you can no longer gain access to it, it's down.

Download	Copy a file from a host system to your computer. There are several different methods, or protocols, for downloading files, most of which periodically check the file as it is being copied to ensure no information is inadvertently destroyed or damaged during the process. Some, such as XMODEM, only let you download one file at a time. Others, such as batch-YMODEM and ZMODEM, let you type in the names of several files at once, which are then automatically downloaded.
EMACS	A standard UNIX text editor preferred by UNIX types that beginners tend to hate.
Email	Electronic mail – a way to send a private message to somebody else on the Net. Used as both noun and verb.
Emoticon	See *smiley*.
F2F	Face to Face. When you actually meet those people you've been corresponding with/flaming.
FAQ	Frequently Asked Questions. A compilation of answers to these. Many Usenet newsgroups have these files, which are posted once a month or so for beginners.
Finger	An Internet program that lets you get some bit of information about another user, provided they have first created a .plan file.
Flame	Online yelling and/or ranting directed at somebody else. Often results in *flame wars*, which occasionally turn into *holy wars* (see entry).
Followup	A Usenet posting that is a response to an earlier message.
Foo/foobar	A sort of online algebraic place holder, for example: "If you want to know when another site is run by a for-profit company, look for an address in the form of foo@foobar.com."
Fortune cookie	An inane/witty/profound comment that can be found around the net.
Freeware	Software that doesn't cost anything.
FTP	File Transfer Protocol. A system for transferring files across the Net.
Get a life	What to say to somebody who has, perhaps, been spending a wee bit too much time in front of a computer.
GIF	Graphic Interchange Format. A format developed in the mid-1980s by CompuServe for use in photo-quality graphics images. Now commonly used everywhere online.

GNU	Gnu's Not UNIX. A project of the Free Software Foundation to write a free version of the UNIX operating system.
Hacker	On the Net, unlike among the general public, this is not a bad person; it is simply somebody who enjoys stretching hardware and software to their limits, seeing just what they can get their computers to do. What many people call hackers, net.denizens refer to as "crackers".
Handshake	Two modems trying to connect first do this to agree on how to transfer data.
Hang	When a modem fails to hang up.
Hotlist	Same as a *bookmark* page (see entry).
Holy war	Arguments that involve certain basic tenets of faith, about which one cannot disagree without setting one of these off. For example: IBM PCs are inherently superior to Macintoshes.
Host system	A public-access site; provides Net access to people outside the research and government community.
HTML	HyperText Markup Language. The coding used on Web pages to define hyperlinks (see entry), graphics and the like.
HTTP	HyperText Transfer Protocol. The system used to connect World Wide Web resources to each other and to users.
Hyperlink	A way to connect two Internet resources via a simple word or phrase on which a user can click to start the connection.
IMHO	In My Humble Opinion.
Internet	A world-wide system for linking smaller computer networks together. Networks connected through the Internet use a particular set of communications standards to communicate, known as *TCP/IP* (see entry).
Killfile	A file that lets you filter Usenet postings to some extent, by excluding messages on certain topics or from certain people.
Log in/log on	Connect to a host system or public-access site.
Log out/log off	Disconnect from a host system.
Lurk	Read messages in a Usenet newsgroup without ever saying anything.

Mailing list   Essentially a conference in which messages are delivered right to your mailbox, instead of to a Usenet newsgroup. You get on these by sending a message to a specific email address, which is often that of a computer that automates the process.

MOTSS   Members of the Same Sex. Gays and Lesbians online. Originally an acronym used in the 1980 federal census.

Net.god   One who has been online since the beginning, who knows all and who has done it all.

Net.personality   Somebody sufficiently opinionated/flaky/with plenty of time on his hands to regularly post in dozens of different Usenet newsgroups, whose presence is known to thousands of people.

Net.police   Derogatory term for those who would impose their standards on other users of the Net. Often used in vigorous *flame wars* (in which it occasionally mutates to net.nazis).

Netiquette   A set of common-sense guidelines for not annoying others.

Network   A communications system that links two or more computers. It can be as simple as a cable strung between two computers a few feet apart or as complex as hundreds of thousands of computers around the world linked through fibre optic cables, phone lines and satellites.

Newbie   Somebody new to the Net. Sometimes used derogatorily by net.veterans who have forgotten that they, too, were once newbies who did not innately know the answer to everything. "Clueless newbie" is always derogatory.

Newsgroup   A Usenet conference.

NIC   Network Information Center. As close as an Internet-style network gets to a hub; it's usually where you'll find information about that particular network.

NSA line eater   The more aware/paranoid Net users believe that the National Security Agency has a super-powerful computer assigned to reading everything posted on the Net.   They will jokingly (?) refer to this line eater in their postings. Goes back to the early days of the Net when the bottom lines of messages would sometimes disappear for no apparent reason.

NSF   National Science Foundation. Funds the NSFNet, a high-speed network that once formed the backbone of the Internet in the U.S.

Offline   When your computer is not connected to a host system or the Net, you are offline.

Online	When your computer is connected to an online service, bulletin-board system or public-access site.
Ping	A program that can trace the route a message takes from your site to another site.
.plan file	A file that lists anything you want others on the Net to know about you. You place it in your home directory on your public-access site. Then, anybody who *fingers* (see entry) you, will get to see this file.
Post	To compose a message for a Usenet newsgroup and then send it out for others to see.
Postmaster	The person to contact at a particular site to ask for information about the site or complain about one of his/her user's behaviour.
Prompt	When the host system asks you to do something and waits for you to respond. For example, if you see "login:" it means type your username.
Protocol	The method used to transfer a file between a host system and your computer. There are several types, such as Kermit, YMODEM and ZMODEM.
README files	Files found on ftp sites that explain what is in a given ftp directory or which provide other useful information (such as how to use ftp).
Real Soon Now	A vague term used to describe when something will actually happen.
RFC	Request for Comments. A series of documents that describe various technical aspects of the Internet.
ROTFL/ROFL	Rolling on (the) Floor Laughing. How to respond to a particularly funny comment.
ROT13	A simple way to encode bad jokes, movie reviews that give away the ending, pornography, etc. Essentially, each letter in a message is replace by the letter 13 spaces away from it in the alphabet. There are online decoders to read these; nn and rn have them built in.
RTFM	Read the, uh, you know, Manual. Often used in flames against people who ask computer-related questions that could be easily answered with a few minutes with a manual. More politely: RTM.
Screen capture	A part of your communications software that opens a file on your computer and saves to it whatever scrolls past on the screen while connected to a host system.
Server	A computer that can distribute information or files automatically in response to specifically-worded email requests.

Shareware   Software that is freely available on the Net. If you like and use the software, you should send in the fee requested by the author, whose name and address will be found in a file distributed with the software.

.sig file   Sometimes, .signature file. A file that, when placed in your home directory on your public-access site, will automatically be appended to every Usenet posting you write.

.sig quote   A profound/witty/quizzical/whatever quote that you include in your .sig file.

Signal-to-noise The amount of useful information to be found in a Usenet newsgroup. Often used derogatorily, for example: "the signal-to-noise ratio in this newsgroup is pretty low."

SIMTEL20   The White Sands Missile Range used to maintain a giant collection of free and low-cost software of all kinds, which was 'mirrored' to numerous other ftp sites on the Net. In 1993, the Air Force decided it had better things to do than maintain a free software library and shut it down. But the collection lives on, now maintained by a Michigan company.

SLIP   Serial Line Internet Protocol. Used to turn home computers into Internet sites over a phone line.

Smiley   A way to describe emotion online. Look at this with your head tilted to the left :-). There are scores of these smileys, from grumpy to quizzical.

Snail mail   Mail that comes through a slot in your front door or a box mounted outside your house.

Spam   Message posted to numerous Usenet newsgroups to which it has absolutely no relevance (also a verb).

Sysadmin   The system administrator; the person who runs a host system or public-access site.

Sysop   A system operator. Somebody who runs a bulletin-board system.

TANSTAAFL There Ain't No Such Thing as a Free Lunch.

TCP/IP   Transmission Control Protocol/Internet Protocol. The particular system for transferring information over a computer network that is at the heart of the Internet.

Telnet   A program that lets you connect to other computers on the Internet.

Terminal	There are several methods for determining how your emulation keystrokes and screen interact with a public-access site's operating system. Most communications programs offer a choice of 'emulations' that let you mimic the keyboard that would normally be attached directly to the host-system computer.
Upload	Copy a file from your computer to a host system.
URL	Uniform Resource Locator. An address on the World Wide Web.
User name	On most host systems, the first time you connect you are asked to supply a one-word username. This can be any combination of letters and numbers.
UUCP	UNIX-to-UNIX CoPy. A method for transferring Usenet postings and email that requires far fewer net resources than TCP/IP, but which can result in considerably slower transfer times.
VT100	Another terminal-emulation system. Supported by many communications program, it is the most common one in use on the Net. VT102 is a newer version.
WWW	World Wide Web or the Web.

Useful Books and other Literature

This document is based on the "EFF's Guide to the Internet" by Adam Gaffin. To obtain a full version of the document (which was 160 pages of very well written text when I printed it) visit URL:

`http://www.eff.org/pub/Net_info/Guidebooks/`

Alternatively send an email message to **netguide@eff.org**
The subject and message are irrelevant, but a copy of the guide will be sent back to you (it takes up about 500K of space).

- "Internet Instant Reference" by Paul E. Hoffman. Published by SYBEX. Good handy manual.

- "The Whole Internet Guide" by Ed Krol. This is the Web Bible.

- "The Official Internet Yellow Pages" published by New Riders Publishing

There are many, many others. See what other people have got, borrow if possible before buying, but eventually you will probably find it worthwhile investing in at least one book about the Web.