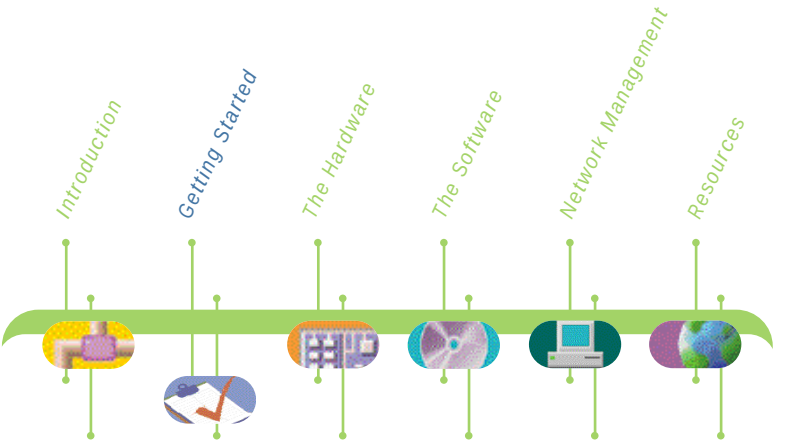

Getting Started

Networking topologies and protocols

WHEN WE TALK ABOUT TOPOLOGY in the networking context, we're referring to the layout of computers, cables and the other gadgets in the network. Topologies can be 'physical', a reference to how the network is connected, or 'logical', in which case we're looking at how the data moves around the network. A network can have one kind of physical topology and a different kind of logical topology.

There are three basic possible network topologies — bus, ring and star. In a bus topology, all of the devices are connected to one cable, which is called the bus or backbone. The ends of the bus don't connect.

If the ends *do* connect, you have a ring network, while in a star network, all of the cables connect to one central 'hub' device. This makes the network easier to manage — since one end of every connection is in one place — but requires more cabling, for the same reason.

Networks which use a physical star topology can be either ring or bus networks — if the ring or a bus is 'folded' into a star shape — the central hub being really just a 'patch bay', made of nothing but sockets and wires. In this case, the hub's only function is to simplify the physical cable layout by giving everyone one place to connect one end of their cable. There are also star-wired ring networks in which the star cables are simply long taps to a central ring in the hub; in this case, you can interrupt any of the star cables and as long as the hub is intact the network will still work.

10BaseT and 100BaseT Ethernet are true star networks — they use electronic repeater hubs which amplify the signals they receive before squirting them out to all of the other ports. They couldn't, even in theory, be 'unfolded' into a ring or bus, at least not without making the hub as big as the whole network. The repeater hub design allows longer cable runs, but makes the hubs more expensive. Basic 10BaseT hubs, though, are very affordable these days. A 'hub' can also be a computer with multiple network cards, though for Ethernet it's generally a stand-alone device.

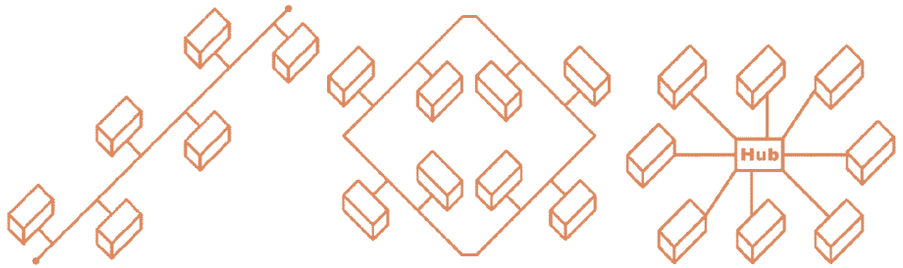

Bus topology | Ring topology | Star topology

Any star topology network, whether logically a star or not, is dependent on its hub to work. If the hub malfunctions, the network goes down. For a ring or bus network, though, the whole network goes down if any cable or connector is faulty. Star networks are thus much more reliable and make troubleshooting much easier.

By using different networking appliances, you can mix networks with different topologies. A business might, for instance, have a high-speed bus network connected to several slower star networks around a building.

Ethernet layout in detail

The basic layout for both 10Base2 and 10BaseT is simple. The older 10Base2 uses the bus topology. A 10Base2 segment (a segment in this case is a network with no bridges or switches or other devices in it) contains two or more computers, each with a network card (NIC) fitted with a T-piece which accepts two network cables, or a network cable and a 50-ohm terminating resistor.

Each end of the 10Base2 network must have a terminator and a cable connected, and everything else has two cables connected. The terminators prevent data on the network from 'bouncing' off the ends of the cable. The T-piece must go right on the network card — you can't use extension cords between the card and the T-piece.

You can take a computer out of the network by disconnecting its T-piece from the network card, leaving the cables connected to the T-piece's two arms. Disconnect in any other way, or remove either terminator, and the network stops working until you plug it back together. Every machine has to be at least 50cm of cable away from every other machine. Viewed as a whole, the 10Base2 segment can be thought of as one long cable with as many little taps coming off it as there are computers.

A 10BaseT/100BaseT hub.

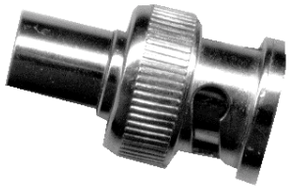

A 10Base2 terminator.

10BaseT, on the other hand, uses a star topology with a hub. You can actually connect two 10BaseT-equipped computers with a simple crossover cable (see page 38), but for any more machines you need a hub. Every machine on the network must have its own single lead to a port on the hub,

which must therefore have enough ports to support the number of machines you want to network. Multiple hubs can be connected to allow larger networks, often using the same crossover cable that can connect two machines.

Popular protocols

A 'protocol' is any agreed-upon format for data transmission between devices. If you think this definition is pretty broad, it is: there are hundreds of protocols. Networking protocols can be thought of as the language that the network devices speak to each other, but again, the term is rather broad. Each version of Ethernet is itself a protocol, and there are further protocols for the information sent using Ethernet protocols. Ethernet puts data onto the cable and takes it off again; 'higher level' protocols determine what that data actually looks like.

To connect computers directly to the one network, without using translation devices, they all have to be using the same 'physical' protocol. This means they all need to be using the same cables, connectors and network cards — say, 10BaseT Ethernet. Within that restriction, though, they can use different higher level protocols without interfering with each other. But only computers that share at least one higher level protocol will be able to communicate with each other.

Most of the time, when people talk about 'network protocols' they're referring to the higher level protocols. We'll get to those in a moment; first, some other physical protocols.

Token ring

Token ring runs second — these days, a quite distant second — to Ethernet in popularity. As the name suggests, it uses a ring topology. Like Ethernet, it's a baseband network on which only one computer can talk at a time. Unlike Ethernet, it hands out talking privileges using a special data packet, the 'token', which is passed around the network from computer to computer.

No computer can send or receive data when it doesn't have the token; if a computer wants to send data to another computer, it has to wait until it has the token. When a computer gets the token, it checks to see if the token contains any data addressed to it, and reads that data if it does. If it has anything to send, it puts the data it wants to send in the token along with the address of the computer it's going to and then relinquishes control.

The token ring system makes collisions impossible, so computers share the available network bandwidth more effectively. This means token ring deals better with high traffic. On the other hand, when traffic is low, token ring is slower than Ethernet for the same nominal bit rate. The whole network can also hang if a computer refuses to release the token. When a computer on an Ethernet network goes berserk and starts 'jabbering' (sending lots and lots of meaningless data), the network will slow down considerably but the other machines *will* at least be able to get a word in edgewise.

Wireless networking

Wireless — that is to say, radio-based — networking protocols have traditionally required expensive hardware and have been of interest only to very specialised markets. This is set to change with products such as Diamond's Homefree and Proxim's Symphony, which are affordable for small office and home users. The most serious downside of these networking systems, besides their price, is that they're slow.

Plain Ethernet runs at 10Mbps, at which speed moving 1M of data from one computer to another on an uncongested network takes a little more than a second. Cheap wireless systems manage 1 to 1.6Mbps, which is OK for many applications but painful if a lot of data has to be moved.

The cheaper wireless networking systems, like Homefree, also suffer from a seriously limited range. Homefree seems to be fine in any situation where you could easily run a 10Base2 or 10BaseT cable; solid walls or floors slow it down or stymie it completely. Symphony uses larger antennas and works more efficiently.

USB networking

The Universal Serial Bus (USB) port on modern computers is a handy way to connect lots of low to medium-speed devices, and with quite inexpensive special hardware it can also be used for networking. There's no standard support for this in any operating systems yet, but there are third-party products such as EZ-Link from Anchor Chips that make it possible. USB networking is a cheap and simple alternative to pricey PCMCIA network cards for laptops, but since the two-PC kit is more expensive than two regular network cards with cables, and is at best half the speed of 10Mbit Ethernet, it isn't very exciting for general networking purposes.

Other networking systems

Various other networking systems vie for the home office market dollar — including those that use house electrical or phone wiring, for instance. These systems are unlikely to be available in Australia because of differences in local wiring and approval systems, but they're generally slow and expensive. Like wireless systems, their advantage is that they don't require new cables to be run.

Popular higher level protocols

When you're setting up a network, you can assign multiple higher level protocols. You don't have to, though; all you need is one protocol that everyone shares, and all of your basic networking functions will work. Some programs, notably games, don't work on some of the popular protocols. But they are the exception, not the rule.

Possibly the most popular office LAN protocol in the Windows world is NetBEUI (NetBIOS Enhanced User Interface, pronounced 'net-booey'), if only because it's often installed by default and doesn't need any extra configuring. NetBEUI is an evolved version of IBM's old NetBIOS protocol, still used on Novell NetWare networks. It often doesn't bog down lower powered computers, and works well on networks with

fewer than 50 machines, which covers most situations.

IPX/SPX is another popular protocol for business networks, and for many network games released before 1998. Short for Internetwork Packet Exchange and Sequenced Packet Exchange, IPX and SPX, are old Novell standards. IPX is 'connectionless'; it simply puts data on the network with an address and hopes it reaches its destination. SPX is another protocol which handles connection-oriented features such as error recovery.

AppleTalk is the protocol used on most Macintosh networks because it's built into every Mac computer and laser printer. PCs can be connected to AppleTalk networks, too (see the Network Management chapter for more information).

TCP/IP, which stands for Transmission Control Protocol/Internet Protocol, is the big daddy Internet protocol set. Like IPX/SPX, it's a portmanteau protocol; in this case, IP is the connectionless component. TCP/IP is the most powerful and flexible high-level protocol, but it is more difficult to set up and is overkill for most networks. Nonetheless, a lot of computers have it installed, because you need TCP/IP to access the Internet.

TCP/IP requires a little extra explanation. Every computer attached to a TCP/IP network, whether it is a LAN or the Internet, has an IP address — four numbers, each of which can be from 1 to 254. This gives a total of more than four billion possible addresses.

Every TCP/IP address also has a second, similar number, called the 'subnet mask'. Computers on different subnets can't see each other.

There are only three ways a computer can get an IP address.

- It can have one assigned to it by another computer. This happens when you connect to the Internet using a modem, which acts as a separate network adapter to your network card, and can be assigned a different IP address.
- It can make one up randomly. This option is rarely employed, but Windows 98 can do it, as will be explained in the Software chapter.
- It can have one assigned to it manually. This is the standard procedure for setting up a small LAN using TCP/IP.

Choosing the network for you

FOR MOST NETWORKING PURPOSES THESE days, Ethernet is the king of the hill. All of the cheap PC network cards and hubs and other gadgets you see in a computer store are Ethernet devices, and there are few networking needs it can't meet.

10Base2 and 10BaseT are the lowest-cost Ethernet options and each has its pros and cons. Both 10Base2 and 10BaseT offer really, really cheap network cards — $40 gets you a plain Plug-and-Play NIC with both 10Base2 and 10BaseT connectors. For 10Base2, that's pretty much where the spending stops; the network cards come with T-pieces, so all you need are enough cables and a couple of terminators, and you're in business.

For 10BaseT, you have to buy a hub as well, and this will set you back about $100 for a five-port hub and around twice as much for 16 ports. 10BaseT hubs are available in various sizes and can be 'cascaded' to add more ports to your network. 10BaseT cables are very cheap, just like 10Base2, and no T-pieces or terminators are required.

So why bother with 10BaseT? Well, in the 10Base2 configuration, one dud cable, dodgy T-piece, duff terminator or poor connection renders the whole network stone dead until the defective component or connection is fixed — or, worse yet, just interrupts the network every now and then. Intermittent problems are the most annoying.

Finding the defective component in 10Base2 is a process of elimination. You just start somewhere, anywhere, and then 'divide and conquer'. You cut the network in two and re-terminate the two halves, then see which half still has the problem and divide it again, and so on, until you locate the source of the failure.

With 10BaseT, on the other hand, one dud cable or network adapter will only remove one machine from the segment, which obviously indicates the location of the problem. A dead hub will kill the network for every machine that is directly connected to that hub, but hubs are much more reliable and are less prone to accidental damage than 10Base2 cables, T-pieces or terminators.

So for practically any small office application, the network to go for is 10BaseT. For home use, 10Base2 is a little cheaper and works alright when you have relatively few machines.

If your needs are more complex, you should start thinking about 100BaseT and/or extra hardware to maximise the available network bandwidth and give a 'thicker pipe' to those devices that need it. 100BaseT network cards aren't much more expensive than 10BaseT, but 100BaseT hubs cost significantly more.

There are dual-mode hubs that work with both standards, but a plain hub cannot run one of its ports faster than another. More complex hardware allows you to mix the networks more efficiently (see the Hardware chapter's section on switches and routers).

Client/server or peer-to-peer?

If you're only using one kind of computer, a peer-to-peer network will work well. Since most offices these days are all-Windows shops, this is not a problem. If you have multiple operating systems — mixed Windows, Macintosh and Unix boxes, say — then a server-based system can reduce the insanity. The big fat server can be configured to speak everyone's language and dole out files, so the desktop machines don't have to all be polylingual too.

If you have lots of computers hammering away at the same data, that data should be on a server. The number of computers there should be in 'lots' is open to debate, but if many people are relying on the same files, obviously those files should be somewhere secure and, preferably, fast.

If you need data security, a server-based system is the only way to go. A peer-to-peer Windows 95/98 network is about as secure as a tent. For many people this doesn't matter much, but you don't have to be a big business to have data you want kept safe and secret.

On the other hand, if you're strapped for cash, then the price of peer-to-peer is right. It's just a bunch of regular computers, after all. Peer-to-peer systems also don't, generally, need a full-time network manager — not because it's *easier* to do advanced things with a peer-to-peer network, but simply because you *can't*.

Plan it out

Most home and office networks are pretty simple enterprises. If you simply want to connect a gaggle of Windows PCs, setting up your network will be easy. A network card for every PC, one protocol for everyone (NetBEUI or IPX will do fine), and a 10BaseT hub or hubs with enough ports for everyone, and you're ready to go. If you're especially short of cash and don't mind troubleshooting, a simple 10Base2 network will do the job too. All-Macintosh networking is even simpler.

If your network is larger, or it has more complex requirements — a combination of Macs, Windows 95/98/NT, Unix boxes or shared Internet access — you'll have more planning to do. The trick is to give everybody enough network bandwidth without going overboard and giving everyone dedicated 100Mbit connections to download 1K emails.

A network switch, which you'll find discussed in more detail in the Hardware chapter, allows you to hook up heavy-traffic machines to a 100Mbit connection, medium-traffic machines to a 10Mbit connection, and plain desktop machines can share a 10Mbit network. This is about as complex as most small business networks get, but the sky's the limit once you start playing with bridges, switches and routers.

Above opposite is an example of such a network. The file server is the only device attached via a 100Mbit connection. There are four power users and a print server that each have a 10Mbit connection to themselves, and another 10 ordinary users running from a cheap 10BaseT hub that has another 10Mbit connection, via the switch, to the other devices. The switch stops the

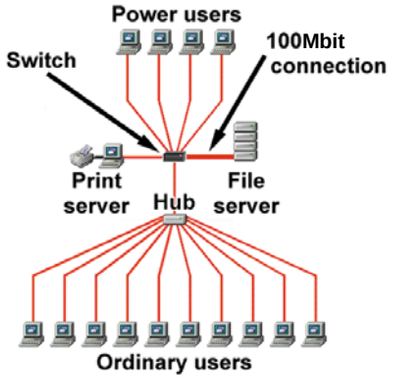

various devices connected to it from using up each other's bandwidth, so no matter how much the 10 ordinary users are chattering among themselves, any of the power users can always obtain a whole 10Mbit connection to the file server.

That said, many small office networks will just look like the bottom part of this example — a 10BaseT hub and a collection of PCs. It's very simple and it works.

What you'll need

Network cards

Shopping for basic network hardware these days is simple. For practically all purposes, one network card with given basic specifications is much like another, so it's safe to buy based on price.

On the next page a few current-model PCI network cards are illustrated. Like many current NICs they have a cut-down shape that gives the impression that circuit board fibreglass must be much more expensive than one would suspect. The one on the left, with only an RJ-45 socket, is a 10BaseT/100BaseT card which sells for about $50. The other, which has RJ-45 and BNC connectors, is a 10BaseT/10Base2 card and costs maybe $25 more.

If you have a PC running Windows 95 or 98 and you want a smoother ride, get a PCI network card. Windows 95 and 98 support Plug and Play, the automatic-configuration system that drags PCs grudgingly towards the ease of setup that Macintoshes and the late lamented Amiga had in the 1980s.

All PCI network cards support Plug and Play. All recent ISA network cards do too, but the zillions of old ISA cards still floating around do not. Any PC with PCI slots should support Plug and Play, and more recent ones support it for ISA as well. If your computer doesn't have PCI slots, it probably doesn't have Plug and Play support.

Cables

Network cables purchased from a store are generally good quality. The only things you have to remember are the magic names:

- 10Base2: requires RG-58 coaxial cable; accept no substitutes
- 10BaseT: will run perfectly well on Category 3 UTP cable
- 100BaseT: requires the slightly more expensive Category 5 cable. Do not accept cable touted as 'Category 5 quality'; it has to be the genuine article.

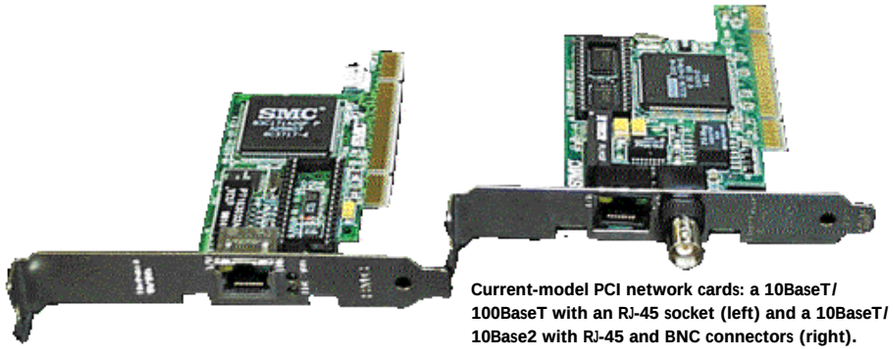

Current-model PCI network cards: a 10BaseT/100BaseT with an RJ-45 socket (left) and a 10BaseT/10Base2 with RJ-45 and BNC connectors (right).

If you're wiring up an office for 10BaseT, it is sensible to buy Category 5 (or 'Cat5') cable, so you can upgrade to 100BaseT without running new cables. You'll get a greater cable range with Cat5 cable, too, though it isn't wise to push your luck. Many places only stock Cat5 cable these days; a 5m Cat5 cable should cost less than $20. For the same money, you can get a 10m 10Base2 cable. 10Base2 terminators and T-pieces cost a few dollars each; if you choose to use 10Base2, buy some spares. T-pieces and terminators don't actually go bad very often, but they do seem to run away to the same place as paper clips and pens.

If you're somewhat handy and choose to make your own cables, you'll be pleased to know it isn't very difficult. There are special crimping tools that allow you to rapidly attach BNC or RJ-45 connectors to cables. The cables and connectors are sold in all good electronics stores.

There is no alternative to crimping for RH-45 connectors, but you can get BNCs in forms you can solder or screw on. Solderable connectors are dependable but more difficult to put on the cables than crimped connectors; screw-on connectors are renowned for their unreliability.

Hubs

A five-port 10BaseT hub like the one on the opposite page, which also has a BNC connector, sells for about $100. The fifth port is also this hub's 'uplink' port, through which it can be cascaded to another hub. If you do this, you lose the use of the fifth port. Many hubs work in this manner.

An additional $25 will buy you another three ports and slinkier styling. You'll have a hard time paying more than $500 for a non-name brand 10BaseT hub; $500 will get you 24 ports. 100BaseT hubs, most of which also support 10BaseT connections, are more expensive. You'd be looking at least $250 for a four-port model and up to $2,000 for 24 ports. Remember that since 10BaseT and 100BaseT hubs are just repeaters, which take whatever comes in one port and transmit it out of every other one, connecting just one 10BaseT or 10Base2 device to a dual-mode hub will force it to run at that speed. They can't store data to feed down the 10BaseT connection as fast as it can take it, so they have to make every port the speed of the slowest one.

Hubs, like network cards, are commodity items these days. You can buy based on the description on the side of the box, and it's safe to buy hubs that aren't made by a big-name manufacturer. You'll pay significantly more for a hub made by Intel or 3Com or one of the other big names, but you won't necessarily get a better performance product.

Switches and routers

These are more advanced network devices, which will be covered in more detail in the Hardware chapter, and vary widely in price. A basic switch with two 100Mbit ports and five 10Mbit ones — which can be used like a dual-mode hub but doesn't choke back traffic to the speed of the slowest port unless all of the traffic actually has to *go* to the slowest port — will set you back $900. You'll pay the same for an eight-port 10/100Mbit switch that can take expansion modules to increase the number and type of ports as your network grows. Routers start from under $1,000 and go to . . . well, think of a number and double it — several times.

If you're wondering whether you need a standalone router, or even a switch, the answer is probably no. Find out in the Hardware chapter.

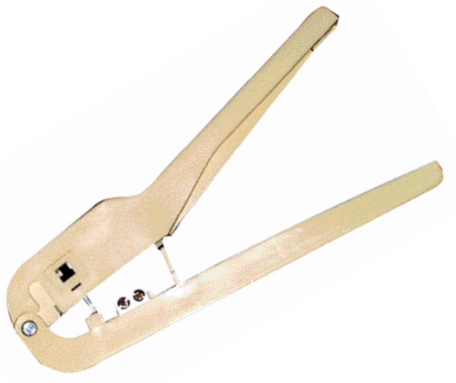

An RJ-45 crimping tool for those who choose to make their own cables.

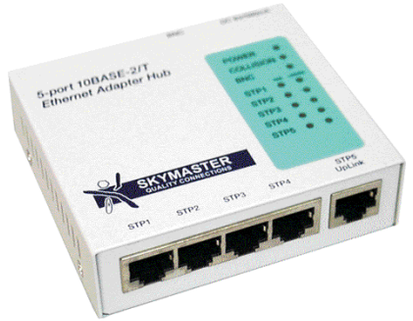

A standard five-port hub.

A more expensive eight-port hub.

Software and drivers

As mentioned above, all current operating systems have built-in support for networking. If you are connecting systems of different types you may need to add some extra stuff, but you will probably just need to know how to use what you already have properly. See the Software chapter for more information.

One thing you *will* need, though, is a proper driver for your network card. Windows 95, and especially Windows 98, are pretty good at providing drivers for all sorts of devices, and almost any network card that says it's 'NE2000 compatible', referring to the old Novell card that set the standard for some years, will work perfectly with a Windows NE2000 driver. New network cards almost always come with a good Windows 95/98 driver, and usually with a Windows NT driver as well. There's a good chance these drivers will work fine; if they don't, the manufacturer's Web site often has a newer one that does. If you run a different operating system, make sure the card will work with it.

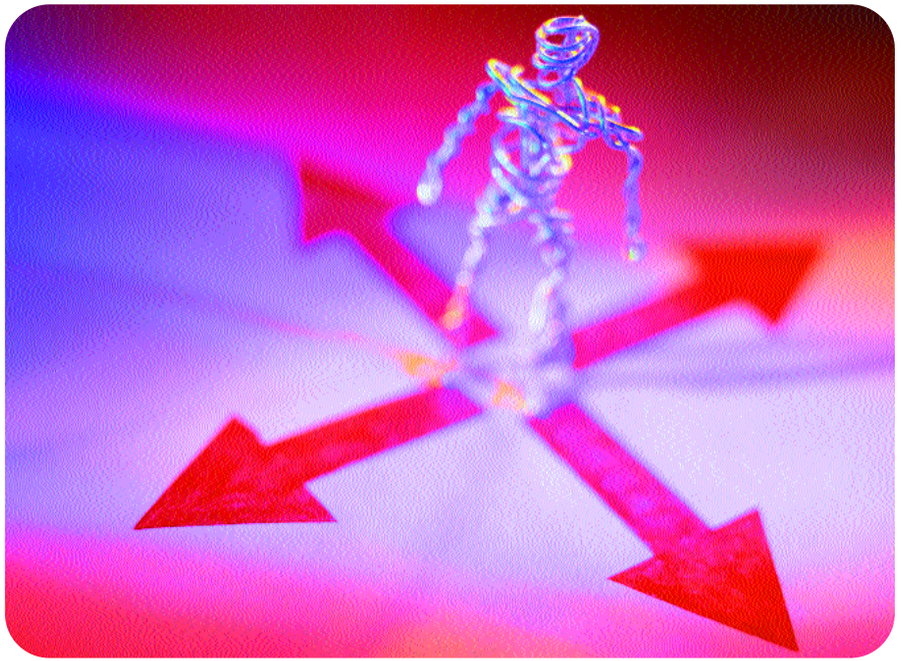

Common terminology

IN THIS POCKETBOOK YOU'LL SEE A LOT of networking jargon used to describe the components of a network and how they operate. Understanding these specialised terms will make it a lot easier to figure out what's connected to what and why. Chances are you've heard some of these terms before, but take the time to find out what they mean and where they fit in the bigger networking picture.

DNS: Domain Name System or Service

DNS is an Internet service which can also be used on TCP/IP LANs. It translates host names into IP addresses. DNS allows you to enter a name that's easy to remember — 'www.yahoo.com', for instance, or 'fred' for a machine on a local network — and have it automatically translated into a numeric address. The DNS system is itself networked; DNS servers which don't recognise a given host name have the ability to ask other DNS servers until an answer is found. DNS doesn't work with dynamically allocated IP addresses (see DHCP below).

A DNS is unnecessary for small networks although many people use DNS every day when using a Web browser. That DNS server is administered by your Internet Service Provider, though, and you don't need to worry about it, beyond telling your computer what the IP address of the DNS server is when you set up your Internet connection.

Larger networks that use TCP/IP can benefit from a DNS server. It's easier, for instance, to set everybody's Internet access proxy to 'bill' instead of '192.168.1.235'.

DHCP: Dynamic Host Configuration Protocol

DHCP is the protocol that allows a network server to dynamically allocate IP addresses to computers on a TCP/IP network. The server doles out addresses from a defined range to any computer which requests one. Computers can have a different address every time they connect to the network or in some cases can even change addresses while still connected. The computers on the network don't need to have an IP address set manually and the network addressing scheme can be changed simply, at will. It's possible to mix DHCP with static, manually assigned addresses.

Most users will only use DHCP when they connect to a dialup Internet Service Provider; most ISPs assign each user an address for the duration of the connection. On larger networks, a DHCP server can greatly simplify network administration, since computers can come and go freely without IP address clashes or endless manual address typing. Most smaller networks don't have a use for TCP/IP, much less DHCP. The chief thing to remember about DHCP is that Windows 95 machines on TCP/IP LANs will try to use DHCP if you don't tell them not to, and will therefore fail to see the network and slow it down (see the Software chapter).

Domain name

Most domain names represent only one IP address, but larger organisations can have

several. The name always includes one or more suffixes — for example, in the URL http://www.apcmag.com/, the domain name is apcmag.com.

Without domain names, ordinary Web browsing and most other Internet use, would be close to impossible. As things stand, domain names connect invisibly to IP addresses via DNS, and changes to the IP addresses cause, at worst, a brief period of inaccessibility before the various DNS servers catch up.

Gateway

In networking, a gateway is a hardware and software combination that connects different kinds of networks. This connection can, for instance, allow users of a network using one protocol to see shared drives on a network using a different protocol, or allow users of one email system to exchange mail with users of a different system.

In old-fashioned parlance, switches and routers could be thought of as gateways, but these days the term is most commonly applied to computers that connect a LAN to the Internet. This is the only kind of gateway you're likely to find on most smaller networks.

Host name

A host name, or site name, is a unique name by which a computer is known on a network. On the Internet, host names are composed of a local part at the beginning and a domain name at the end, as in fred.company.com .au, where 'fred' is the local part and 'company.com.au' is the domain name. Host names have to be translated into an actual network address (in the case of the Internet and other TCP/IP networks, an IP address) by something — a DNS, in the case of the Internet. One computer can have several host names, but multiple computers on a network cannot share one host name.

Most small networks have host names set up automatically; in Windows networking, the host name is the 'Computer Name' set in the Identification tab of Network Properties. If you know this name, you can use the 'Find —> Computer' option from the Start menu to locate the machine even if it's only just been connected to the network and hasn't yet appeared on the list of connected computers.

Name server

Software that translates names from one form to another. A DNS server is a name server.

NetWare

Novell NetWare is Novell's proprietary networking operating system for IBM-compatible computers. NetWare uses the NetBEUI, IPX/SPX or TCP/IP protocols, and works with MS-DOS, Windows, OS/2, Macintosh and Unix clients. Windows can now do everything NetWare can do, and more simply, but NetWare survives in many current systems.

Route

A route is the series of devices that network traffic passes through on its way from source to destination, or a possible path that data *could* take from one place to another. In Windows, you can see the route your data takes to a given host by opening a DOS window and typing 'tracert [address]', where [address] is the host name or IP address of the destination.

Subnet

A subnet is a portion of a network which works independently of the rest of the network, and shares a common address component. In TCP/IP, subnets have the same prefix in their IP address. It may actually be physically separate from the rest of the network and therefore technically a complete network in and of itself, but physically connecting it to the network shouldn't make any difference to its operation. Hardware considerations may make the whole network slower when subnets are physically connected, but network transactions should work in the same way. See also 'subnet mask'.

If you're running a TCP/IP LAN, you can set up multiple subnets to make sections of the network invisible to each other.

Subnet mask

On an IP network, subnets are divided with a mask that takes the same x.x.x.x four-number form as an IP address, but has only 255s or 0s as the numbers. The subnet mask takes the form it does because IP performs a 'bitwise AND' operation on both the mask and on a given IP address, to determine what subnet the address belongs to. The 255s in a subnet mask 'let through' the corresponding numbers in an IP address, and the 0s 'block' the corresponding numbers.

Thus, if the subnet mask is 255.255.255.0 (a 'class C' network), and a couple of sample IP addresses are 192.168.1.101 and 192.168.1.145, the two addresses are defined by the subnet mask as being on the same subnet. 192.168.10.139, for instance, wouldn't be on the same subnet as the two examples with this subnet mask, but would be with the subnet mask 255.255.0.0 (a `Class B' network). The portion of the IP address that's passed by the subnet mask is the 'network address', and the other portion is the 'host address'.

If you're running a TCP/IP LAN, you must set a subnet mask. It's worth remembering, though; if your TCP/IP network isn't visible to some users, and their IP address is correct, double-check their subnet mask and make sure it's the same as everyone else's.

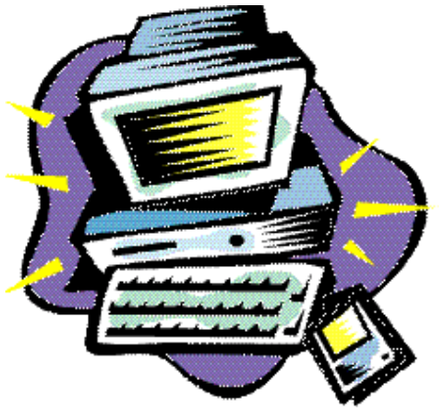

WINS: Windows Internet Naming Service

This is Microsoft's system for resolving NetBIOS names into IP addresses on a LAN. WINS has the ability to deal with dynamically allocated IP addresses (see 'DHCP'), which DNS cannot. It achieves this by distributing a database of machine names that are currently connected and addresses.

This system would become clumsy on a network with a very large number of nodes, such as the Internet, which is why computers with dynamically allocated IP addresses (for example, dialup Internet users) cannot be accessed via a host name by others on the Internet. WINS itself is unsuitable for Internet use in any case, because it handles only NetBIOS services, like shared drives and printers on a Windows machine.

A WINS server can do for the Windows machines on a TCP/IP network what a DNS server does on any TCP/IP network, but it can also handle DHCP. It's therefore a more attractive prospect for larger Windows-only networks or networks where it's OK for only the Windows machines to receive the benefits of name resolution.

Workgroup

Under Windows, the 'Workgroup' entry in the Identification tab of Network Properties works as a kind of subnet identifier; only computers with the same workgroup name can see each other. If you're having problems seeing other machines, check the workgroup names. In general networking terms, it refers to the computers people are using and the network that connects them, so they can exchange email and files.