

8

The Internet

8.1 Background

The Internet is the world's largest and best known computer network. It is often referred to simply as *the net*. It spans the globe and is readily accessible in all developed countries. Access in the developing world is, however, quite poor, particularly in large portions of the African continent.

The Internet evolved out of a number of separate networks, notably the *advanced research project agency network* (Arpanet), a research network that grew up in the 1970s. In the 1970s and 1980s, the principal users of the Internet were universities and the U.S. Department of Defense. In recent times, the Internet has been commercialized. Commercial organizations use it for promotion of their products, sending email, and delivering technical support.

A very important factor in driving the commercialization of the Internet was the development of a user-friendly interface (the web browser) in 1992. Not only are web browsers easy to use, they allow the same information to be viewed from a variety of computer types (i.e., PCs, Apple Macintoshes, and UNIX workstations).

In the period 1993 to 1997, the main commercial activity on the net has been the promotion of products and services and the use of email. There is great interest in electronic cash and secure submission of credit card numbers over the Internet. The take-up of this type of service has been low to date; however, this situation is expected to change radically over the next few years.

Estimates of the number of computers with Internet access varied from 35 to 100 million in early 1997. The exact number is considered impossible to determine because an organization only needs one connection to the Internet to

give access to all computers on its internal network. There are also differences surrounding what exactly constitutes Internet access. Lower figures will be obtained if you include only computers with World Wide Web access, while a much greater figure will be obtained if you attempt to estimate the number of people with Internet-style email addresses.

8.2 Connecting to the Internet

Today anybody with a computer, Internet software, a modem, and a phone line can connect to the Internet by paying a fee to a local Internet service provider (ISP). It helps if the ISP has a *point of presence* (POP) close by because this keeps the call charges down. Large organizations find it more economical to have a dedicated communication link to their ISP.

Internet software can be subdivided into communication software and application software. The communication software includes a TCP/IP protocol stack—the protocol stack used on the Internet (see Sections 7.4.3 and 12.3.2). The TCP/IP protocol stack is built into many desktop operating systems (such as UNIX or Windows 95).

There are many Internet applications; most common are the web browser, email software, telnet software, a newsreader, and phone software. Much of this software is either low cost or free.

8.2.1 The Single User Connection

A dialup connection over the PSTN or ISDN is an economic way for a single user or small business to connect to the Internet (Figure 8.1). The minimum requirements are a modem, a telephone line, and a terminal (usually a PC).

ISDN access is supported by many ISPs. Its advantages over the PSTN are twofold: first, you can connect in about 4 to 10 seconds, compared to 20 to 30 seconds using a modem; and second, your connection will have two to four times more bandwidth. In the case of ISDN, a terminal adapter is used in place of the modem.

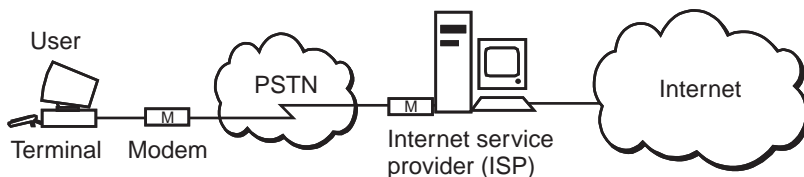


Figure 8.1 A single user dialup connection to the Internet.

The software at the user's end can be configured to dial up periodically and check for email. Alternatively the ISP may offer a service whereby it dials into your computer whenever you have mail waiting. You, of course, have to pay for these calls. It is important to ensure that some security mechanism is also offered to avoid someone else dialing into your computer.

Some ISPs offer a dedicated modem service. This is a modem on the ISP's end reserved for your connection, meaning that you will never get busy tone when you dial in.

People who travel a lot can save on international call charges by subscribing to an ISP that has POPs in each country they visit. In this way, they will be able to keep up to date with their email without having to make international phone calls. Large multinational ISPs and some smaller ISPs, which have formed alliances, can offer this type of service.

8.2.2 Corporate Connections to the Internet

Large businesses can connect their LANs to the Internet via dialup PSTN or ISDN, frame relay, or leased lines. They can thereby give everybody in the organization access to the Internet. They can restrict access to certain applications while allowing access to Internet email (browsing the Internet is often considered to be a waste of time for particular employees).

A router is required to connect to the ISP, and a firewall (see Section 8.6.1) is required to protect your corporate data against hackers. Many ISPs offer *managed* routers on the customer's premises and firewall services. This is particularly attractive for smaller organizations that do not have the technical expertise to purchase and manage their own router or firewall. A typical setup is shown in Figure 8.2.

The decision as to which access services to use (PSTN, ISDN, frame relay, leased line) will depend on economics, traffic levels, and whether or not you will

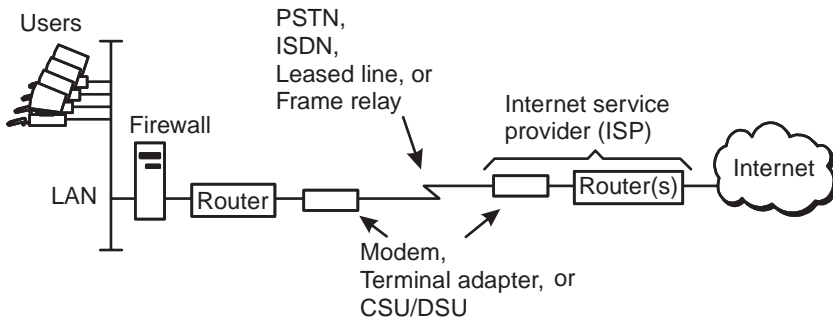


Figure 8.2 Connecting a LAN to the Internet.

be hosting information from your end of the connection. *Hosting* means setting up a server with online information and possibly the means to order products online. It is not necessary to host information from your end of the connection. It is often cheaper and better to pay your ISP to host your web site on one of its servers (see Section 8.4.1).

Permanent Connections to the Internet

If you do intend to host, then a fast, 24-hour-a-day connection is required. ISDN and PSTN connections tend to be expensive under these conditions, although there are exceptions—where ISDN or PSTN are offered at a flat rate for local calls.

Leased lines are probably the most common means of providing high-bandwidth permanent connections to an ISP. They are likely to be the only option available if you require an access speed greater than 1.5 Mbps.

Frame relay will often work out to be more cost effective, particularly if you are already using it for other purposes. Consider, for example, a business that has interconnected its LANs using frame relay. It has already invested in the frame-relay access at each site. To add a connection to its ISP, all that is required is an additional PVC from one of its sites through the frame-relay network to the ISP (see Figure 8.3). The charge for this additional PVC will be a fraction of the cost of a leased line. Care will need to be taken, however, that there is sufficient bandwidth in the frame-relay access link at the particular site connected to the ISP.

If you are not hosting information, then you will probably only require a connection during working hours and possibly only for a small number of hours per day. In addition, if you do not have too many employees browsing the web, then your bandwidth requirements will likely be modest. In these situations, ISDN and even PSTN connections should be given consideration. ISDN in particular is a very likely candidate for organizations with under 100

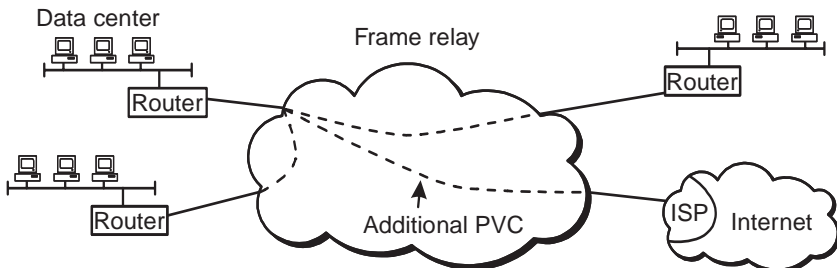


Figure 8.3 Adding an additional PVC to connect to an ISP over frame relay.

users—particularly if email is the primary application. A lot will depend on the amount of usage and on how ISDN tariffs compare to frame relay and leased lines.

8.3 Network Operation

8.3.1 A Network of Networks

Technically speaking, the Internet is an internetwork (or internet with a small “i”)—a network of networks, or a network of subnetworks. In most cases the subnetworks (subnets) are LANs or groups of interconnected LANs. The subnets are linked using routers. As the name suggests, routers are used to send the data packets in the right direction.

Traffic handling on the Internet is based on the TCP/IP protocol suite. These protocols are responsible for routing traffic, avoiding congestion, and correcting errors. TCP/IP is a packet-switching protocol ideally suited to bursty applications. It differs from typical X.25 implementations (see Section 7.5) in that it allows larger packets. This makes it better suited to the interconnection of LANs. Internetworks, routers, and TCP/IP are dealt with in more detail in Chapter 12.

8.3.2 ISPs and Backbones

The commercialization of the Internet during the 1990s has radically changed the structure and ownership of the Internet. Once a network with large government sponsorship, the Internet is now dominated by ISPs and multinational giants such as AT&T, MCI, and Sprint. These companies are ISPs in their own right but also act as service providers to the smaller ISPs.

Large ISPs have high-speed backbone networks linking their various POPs [1]. Backbone bandwidth can be anything from 1.5 Mbps up to 622 Mbps (MCI, early 1997). ISPs have to link their networks together. In the United States, the biggest interconnection points are called *network access points* (NAPs). There are similar interconnection sites around the globe (but they are not always called NAPs).

The interconnection between ISPs plays a very important role in producing an efficient and robust Internet. Unfortunately, however, commercial and competitive pressures can result in inefficient arrangements. For example, at the time of this writing, in Ireland not all service providers use the same interchange point. Thus, a proportion of traffic that should stay within the country actually gets to its destination via the United States or mainland Europe.

8.3.3 Bandwidth

Bandwidth on the Internet is a problem for users and ISPs alike. This is because Internet traffic is growing so fast that any increase in bandwidth quickly becomes saturated by an increase in traffic. MCI, for example, experienced traffic growth of 30% per month during 1996. In Europe, it is quite noticeable that the Internet slows down in the afternoon when users in the United States start to go online.

The important parameter is end-to-end bandwidth, not the bandwidth of individual links. Congestion from other Internet users is a critical factor here. End-to-end bandwidth is governed by the weakest link between you and the other end. If you dial up with a 14.4-Kbps modem, this will often be the limiting factor. If, however, a particular site is very busy, no increase in the bandwidth at your end is going to compensate for the congestion at the far end. The same is true when there is congestion on the backbone.

Response time from a particular web site can also be adversely affected by the web server at that site. If a lot of people are accessing a server at a given time, the server can become a greater bottleneck than the Internet connection.

The Internet does not give any quality of service guarantees, but individual ISPs may be able to offer certain guarantees, provided your traffic remains within their network. This principal can be further extended to ISPs which cooperate closely with one another. In the future, new technologies and protocols, such as *Internet protocol version 6* (IPv6) (see Section 12.3.3), will assist ISPs in offering *grade of service* (GOS) guarantees.

8.3.4 Other IP Networks

Some service providers offer connections to an IP network that is separate from the Internet. These networks offer higher end-to-end bandwidth and lower security risks. They use the same Internet software on the user's computers. Sprint Corporation, for example, offers such a service and through the use of gateways allows its customers the ability to connect to either this network or the Internet over the same access line.

8.4 Internet Applications

8.4.1 World Wide Web

The WWW or the *web* is a name used to describe the Internet as seen through web-browser software, such as Netscape Navigator or Microsoft Internet Explorer—the two most popular web browsers. Viewed this way, the Internet

looks like a web of interlinked documents or pages. There is more to the Internet than the web; however, more and more Internet services are becoming accessible via web browsers, and the distinction between the web and the Internet is diminishing.

A web page is a *hypertext* document containing built-in links to other documents. Alternatively, the links can be to downloadable files, sound or video clips, email addresses, or Internet phone connections. Web pages are stored on servers known as *web servers*. Forms can be included in a web page to allow the user to send information to the web server (Figure 8.4). This information could be an order for equipment, a response to a survey, a request for technical support, or the criteria for querying a database.

WWW Forms and Common Gateway Interface

When a user submits a form, information is transmitted back to the server and the server executes a *common gateway interface* (CGI) script using this information as input data. The server then delivers an appropriate response to the user. In the case of a database query, for example, the server will generate a web page on the fly containing the search results. The server may also perform some actions that are hidden to the user, such as storing a record in a database, transmitting information to another computer system to process an order, or emailing order information to a sales team.

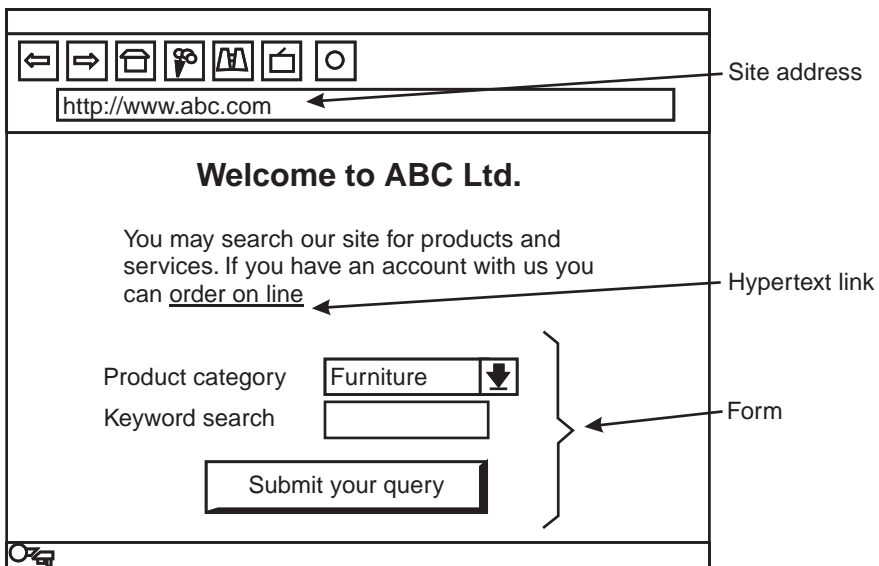


Figure 8.4 Web page viewed with an Internet browser showing a link and a form.

CGI scripts, while very useful, can potentially contain security holes if poorly written. For example, a poorly written script could be used by a hacker to obtain the password file from a web server. They should be only be written by somebody who fully understands the potential problems [2].

Secure Transactions on the Web

Most of the information transmitted between a web server and the browser is sent in plain text. There is, however, an optional encryption scheme, which is implemented in some servers and most browsers available today. This is desirable for the transmission of credit card numbers and other confidential information.

The user is normally given some indication that a secure transaction is taking place. In Netscape Navigator, for example, part of the screen changes color and a key symbol is displayed at the bottom of the screen. (The key symbol is always there but it is “broken” when not working in encrypted mode.) There are two levels of security available; 40-bit keys and 128-bit keys. This is further indicated by the key symbol having one or two prongs, respectively. The 40-bit key is considered insecure.

A new encryption key is automatically generated for each session with the secure server in a pseudo-random fashion. This key is transmitted to the user using public key encryption. This process is transparent to the user.

To protect against a hacker masquerading as a secure server and stealing your credit card number, the operator of a secure server can obtain an authentication certificate from a trusted third party (e.g., a *certification authority* such as VeriSign). This certificate is transmitted to the user at the start of a secure session. The browser can be configured to only accept certificates signed by specified certification authorities. Finally, the user can view the certificate with his or her browser, and if it contains the name of the company that he or she has connected to, he or she can be confident that the transaction will be safe [2].

Hosting Web Pages

It is very easy to advertise your company on the Internet. All that is needed is to pay an ISP to host your web pages on one of its servers. Creation of simple web pages is straightforward—most word processors have the capability to save files in web format. If however you need something more sophisticated, such as search facilities or online-ordering, you may need to get help. Many ISPs and other third parties offer such a service.

A collection of web pages owned by a particular company or organization is normally referred to as a *web site*. These pages are linked together using hyper-text links. The starting point for a particular site is called a *home page*.

As already mentioned, you have the option of hosting your web pages on your own server or on an ISP's server. The relative advantages of each option are given next [3].

Hosting on an ISP's server:

- Will usually cost less, particularly for small businesses;
- Will usually have more bandwidth available;
- Will not need a permanent connection to the Internet to support the site;
- Will usually be professionally maintained with regular backups to tape and administrative staff available outside normal working hours;
- Will make it easier to protect a corporate network against unauthorized access (because access to the corporate network is not necessary).

Hosting on your own server:

- Allows you greater ease in creating links between your web server and information on other computers on your internal network;
- Allows you more control in protecting your web pages against hackers by putting your web server behind your firewall. The type of attack referred to here is where the hacker replaces your web pages with uncomplimentary or misleading pages.

Only organizations with considerable IT experience should consider hosting web pages from their own site. If you choose to use an ISP's server, you will need to check out exactly what you are getting:

- Who is sharing the server with you (if your web site shares a server with one or more very popular sites, they may slow down the web server)?
- How often is the server backed up, and are backups stored off site?
- Can you publish pages on the server without assistance from the ISP?

Search Engines

There are quite a number of powerful search engines on the Internet. These work by automatically gathering information from the web, indexing it, and storing this index on a powerful computer. Thus, when you use these search engines, you search this index, not the web itself. Search results can be returned

within a few seconds. If you want to guarantee that your web site is included in these indexes you should register your home page with the search engine. This registration can normally be done for free over the web by filling out the appropriate form on the search engine's web site.

8.4.2 Email

Internet email is discussed in Section 10.3.1.

8.4.3 Telnet

Telnet is an application that allows you to connect to a computer (usually a UNIX host) in text mode. It can be used, for example, to access a text-based database that could not be accessed using a web browser. It is also a convenient way of allowing a network administrator to manually configure routers and servers. Because of telnet's text-only interface, most businesses prefer to make information available via a web browser.

8.4.4 News Groups

News groups are an important part of the Internet, though often neglected by newcomers. They are discussion areas where individuals pose questions or express opinions and others answer them. There are over 20,000 groups (in 1997), and a small proportion will be of interest to a business user. Communication and IT managers will find technical newsgroups covering most areas of interest. Depending on the nature of your business you may also find some news groups directly related to your core business.

News groups can be used to advertise your business, provided you go about it in a sensible way. It is generally acceptable, for example, to embed a link to your web page beside your signature in a message that contributes some useful information to the group. Blatant advertisements posted to a serious news group can result in the advertiser getting *flamed* (hate mail) or, worse still, mail bombed. Mail bombing means sending so many mail messages that the recipient is unable to use their mail account.

Many news groups are part of the *Usenet news* system. In Usenet, messages are replicated on a large number of *Usenet servers* around the world. Most ISPs, for example, maintain a Usenet server. This allows fast access to the messages but—due to space constraints—messages are only held on these servers for a couple of weeks. Archives exist for most news groups if you wish to look for older messages. Some Usenet groups are *moderated* (i.e., messages must pass

through a person known as a *moderator*). The moderator may add pertinent comments and acts as a mechanism for keeping messages to the point.

Private news groups are also possible. These can be either open to all comers but not part of Usenet or they can have access restricted to those who pay a membership fee or to employees of a particular company.

8.4.5 Voice and Video

Voice and video can be transmitted across the Internet using a *real-time protocol* (RTP). While this is a very cheap way of communicating over long distances, there are two significant problems. First, there can be a significant delay (sometimes two or three seconds) in the communication path; second, the call-setup mechanisms are poor (at least at the time of writing calls are often established by prior arrangement or by using conventional telephony to inform the other party that you wish to converse over the Internet).

Voice is transmitted at about 8 Kbps, and speech quality deteriorates if there is insufficient bandwidth. Quality video requires more end-to-end bandwidth than the Internet can provide. A compromise is to reduce the picture size and reduce the number of frames per second.

8.4.6 Fax

Fax messages can be sent over the Internet using special fax servers. These servers can be connected to the PSTN as well as to the Internet, thus allowing faxes to be sent to conventional fax machines. An organization with Internet connections in different countries can save on international call charges for fax by installing an Internet fax server in those countries. Figure 8.5 shows the routing of a fax message from a PC on a LAN in one country to a fax machine connected to the PSTN in another.

As noted elsewhere in this book, email can, in many cases, replace fax. Even if the document does not exist in electronic format, it can be scanned. The scanned image can then be sent by email as an attached file.

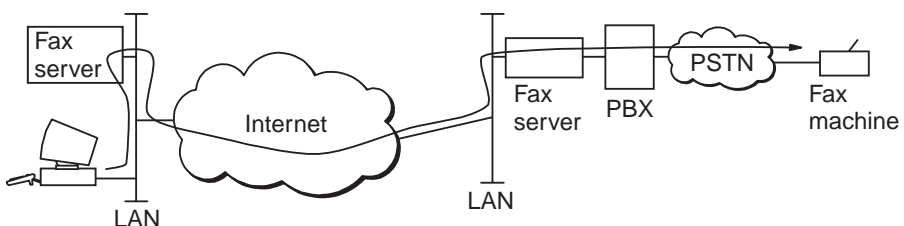


Figure 8.5 A fax message sent over the Internet.

8.4.7 Internet VPNs

The Internet can be used to interconnect LANs in different parts of the world. By using encryption, a *virtual private network* (VPN) can be created. The encryption is implemented in the firewalls, routers, or in a dedicated device. With some products, remote users, with the requisite software, can make a secure encrypted connection to the VPN. *Authentication* (I am who I say I am) is just as important as encryption when building Internet VPNs.

8.4.8 Intranets

An *intranet* is a corporate network based on Internet protocols and applications. Many companies have transformed their corporate networks into intranets because of the considerable price and usability advantages of Internet software.

Here are some examples of what can be done with a corporate intranet:

- Documentation can be published on the internal web.
- A search engine can be added to allow keyword searching of internal documents.
- Internal newsgroups can be established.
- Internet phone, fax, and video technology can be used to cut the communication bill.

Building an Intranet

An intranet consists first of an infrastructure of one or more interconnected LANs. Many organizations already have such an infrastructure in place and simply need to ensure that it is capable of handling the TCP/IP protocols. Chapter 12 describes the use of IP addressing and routing on such a network. The Internet can be used to provide the long-distance links for an intranet using encryption as described in Section 8.4.7.

Next you will need to establish at least one server to carry the web pages, search engines, name servers, and other server applications.

Finally, the desktop computers will need TCP/IP software with IP addresses and a web browser at a minimum. Depending on your applications, they may optionally require software for voice or video applications.

Extranets

The intranet concept can be further extended to the *extranet* concept, where trading partners agree on mutual access to servers across an IP network or the Internet to automate certain business transactions.

8.5 Addresses and Domain Names

Any medium to large business with a presence on the Internet will need a *domain name* such as `ibm.com` or `telecom.ie`. Such a name is somewhat analogous to a postal address. It allows data to find its way to your computers. Domain names need to be chosen carefully so that they are easily remembered (or better still can be guessed) and they create a favorable impression of your company. The latter point needs some explanation. A company called, say, ABC Ltd. in Ireland could opt for three different types of domain name:

- `abc.ie`—*ie* stands for Ireland;
- `abc.com`—*com* stands for commercial organization;
- `abc.tinet.ie`—*tinnet.ie* is the domain name of their ISP.

Your chosen domain name will become part of your web address (e.g., `www.abc.ie`) and also part of your email address (e.g., `sales@abc.ie`).

The `abc.tinet.ie` option is the least advantageous. Besides being difficult to remember, it suggests that you are doing things on the cheap by piggybacking your domain name onto your ISP's domain name. It has the further disadvantage that if you change ISP, you will also have to change domain name. This could cause a disruption to your electronic communications similar to the disruption a change in your main telephone number would cause to your voice communications.

The `abc.ie` and `abc.com` names are the superior choices. They must, however, be registered for a fee with a registration authority. The `.ie` ending is the Irish *top-level domain* (TLD) and must be registered with the Irish registration authority. Each nation has its own national TLD. The `.com` ending is a *generic top-level domain* (gTLD) and should be used if your company engages in business around the world. Currently, it must be registered with the InterNIC in the United States. There are several gTLDs, most of which are self explanatory (e.g., `.org` [non-profit organization], `.firm`, `.store`) [4].

The registration process involves sending your chosen domain name along with details of your organization to the registration authority [5]. Names are allocated on a first come—first served basis. There is normally a dispute procedure that can be invoked for example when a trademark is usurped.

As a general rule it is best to hand over the registration process to your ISP for a small fee in addition to the registration authority fee. They must be involved in part of the process, and they will be experienced in the pitfalls. You should, however, ensure that you are named as the administrative contact [6].

Finally, domain names are not essential to have a presence on the web, and very small companies might prefer not to invest in the (modest) registration fee. In this case the web address of ABC might be *www.tinet.ie/companies/abc/*. This, of course, has the same disadvantages mentioned above for the *abc.tinet.ie* domain name.

At least one IP address will be required for a connection to the Internet. Your ISP should give you these addresses.

8.6 Security

There are a number of serious risks associated with connecting a corporate network to the Internet. The most important risk is from computer hackers who may have access to some or all of your computer systems via the Internet. The simplest protection against hacking is not to connect to the Internet. If you must connect, then you should implement a well-thought-out security policy including the use of a firewall at the point of access to the Internet.

8.6.1 Firewalls

A firewall is a computer or router connection between your network and the Internet. The firewall's function is to block unwanted traffic while allowing legitimate traffic through. It can block traffic on the basis of information in the packet headers, such as source or destination address, or the protocol in use (e.g., email, telnet, or web). So, for example, a firewall can be configured to block all access from outside to a particular host on the inside based on the IP address of that host.

The cheapest form of firewall is one implemented on the router that connects your network to the Internet. This type of firewall will perform the basic filtering required to keep most intruders out. It requires quite a bit of router configuration know-how, and expertise often has to be bought in.

A more comprehensive firewall is implemented in software on a dedicated computer attached to the same LAN as the router connected to the Internet. This type of firewall has the added advantage that it can detect suspicious activity (i.e., activity associated with known hacking techniques), raise alarms, and keep a log of particular events [7].

Care should be taken to install this second type of firewall on a computer that will be able to handle the traffic load. Remember that all traffic to and from the Internet must pass through this computer.

The firewall can also be used to enforce access rights for employees. Filters can be established in the firewall to specify what each computer on the

internal network can and cannot access. This technique can be used, for example, to allow specific employees access to the web while restricting others to email only.

8.6.2 Access Control

No firewall system is perfect, and there is always a risk that an experienced hacker will penetrate your firewall. Further protection of your data can be achieved by logical access control to your computers. Passwords are a minimum, and their limitations should be understood. They can be hacked by eavesdropping, by obtaining password files, and through social engineering such as masquerading as a technical support person on the telephone and asking someone for their password. Other forms of social engineering include hackers taking temporary work with the company in question or with cleaning contractors.

Educating users can go some distance, but more sophisticated systems should be considered for highly sensitive data. One such technique involves the use of a device that looks like an electronic calculator, which generates an access code that changes every 60 seconds. The code is only valid during those 60 seconds. The user reads the code from an LCD display. The main point is that you have to have the physical device before you gain access [8]. If you lose it, you inform the network administrator, who cancels access from this device. Social engineering is thus rendered much more difficult. Authentication schemes using digital certificates (see Sections 8.4.1 and 8.6.3) can also be used.

8.6.3 Encryption

Any information sent over the Internet can potentially be intercepted along the way. This interception can be passive (read only), or if someone is really out to do damage they could possibly intercept and modify your information. Encryption is the best way to protect against this form of attack. There are three common encryption scenarios. First, email messages are encrypted in the sender's machine and decrypted at the recipient's machine. Second, transactions with secure servers are encrypted in the end machines. And third, when virtual private networks are created using the Internet, encryption is usually performed in the routers or in a dedicated encryption box connected alongside the routers.

There are two types of encryption in use today: symmetric key and public key. They have different characteristics and are often combined to get the best of both worlds.

Symmetric key is the simplest. A secret number called a secret key is used to encrypt the information using a scrambling procedure known as an

encryption *algorithm*. The information can only be decrypted using the same key and algorithm. The algorithm is usually well known, and it is the secret key that provides the security. Both parties must know the key.

Public key encryption is more sophisticated. Each party has two keys, a public key and a private key. The public key can and should be freely advertised. The private key is never divulged. If Bob wants to send a message to Alice, he will encrypt it with Alice's public key. The only way to decrypt this message will be with Alice's private key.

It should be clear from the above that public key encryption simplifies the transfer of keys. It does have the significant drawback, however, that it is about 100 times slower to perform in software than symmetric key encryption. Most encryption schemes therefore combine the two methods as follows. Bob's software generates a random symmetric key and encrypts his message with it. It then encrypts this symmetric key using Alice's public key and sends it to Alice along with the encrypted message. Alice now has all she needs to decrypt the message. As added security, the symmetric key will be different for each message.

The strength of an encryption scheme depends on the length of key used. For symmetric key encryption, keys of 80 bits or more are recommended for valuable information. For public key encryption, the same protection requires keys of 768 bits or more [9]. As computer power increases, longer and longer keys will be needed.

Public key encryption can also be used to generate *digital signatures*. A digital signature is a data stamp appended to a message or file which quasi-uniquely links that file or document to the *private* key of the signer. Put another way, using a very powerful computer, it would take years to forge a digital signature. The digital signature can be verified by anybody using the *public* key of the signer. This verification not only verifies the origin of the message, it also proves the *integrity* of the message. Digital signatures are not yet legally binding in the United States but can still provide a high degree of confidence. Bilateral agreements can be used to provide legality in the absence of a generic law.

One problem with public key encryption is verification of the public keys. If Bob sends an encrypted and signed message to Alice looking for her credit card number, Alice must be sure that the public key that she assumes belongs to Bob actually does. One method of providing this guarantee is with a *digital certificate* provided by a trusted certification authority (as mentioned in Section 8.4.1). A digital certificate contains the name of the holder, his or her public key, an expiration date, and a digital signature signed by the certification authority. The public key of the certification authority can be trusted because it will be widely publicized by them. Hierarchies of trust can be established

whereby, for example, a company certifies its employees and the company is in turn certified by the certification authority [10].

References

- [1] Gareiss, Robin, "The Online Corporation: Choosing the Right Internet Service Provider," *Data Communications International*, November 21, 1995.
- [2] Stein, Lincoln D., "The World Wide Web Security FAQ" Version 1.3.7, <http://www-genome.wi.mit.edu/WWW/faqs/www-security-faq.html>, May 7, 1997.
- [3] Gareiss, Robin, "Web Hosting Services: No Mess, Less Stress," *Data Communications International*, November 21, 1996.
- [4] International Ad Hoc Committee, "Recommendations for Administration and Management of gTLDs," <http://www.iahc.org/draft-iahc-recommend-00.html>, February 4, 1997.
- [5] InterNIC, "New Domain Name Registration: Frequently Asked Questions," <http://rs.internic.net/domain-info/registration-FAQ.html>, May 1996.
- [6] Wilson, Ralph F., "Does Your Business Need a Custom Domain Name?" *Web Marketing Today*, <http://www.wilsonweb.com/articles/domain.htm>.
- [7] Newman, David, "Can Firewalls Take the Heat?" *Data Communications International*, November 21, 1995.
- [8] Johnson, Johana Till, "Enterprise Security: Better Safe than Sorry," *Data Communications International*, March 1995.
- [9] Schnier, B., *Applied Cryptography*, 2nd ed., New York: John Wiley and Sons, 1996, quoted in Chuck Shih, Mats Jansson, Rik Drummond, and Lincoln Yarbrough, "EDIINT Functional Specification—Requirements for Inter-operable Internet EDI," EDIINT Working Group, <http://www.imc.org/ietf-ediint>, March 1996.
- [10] VeriSign, Inc., "Frequently Asked Questions about Digital IDs," <http://www.verisign.com/repository/digidfaq.html>, 1997

Further Reading

Cheswick, William R., and Steven M. Bellovin, *Firewalls and Internet Security—Repelling the Wily Hacker*, Reading, MA: Addison-Wesley, 1994.

Chapman, Brent D., and Elizabeth D. Zwicky, *Building Internet Firewalls*, Sebastopol, CA, O'Reilly and Associates, 1995.