

## Connecting to the Internet

One of the reasons for implementing TCP/IP on an internal network is to be able to connect it to the Internet. This section has a general look at various methods of obtaining an Internet connection, and the important security issues that arise. Improving the performance for users will also be examined. With the increasing use of intranets and extranets, many of the security and performance issues covered will be important, even if you are not planning to connect to the Internet itself.

### Choosing a Connection

Getting an Internet connection is reasonably straightforward, as long as the correct sort of connection has been chosen at the outset, and the Internet Service Provider (ISP) is sufficiently competent.

As well as a physical and logical connection, you will need to establish whether you need to register a domain name (and if so, which one) and whether you need to have a set of IP addresses issued to you.

### Connection Types

These may be divided in to two broad categories: permanent and dial-up. Permanent connections are usually implemented using a leased line. Dial-up connections made be implemented using either a Plain Old Telephone Service (POTS) or Integrated Services Digital Network (ISDN) line. ISDN connections may be further sub-divided in to “ordinary” dial-up and routed connections. Which type of connection to choose will depend primarily upon requirements and budget.

The only hardware required to implement a POTS dial-up connection is a modem (short for modulator/demodulator – it converts the digital computer signals to the audio tones that the phone system can handle) to sit between the computer and telephone line. Given a suitable combination of computer, modem and ISP, connection speeds of up to 56Kbps should be manageable; although the speeds achieved in practice are often a little less. Although primarily suited to the individual, it is possible to use some extra hardware and software to implement a proxy (covered later) to allow multiple users to access the account over a local area network. Because the only extra item required to implement this is a modem, this solution is generally quite cheap.

Implementing a dial-up connection using ISDN requires the addition of an ISDN line and a terminal adapter (TA). In this situation, the TA appears, on brief inspection, to be acting just like a faster modem. This sort of connection usually uses the Basic Rate Interface (BRI) version of ISDN which provides 2 64Kbps ‘B’ (bearer) channels plus one 16Kbps ‘D’ (delta) channel, which is used by the phone system for its own use (call set-up and tear-down, etc.). Some TAs may only support one of the channels, and some ISPs may not allow the use of both channels at once without paying extra. To use both B channels requires the availability of channel aggregation so that both ends know that the two channels are related to each other. Support for multiple users may be added in a similar manner as for a POTS dial-up.

Routed ISDN requires an ISDN line and a router suitable for connection to the LAN and compatible with the ISP's systems (which isn't normally a problem these days). There are a number of subtle options available depending on exact requirements. For example, the connection could be used for incoming traffic as well as outgoing (so incoming mail would be delivered immediately) or a company web site could be hosted internally instead of on the ISP's systems.

The big problem with this is that call costs are totally outside the control of the company, since the line is brought up every time there is incoming traffic. Routed connections such as this are generally implemented using Calling Line Identification (CLI). The router at the ISP end, on receiving traffic for the client, makes a call to the remote router, but does not wait for it to connect. The remote (company) router spots the incoming call and who it came from, makes a call back to the ISP (which uses CLI and other authentication to know where the call is coming from) which then delivers the data. If the web site becomes popular, the call charges can mount up very quickly!

Another possibility, more usually implemented as standard with routed ISDN connections, is bandwidth on demand. Essentially, when the level of traffic on the first channel starts to reach a limit (either of the channel or a defined threshold), the other channel is brought up as well. This of course incurs two call charges.

Give consideration to the type of security to be implemented and the possibility of using a proxy to try to reduce the amount of usage the line gets. Also, it may be necessary to use something (hardware or software) which provides network address translation (NAT), otherwise you will have to renumber the network using legal IP addresses.

When installing any form of WAN link (as an Internet connection, in effect, is), especially when using ISDN, it is *vitaly important* to ensure that everything is set up properly. There are many recorded cases of companies suddenly receiving quarterly bills of thousands of pounds because a poorly configured system was trying to broadcast a routing update or perform a DNS lookup that it shouldn't have been. A few hundred bytes of data would incur a minimum call charge; every 30 seconds in the case of RIP, which could easily end up keeping a line open the whole time.

Although leased lines are generally used for the tail end of a permanent Internet connection, there are a number of other technologies that may be used between the client and the ISP. However, the in-house implementation will generally be the same. The leased line provider will terminate their line with a small box (sometimes called a Network Termination Unit or NTU but sometimes called other things, the function is the same though), which will have a serial connector on it (and maybe some others but they're not usually relevant). This serial connector will need to be connected to a compatible connector on the back of your router. Make sure, if you are buying the router yourself, that you order the right cable to go with it – they're often not supplied.

Some routers are configured using a dumb terminal (e.g. a PC running some terminal software) plugged into a serial port. Others are configured using GUI software running on a machine on the local network (hopefully on the same platform as you have available). Others support both. Configure the routers' interfaces with the correct IP addresses and subnet masks (the ISP may assign an IP address to your router dynamically), add a route from the local network interface to the WAN port (if necessary, you may need the IP address of the router at the ISP end, or it might work it all out for itself), and tell all your local machines to use the router as their default gateway. That's assuming that your local computers are all using legal, public, IP addresses; if not, you will need something to provide network address translation (NAT) (or renumber your network!). You *will* also want to implement some form, or (better) forms of security.

Don't assume that because your connection doesn't work that you have done something wrong. Even if your ISP tells you that everything is set up correctly at their end, have them check again; as often they find it easier to blame problems on a user error.

### Choosing an ISP

There are now a number of magazines that provide details of the levels of service provided by ISPs, usually updated on a monthly basis. Apply the same criteria to any potential ISP as you would to any other company providing you with an important service. Make sure that they can answer any concerns you may have. Check how clued-up the people you talk to are. The people providing support once you're a customer will probably be different, but if an ISP has good customer-facing staff, the back-office technical staff have a better probability of being competent too.

If possible, ask around any other people you know with a similar type of connection to the one you are planning, and see how they feel about their ISP. In addition, ask potential ISPs for reference sites.

### Firewalls


Firstly, whatever else you may be told, a firewall is not the sole solution as far as securing your network is concerned. It is vitally important that the correct design be used, and the correct product chosen to implement whatever security policy has been decided upon. However, it is also vital that the internal users of the network are aware of what access they have, what they can't do, and, most importantly, why. Internet security should be implemented as part of a much wider process of educating users as to their responsibilities for maintaining the security and integrity of the information and systems they have access to.

Without a firewall in place, all the hosts on a local network are open to attack from outside. The idea of a firewall is to allow users access to resources both on their local network and on parts of other networks such as the Internet, whilst at the same time preventing anyone outside the local network from gaining access to it without authorisation. This is done by controlling all the traffic between the local and external networks, and blocking (and logging if required) anything which is not allowed by the security policy for the network or organisation.

Traffic between the local and external network may be carried out at any of the network layers but, typically, devices operate at the Application or Network Layers (in terms of the ISO model). In the context of TCP/IP, firewalls at the application layer are generally termed **Application Gateways** or **Proxies**. Firewalls at the network layer are **Filtering Routers**, **Screening Routers** or **Packet Filters**.

### Proxies

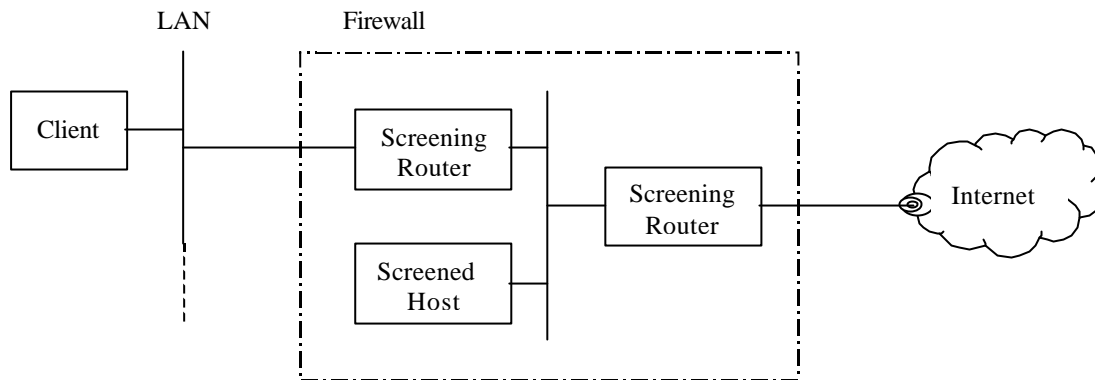
A proxy prevents direct connections between a piece of client software on the local network (such as a Web browser) and a server on the external network. Instead, two separate connections are used, one from the client to the proxy and one from the proxy to the target server.



Instead of connecting directly to the required server, the client program makes a request to the proxy server. As long as the particular client is allowed to carry out the action requested (e.g. visit a given Web site) the proxy client will open a connection to the target server. Traffic from the server is then relayed indirectly back to the originating client.

## Filtering Routers

Various architectures of increasing complexity are possible, using routers that filter packets. The diagram below shows a more complex arrangement, complete with a host that may be running proxy software.



The firewall area in this case is sometimes known as a DMZ or Demilitarised Zone. The screened host may also be called a Bastion Host.

A proxy is not generally used or useful by itself, since only those applications which are understood by the proxy software may be used, requiring upgrades if additional functions or new applications become available. A packet filter is of little use by itself, since although it can block packets to specified hosts and services, they are not capable of assessing the content of the packet to see if it is likely to do any harm to the destination host or otherwise cause a breach of the network security policy.

The level of investment in the design and implementation of the security for a network should be at least proportional to the value of the systems (hardware, software **and** data) being protected.

### Other Performance Issues

One of the biggest performance limitations with the Internet is all the other people using it. This is particularly true of interactive services such as the World Wide Web. Careful choice of ISP may help, depending on the sites that are being accessed and the ISPs routing to them. Best results are likely to be obtained in the morning (UK time) when US students and business people tend to be asleep, thus reducing overall loading on the Internet enormously.

Use Performance Monitor, described in the *TCP/IP Utilities* Section, to keep track of the resources being used on the system. Performance tuning a *Windows NT* system is outside the scope of this course. Network Monitor is also useful, but is obviously biased towards measuring network parameters, whereas in most instances a more holistic approach is required to get the best out of a system. Remember to take baseline measurements so that you have something to compare your network to when new devices or services are added.

When operating *Windows NT* in a multi-protocol environment, you may find that other protocols, IPX and NetBEUI in particular, may be using more than their fair share of processor time, causing the TCP/IP stack to drop packets. If possible, remove any unnecessary protocols. Otherwise, go to the Control Panel → Network applet, select NetBIOS Interfaces on the Bindings tab, and move TCP/IP to the top of the list, if it isn't there already.

Many problems where the solution might appear to be “upgrade the processor” or “upgrade the disk sub-system” actually turn out to be “upgrade the memory” in the first instance.

## Recommended Reading

Practical UNIX & Internet Security, 2nd Edition  
By Simson Garfinkel & Gene Spafford  
ISBN: 1-56592-148-8

Building Internet Firewalls  
By D. Brent Chapman & Elizabeth D. Zwicky  
ISBN: 1-56592-124-0

Web Security & Commerce  
By Simson Garfinkel with Gene Spafford  
ISBN: 1-56592-269-7

Windows NT TCP/IP Network Administration  
By Craig Hunt & Robert Bruce Thompson  
ISBN: 1-56592-377-4

Getting Connected  
The Internet at 56K and Up  
By Kevin Dowd  
ISBN: 1-56592-154-2 (pig cover) 1-56592-203-4 (lion cover)

All published by O'Reilly.

Firewalls and Internet Security: Repelling the Wily Hacker  
By William Cheswick and Steven Bellovin  
Published by Addison-Wesley

# Summary

- Main connection types: Dial-up (POTS or ISDN) and Leased Line.
- Firewalls used to protect internal network from outside.
- Proxies used for protection and to improve performance.
- Firewalls and proxies only assist in securing a system, they are not a complete solution.