Post restante

Bob Walder shows how you can pick up your email anywhere, from a standard browser.

Last month, I told you how to get your Exchange Internet Mail Connector (IMC) up and running and so by now you should be merrily sending and receiving mail via your very own SMTP/POP3 mail server. Many Exchange users will use the Outlook client, too, although with the internet mail support we installed last month that is by no means compulsory as any POP3/IMAP4 client can now be used to send and receive mail via your Exchange Server.

No matter which client you use, though, you can frequently find yourself incommunicado when you are out of the office without your desktop or notebook PC. You cannot simply walk up to any old PC and check your email using someone else's client. Well, you could but it would involve some serious reconfiguration as well as inconvenience to the owner of the PC you are borrowing.

There is a way around the problem, though. While you do not pick up your email from the mail server it simply resides there in limbo until you return to your office and download it again. By installing Outlook Web Access (OWA), you gain access to your Exchange mailbox over the internet using nothing more than a standard browser [Fig 1]. With Exchange 5.5 came additional support for OWA including contacts and calendaring support over the web, plus access to Public Folders.

The first step is to ensure that Internet Information Server (IIS) and Active Server Pages are installed and working correctly. The next is to ensure that the OWA components have been installed. OWA must be installed on the same server as IIS and Active Server although this does not need to be the same machine which is hosting the Exchange Server accessed with these scripts.

Note that IIS 4.0 is not supported by the Exchange Server version 5.0 Active Server Components [Fig 1] installing them both on the same computer results in error messages and the OWA client returns error messages. Exchange Server 5.5 does support Internet Information Server 4.0, though. To use Exchange with the OWA client you must install Exchange Server 5.5 before installing IIS 4.0.

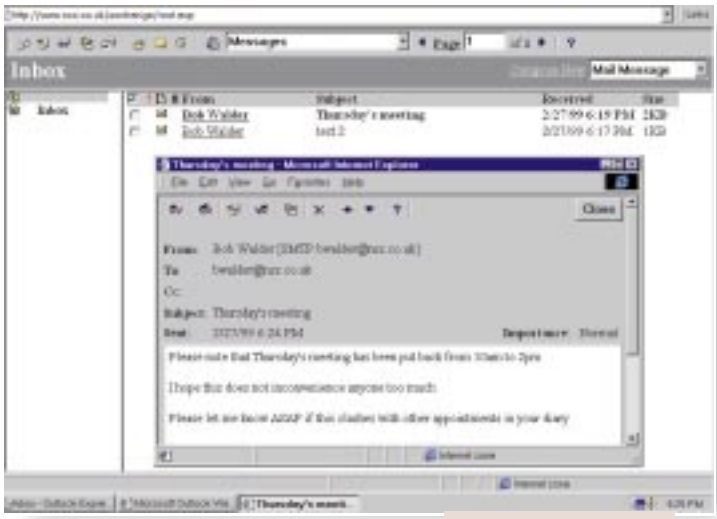

▲ Fig 1 Reading a message using OWA

■ Installing OWA

To install OWA on your IIS/ASP server, first insert the Exchange Server 5.x CD and then run server setup.

- **Select** custom install and the OWA option. If this is an existing installation, simply run the Set-Up program and select the Add/Remove Components option.
- **Select** OWA in the Options list and click Continue. This will install the Collaboration Data Objects (CDO), the CDO Rendering library and various Active Server Pages and associated script files.

The next step is to configure Exchange Server to support the necessary protocols.

- **Log on** to the server as Administrator and start the Exchange Administrator program.
- **Expand** the Site container of the organisation tree and select the Protocols container. The key one here is the HTTP (web) Site Settings: double click on it.
- **Make sure** the Enable Protocol box is checked [Fig 2].
- **Confirm** the settings in the Anonymous Access section if you wish to allow anonymous users to access the public folders or browse the global address list.
- **Click** on OK. Double click on the LDAP (Directory) Site Defaults.
- **The port** number should be OK. Check the Enable Protocol box.
- **Check** the Allow Anonymous Access box on the Anonymous tab. Strangely enough, even if you are not planning to allow users to browse your Exchange Directory using LDAP you still need to perform these steps to get OWA working.
- **Double click** on POP3 (Mail) Site Defaults.
- **Check** the Enable Protocol box. And do the same for IMAP4 and NNTP if you

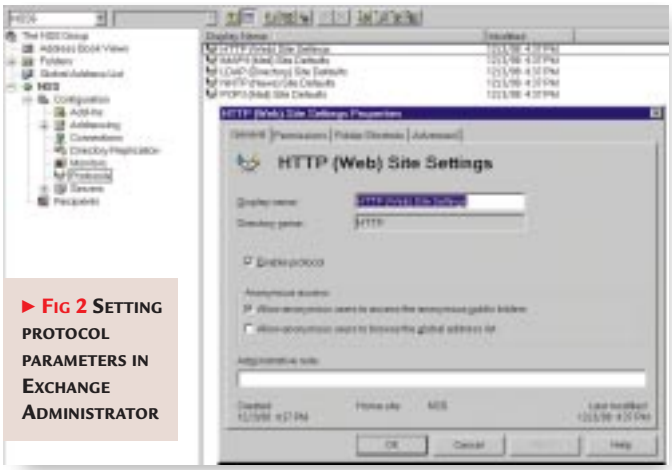

▶ Fig 2 Setting protocol parameters in Exchange Administrator

▲ FIG 3 OWA LOGON PAGE

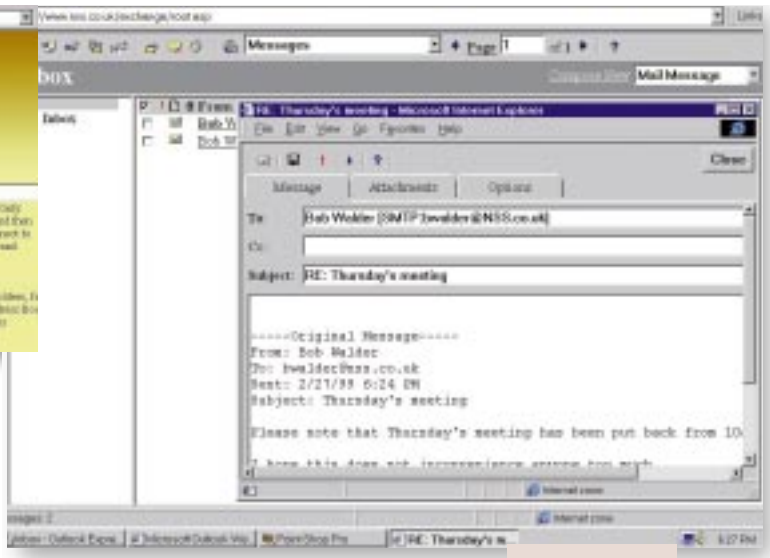

▲ FIG 4 COMPOSING/ REPLYING TO A MESSAGE USING OWA

wish to use these protocols. Although the HTTP protocol applies to the entire site it is possible to override the settings for IMAP4, LDAP, NNTP and POP3 on a per server, or per user basis. Simply expand the Servers container, expand the server in which you are interested and then select the Protocols container. To override settings for each user simply select the mailbox from the Recipients folder and select the Protocols tab. This provides a fine-grained level of control for web access.

If you have done everything right, and your web server is behaving itself, you should now have web access enabled. To try it out, fire up your web browser and type in the address of the OWA directory on your web server. If your normal web server address is www.mycompany. co.uk, you will need to enter www.mycompany.co.uk/exchange. This will bring up the OWA logon screen [Fig 3] where you enter the name of your Exchange alias and click on the 'click here' prompt. If the mailbox is found, you are then prompted for your network login name and password: these should be entered just as if you were logging on to the network back at the office. Once past the logon screen you have something similar to Fig 1, with the toolbar down the left and the inbox full of outstanding messages on the right-hand side. Each message in the inbox is a hyperlink which can be accessed via a single click and from here you can compose, read [Fig 4], forward, reply to or delete your messages, as well as access calendar and contact details. You can also search the Exchange directory for email addresses and phone numbers and browse the public folders.

You can now sit at any PC in the world that is connected to the internet, and has a standard web browser, and gain instant access to your email, calendaring and contacts data. There is even a cut-down version of OWA for use with Pocket Explorer on CE devices.

☛ *In a future column, I will cover how to make your Exchange Server function as an Internet News server.*

THE ZERO COST OPTION

In my recent remote access workshop piece I mentioned the Dial-Up Server for Windows 9x and said that although it was included in Windows 98 you would need to acquire the Plus! Pack for Windows 95 to get hold of it. Although this is perfectly true I neglected to mention the 'zero cost' option (shame on me). John Robson wrote in to put me right.

'Thanks for the "Home and Away" remote access workshop in the April issue. I've just tried Dial-Up Server. It worked a treat and promises to be very useful. However, you comment that "Windows 95 users will have to purchase the optional Plus! Pack to acquire it".

'Actually, you can download DUN 1.3 from www.microsoft.com/windows/downloads/contents/communications/dun13win95/default.asp and this includes Dial-Up Server. This could save readers the cost of the Plus! Pack or an unnecessary upgrade to Win98, the two options I had considered until I found DUN 1.3 on the web.

'I still run Win95 OSR 2 with the USB supplement and have yet to see anything about Win98 to persuade me that it's worth shelling out to upgrade. Mind you, I got close this time.'

Thanks for the tip, John. For the record, I think there are enough new features in Windows 98 to warrant the upgrade although it depends on what you want from your system, and I suppose your OSR2 release with USB will give you most of them! Those with an earlier release of Win95 might be interested in the increased range of updated drivers conforming to the new Windows Driver Model, improved NetWare client with support for NDS, ISDN support built in, client support for PPTP, and FAT32. The last is especially interesting, being an improved version of the FAT file system that allows disks over 2Gb to be formatted as a single drive. It also uses smaller clusters than do FAT drives, which results in a more efficient use of space on large disks. FAT32 alone was worth the upgrade cost to me.

PCW CONTACTS

Bob Walder welcomes your comments on the Networks column. He can be contacted via the PCW editorial office (address, p14) or email networks@pcw.co.uk

Mail model

Bob Walder looks at Exchange Server, showing how it can pick up and distribute your email.

Last month, I looked at internet mail and the differences between SMTP, POP3 and IMAP4. I covered the most flexible method for an organisation to handle its email: registering a domain and using its own SMTP server. I ended by taking a look at two examples of widely available SMTP servers; NT Mail and Exchange Server.

I've said before that I will concentrate mainly on Microsoft products for this series, so my apologies to anyone with another SMTP server but a lot of the principles apply even if the details differ. There are good reasons to concentrate on Exchange apart from the fact that Microsoft's marketing clout will ensure its success whether we like it or not.

Compared with other products it is very flexible. Sure, if you need nothing more than a simple SMTP server it is overkill but most organisations will use Exchange as their internal mail server, group scheduler, bulletin board and news server, as well as for internet mail. One of the features Exchange offers, which I find extremely useful when I am out and about travelling, is the web access. This allows you to access your mailbox and process your mail using any standard web browser from anywhere on the internet. It is ideal if you are travelling without your laptop machine.

The price of flexibility, though, is complexity and Exchange is quite capable of confusing you with a whole shed-load of options that you simply do not need. Here, I will look in more detail at how to persuade it to pick up and distribute your internet mail. I will assume that you already have Exchange installed and working as an internal mail system, and that your network has a routed connection to the internet.

In order to have Exchange handle your internet mail, too, you need to install and configure an Internet Mail Connector (IMC). This comes as part of the Enterprise Edition or can be purchased as an optional extra to the Standard Edition.

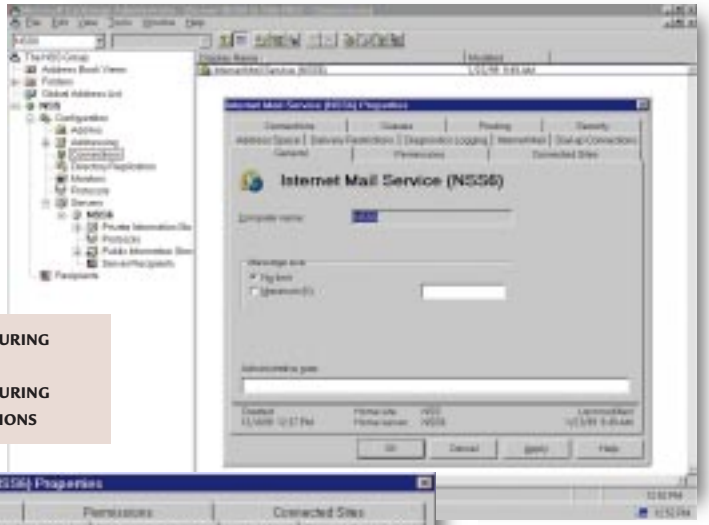

►Fig 1 Configuring IMC: General
▼Fig 2 Configuring IMC: Connections

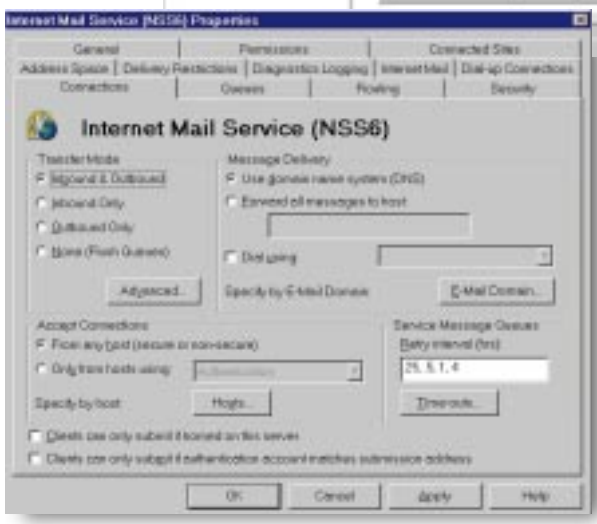

Installing and configuring the IMC is fairly simple.

1 Log on to the server as Administrator and start the Exchange Administrator program.

2 Expand the Site container of the organisation tree and select the Connections container.

3 Select New Other from the File menu, and then select Internet Mail Service from the drop-down menu. This fires up the Internet Mail Wizard that provides a check-list of tasks to complete before installing and configuring the connector.

4 As we mentioned last month, you should make sure that you have your domain name registered and that your ISP has a Mail Exchange (MX) record in its DNS server to point to your mail server. If you have your own DNS server, the MX record should be created there.

5 Stepping through the Wizard you are asked for the name of the Exchange Server on which to install the IMC. You are also offered the option to configure it to send internet mail via a dial-up connection through your ISP using RAS (Remote Access Service). Here, however, we will assume you have a dedicated routed connection to your ISP.

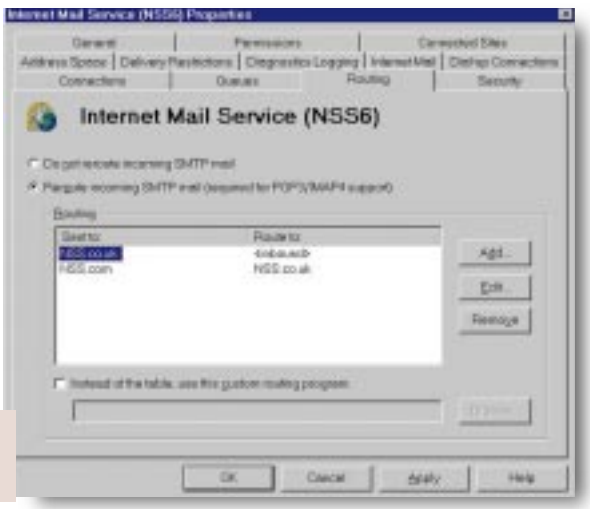

►Fig 3 Configuring IMC: Routing

6 Once the IMC has been installed, you can double click on the entry in the Connections folder to see a number of configuration options [Fig 1]. Most of these can be left as the defaults.

7 The first one to check is the Internet Mail tab. Click on the Change button in the Administrators Mailbox section and select the user who should receive all mail notifications. Clicking on the Notifications button allows you to specify which types of non-delivery reports will generate notifications to the administrator.

8 Ensure that the MIME option is selected in the Message Content/ Attachments (outbound) box, and if you want to use digital signatures ensure that the Clients support S/MIME signatures option is selected.

9 You may want to check the Advanced button where you can enable 'out of office responses' and 'automatic replies' to the internet if required.

10 Next, select the Connections tab [Fig 2]. Here you can set the Transfer Mode to process Inbound and Outbound messages, Inbound Only or Outbound Only. The Message Delivery option should be set to Use Domain Name System.

11 Select the Routing tab [Fig 3]. If you just want to accept inbound mail and have it distributed to Exchange mailboxes, then select the Do Not Reroute Incoming SMTP Mail option. However, if this machine is to operate as a POP3/IMAP4 server for internet mail clients you will need to select Reroute Incoming SMTP Mail. You then need to add at least one entry to the Routing box. Click on Add, enter your domain name (NSS.co.uk in our case), and click on Should Be Accepted As Inbound.

Now we have defined an IMC that will process all inbound and outbound mail, and where all inbound mail for the domain NSS.co.uk will be routed to POP3/IMAP4 mailboxes for users to collect.

☛ **Exchange now needs to know** how to route those mail messages. Bear in mind that an Exchange Server will have mailboxes set up as 'Bob Walder' or something similar, so it needs some means of translating an internet address of bwalder@NSS.co.uk in order to place the message in the correct mailbox. This is achieved via the Site Addressing facility.

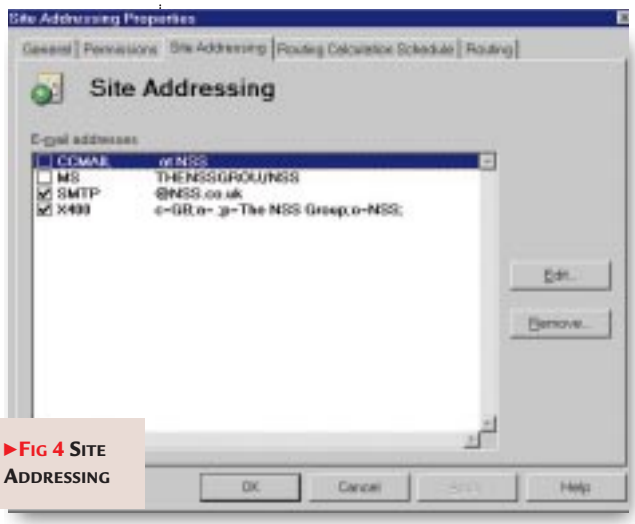

▶FIG 4 SITE ADDRESSING

1 Select the Configuration container and double click on the Site Addressing option.

2 Click on the Site Addressing tab. See Fig 4.

3 Ensure that the SMTP addressing option is checked. Select it and click on Edit.

4 The type is fixed as SMTP and the Address format is shown as '@NSS.co.uk'. This means that all SMTP addresses for this site will be created as 'alias@NSS.co.uk' by default, where the alias is determined by the user properties. Yours should be in the format '@domain' where domain is your registered domain name.

5 Now go to the Tools menu (from the Exchange Administrator tool bar) and select Options, then select the Auto Naming tab. This is the section that determines how mailbox display names and aliases are generated for new mailboxes (the defaults here can be overridden).

6 If you want the display name to be 'Bob Walder' select the 'First Last' option. For an alias, we use first initial followed by the surname so the 'First Initial And Last Name' option is checked. When we create a mailbox for Bob Walder, this configuration results in a display name of Bob Walder and an alias of bwalder.

If you then click on the E-Mail Addresses tab you will see that the SMTP address has been generated automatically as 'bwalder@NSS.co.uk' [Fig 5]. This is how Exchange knows where to route incoming messages destined for the NSS.co.uk domain.

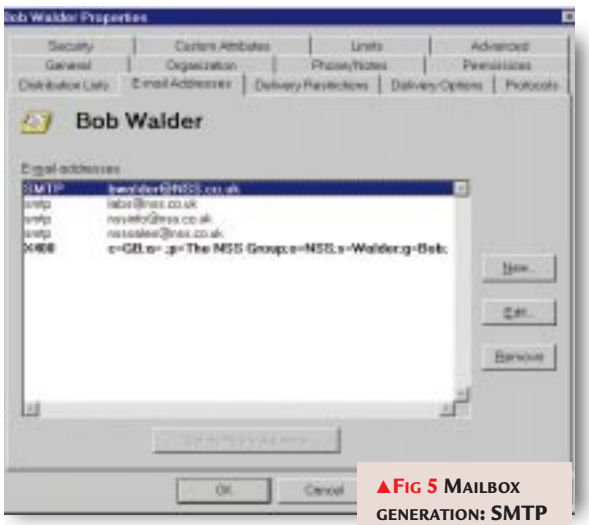

▲FIG 5 MAILBOX GENERATION: SMTP ADDRESSES

Look again at Fig 5. Apart from the Auto Naming entries for SMTP and X.400 we saw in the Site Addressing section, there are a few other addresses for this user, too. These are commonly known as 'aliases' in other SMTP mail servers. They are created by clicking on New then on Internet Mail Address, then entering the address in full. In this particular case, any messages addressed to labs@NSS.co.uk will automatically be routed to the mailbox of Bob Walder.

☛ **Next month** I will tell you how to get your Outlook Web access working so you can access your email from anywhere on the internet using only a standard web browser. So, no more having to lug your laptop around with you when you travel, just to access email.

PCW CONTACTS

Bob Walder can be contacted via the PCW editorial office(address, p14) or email networks@pcw.co.uk

First class post

Use the power of the net to run your own email server, in-house. Bob Walder shows the way.

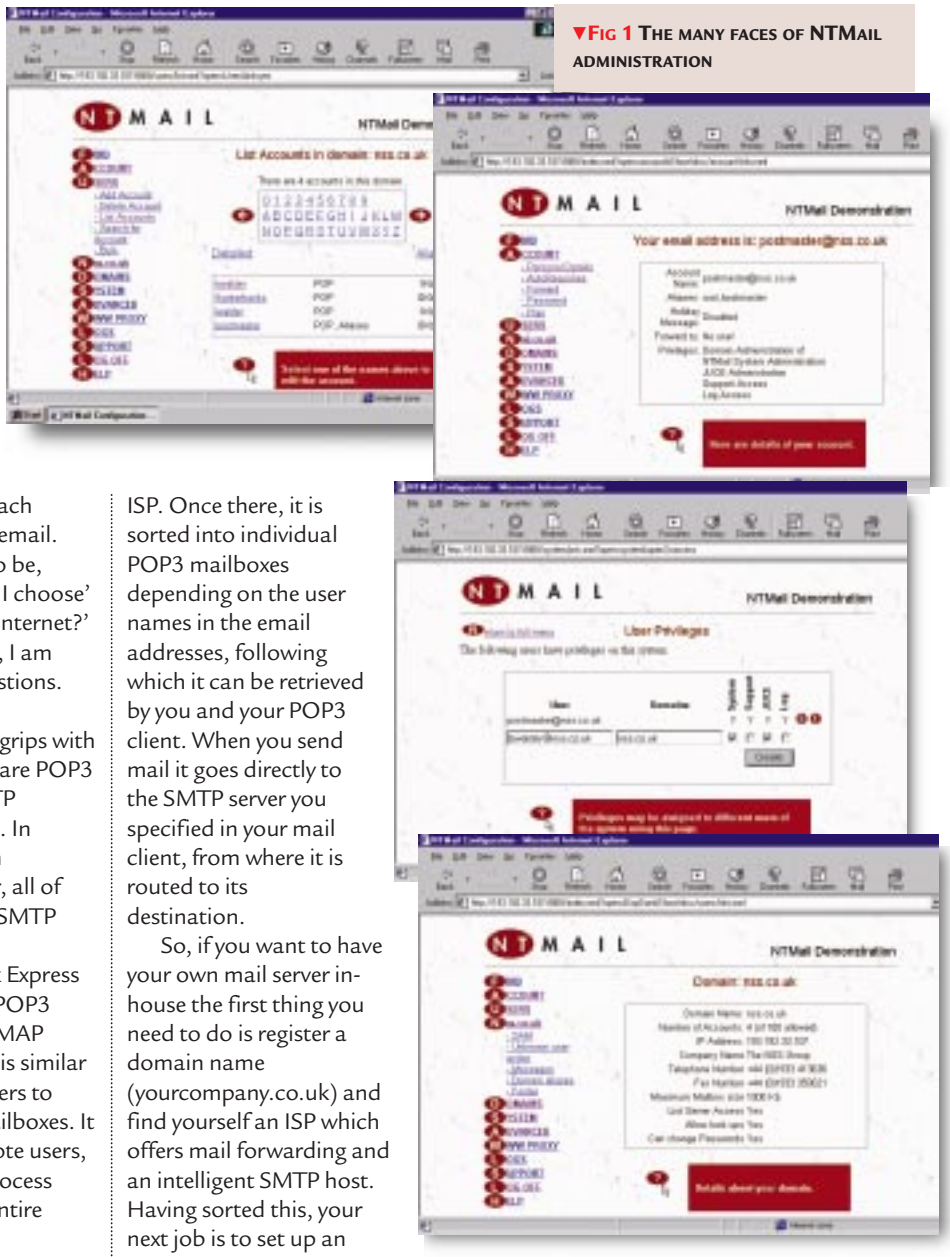

▼FIG 1 THE MANY FACES OF NTMAIL ADMINISTRATION

These days, few organisations are without some form of email communication, whether they operate their own mail server or rely on an Internet Service Provider (ISP). Some are still using email purely as an internal messaging medium while the vast majority have recognised the power of the internet to act as a global transport mechanism for their inter-company electronic mail.

I receive numerous queries each month regarding the subject of email. The two main questions seem to be, 'what sort of mail server should I choose' and 'how do I connect it to the internet?' Over the next couple of months, I am going to try to answer both questions.

The main acronyms to get to grips with when it comes to email services are POP3 (Post Office Protocol) and SMTP (Simple Mail Transfer Protocol). In simple terms, POP3 provides an individual mailbox for each user, all of which are usually hosted on an SMTP mail server at your ISP.

Mail clients such as Outlook Express allow you to retrieve mail from POP3 mailboxes on an SMTP server. IMAP (Internet Mail Access Protocol) is similar to POP3 in that it allows end users to retrieve mail from individual mailboxes. It provides more facilities for remote users, though, such as the ability to process headers without retrieving the entire message, and so on.

An SMTP server is always required somewhere along the line since this is how the internet ships its mail around: from client to SMTP server and between SMTP servers whenever necessary.

What sort of mail server should I choose?

The normal scenario is that mail for your account (or domain) is forwarded to a specific SMTP mail server at your ISP. Once there, it is sorted into individual POP3 mailboxes depending on the user names in the email addresses, following which it can be retrieved by you and your POP3 client. When you send mail it goes directly to the SMTP server you specified in your mail client, from where it is routed to its destination.

So, if you want to have your own mail server in-house the first thing you need to do is register a domain name (yourcompany.co.uk) and find yourself an ISP which offers mail forwarding and an intelligent SMTP host. Having sorted this, your next job is to set up an SMTP host at your end to talk to it.

In other words, you no longer want your Information Service Provider to sort mail into mailboxes for you. Instead, you want *all* the mail for your domain to be forwarded *en masse* to your own SMTP mail server in–house. In effect, you become an ISP for your internal users. There are several ways to go about this.

NTMail is one inexpensive option designed to drag mail from your ISP's SMTP host and hold it on your server until someone with a POP3 mail client connects to inspect their mailbox.

NTMail was written from the ground up as a Windows NT package and was not ported from an older 16-bit Windows or Unix application. This means that it is designed to integrate completely with the NT Operating

System and can thus take advantage of OS features such as multi-threading and multiple processors to provide excellent performance.

NTMail operates as a number of native NT services: one for the Configuration Server and another for each of the other types of service supported such as POP, IMAP, LIST, POST and SMTP. It also integrates closely with NT in its use of the Performance Monitor and also in terms of user security where the native Windows NT system database may be employed with or without NTMail's own user database. Administrators can use the familiar NT Performance Monitor to keep track of critical stats such as the number of messages per second, how many messages are queued, number of posts to lists and so on.

Installation is straightforward. With version 4, all maintenance operations can be performed via a simple web browser-based interface. This is provided by the Configuration Server which runs on a different port from any existing web servers allowing it to coexist happily with IIS, Netscape Commerce Server and the like.

A series of simple forms allows users to be added and removed, mail to be read, domains to be added and configured, and so on. In addition, specific users can be granted the right to administer particular domains.

Multiple domain support is one of the neat features of NTMail, since it is far simpler to implement than some of its rivals and each domain appears to the outside world as a completely separate system, even though they all reside on the same physical box.

As you would expect, NTMail complies with the relevant internet standards including POP3, IMAP4 and MIME (Multipurpose Internet Mail Extensions). Wherever possible, non-compliant email is made compliant by NTMail so that other servers have no problems processing it. Clearly, there is a huge range of clients that will work with NTMail.

There are two major types of log within NTMail; activity, and the email itself. It will log each inbound or outbound message as it passes through the server, recording the message itself, the time, its origin and destination. All logs 'roll over' at the end of each day, making it easy to trace what happened and when.

Other useful features include a simple web proxy, Auto Responders, which returns a standard response when someone sends an email to them, and 'robots' which will accept and process all email for a single domain. Optional extras include the JUCE anti-spam filter and an integrated anti-virus scanner.

Prices range from around an amazing £49 for five mail accounts — ideal for small offices — to about £895 for 250 accounts, or £2,395 for the unlimited version (see the *PCW Contacts* box, below).

Microsoft's Exchange Server is another option, for those with more advanced messaging needs — and bigger budgets! This has undergone several major revisions since its release at the beginning of 1996.

Version 5.0, released in early 1997, was a significant upgrade. Chief amongst the new features was the boosting of internet support. POP3 support appeared, allowing Exchange to host a multitude of internet mail clients not restricted to Microsoft's own. News support (NNTP) was also introduced, allowing Exchange to serve as a true NNTP host for corporate NNTP clients as well as communicate with other internet-based NNTP hosts to replicate newsgroups.

LDAP (Lightweight Directory Access Protocol) and web integration were other significant additions which allowed remote web-based clients to access the Exchange directory and their mailboxes over the net using a standard browser. And, the Outlook client was upgraded to Outlook 97 and shipped with Exchange Server as the preferred client for both internet and corporate mail.

Less than a year later, at the back end of 1997, we saw yet another major upgrade in the form of Exchange Server 5.5 [Fig 2] which, at the time of writing, is currently the latest version. This carried further improvements to the performance and scalability of Exchange Server, increasing backup performance and removing the limit on the size of the message store. Once again, there were additions to internet support with the inclusion of a Chat Service, LDAP3 and IMAP4, further expanding internet functionality and the provision of support for a wider variety of clients.

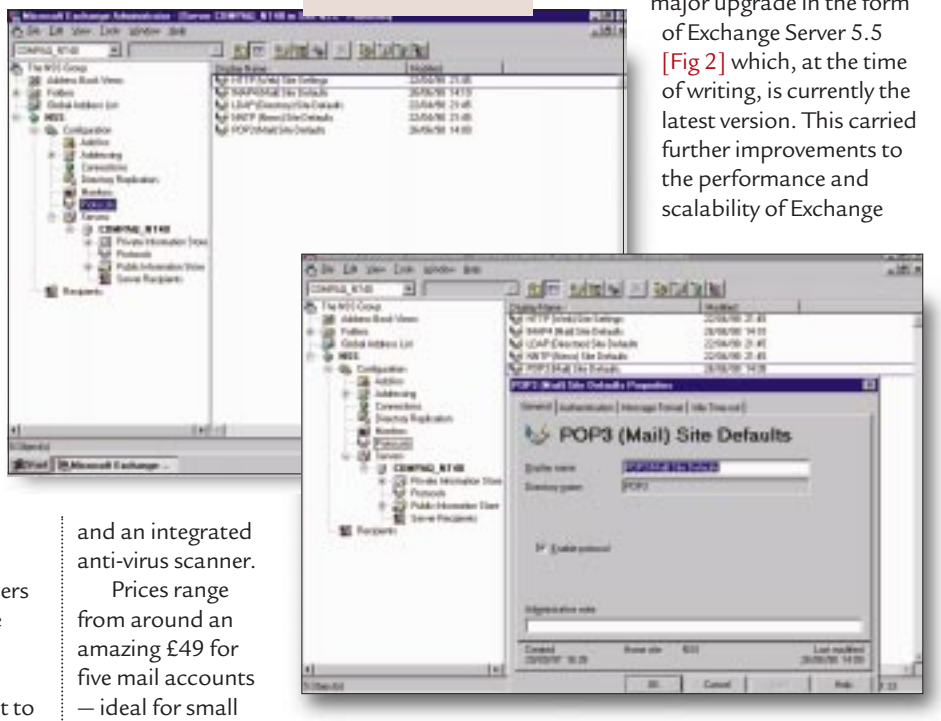

▼FIG 2 THOSE INTERNET PROTOCOLS, IN SUPPORT OF EXCHANGE SERVER 5.5

Next month, I will look in more detail at how to persuade Exchange Server to pick up and distribute your internet mail, and how to access your mail via the web using a standard browser.

PCW CONTACTS

Bob Walder can be contacted via the PCW editorial office(address, p14) or email networks@pcw.co.uk

NT Mail from Internet Shopper 01275 340333, www.ntmail.co.uk

Route it out

Routers are the key to transporting data over a network. Bob Walder hitches a ride.

In past issues I've covered the mysteries of IP addressing and the differences between bridges, routers and switches. This month I'm going to try and pull things together by explaining how routers use IP addresses and subnet masks to determine where your data packets finish up. We'll also learn the significance of that magic parameter in your Windows protocol set-up, the default gateway.

Take a look at our network diagram [Fig 1]. This basically attempts to show a corporate network with two segments joined by router A; segment B uses the addresses 170.1.2.x and segment A uses addresses 170.1.1.x. The netmask for all of these is 255.255.255.0, and if you cast your mind back to the previous piece on IP subnetting, you will recall that segment A can consist of addresses 170.1.1.1 to 170.1.1.254, while segment B addresses the range from 170.1.2.1 to 170.1.2.254. Only the final octet changes within each subnet with this netmask (as indicated by the final 0), so as soon as the third octet changes (from 1 for subnet A to 2 for subnet B) it indicates to any router that the device is on a different subnet.

▼ **Fig 1** Network diagram displaying the role of a router

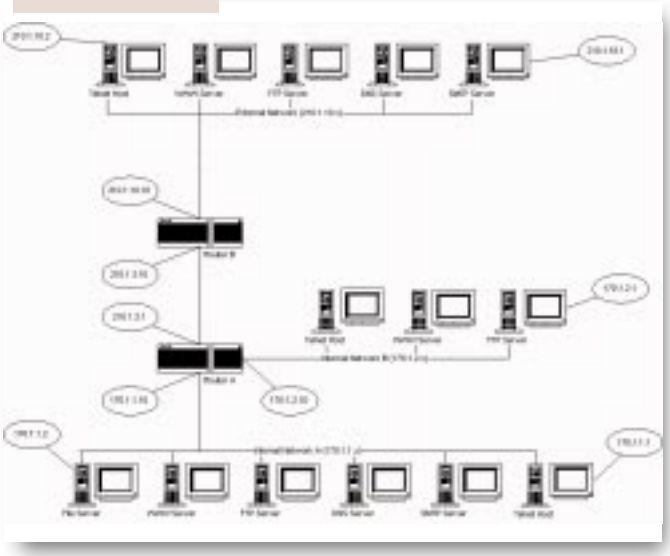

◄ **Fig 2** Setting the default gateway through Windows

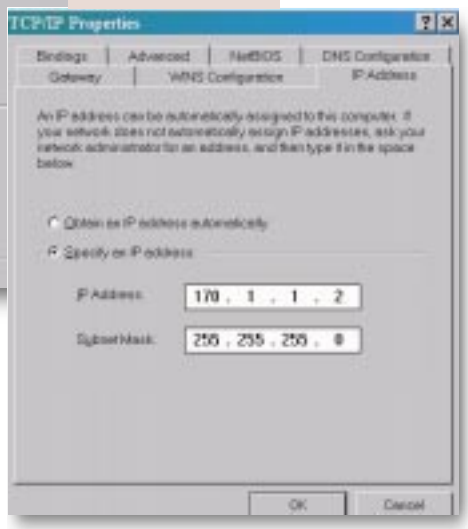

▼ **Fig 3** Setting up an IP address and netmask

Configure it out

So, how do we configure devices on segment A? The file server at the end would have an IP address of 170.1.1.2, a netmask of 255.255.255.0, and a default gateway of 170.1.1.10 [Fig 2]. The default gateway points to the port of the router that is attached to the local subnet. If the file server (170.1.1.2) communicates with the Telnet host (170.1.1.1), the fact that the first three numbers are the same tells it that the destination machine is on the same subnet, and it is not necessary to bother with the default gateway. Our source machine thus uses something called ARP (Address Resolution Protocol) to determine the MAC address of the destination machine (it needs this in order to transmit the packet). The file server sends out an ARP broadcast with the IP address of the Telnet host, and the Telnet machine responds with its own MAC address. To prevent too much of this sort of traffic, each machine keeps such resolved addresses in an ARP cache for future reference. Once the MAC address has been resolved, direct communication can begin between the machines.

Now say our file server wants to communicate with the FTP server on segment B. This time it compares IP addresses and sees that 170.1.2.1 is actually on a different subnet to 170.1.1.2, because the third number of the address is different. At this point, it knows the packet needs to be handled by a device that knows more about the network topology than it does — the router (or default gateway). So, it ARPs for the MAC address of the default gateway, which is 170.1.1.10 in this case, and sends the packet on its way.

Routers are incredibly complex devices, but what they do can be boiled down quite simply: they direct traffic. Like a traffic cop, the router takes packets in from its various ports, checks on their destinations, and sends them off to the appropriate outgoing port. In our case, the router spots that 170.1.2.1 is actually on its second port

by comparing addresses and netmasks again, and so it sends the packet directly to the FTP server.

Two's company

OK, so that has dealt with routing packets within an organisation where the addresses are known. When it comes to connecting our network to the internet, however, there is more than one router involved. In fact, there are millions of the things spread around the world, but we are particularly interested in two: the one at our site and the one at our ISP. Of course, if you wanted to attach a small network to the internet, you could use a proxy server and have all your users go through that: this way, you can often get away with a single-user dial-up account to support a small network. However, we are assuming here that you have opted for a full-routed connection, and so you have to make sure that your router and the ISP router know about each other.

Your ISP will provide you with an address and netmask for the router at its own site [Fig 3, p271] and this is what you need to enter into your own router to enable it to speak to the outside world. What you do, in fact, is create another subnet, this time between the external port of your router and your ISP. In our diagram, this small subnet consists of just two devices: the external port of our router is 210.1.3.1, and the appropriate port of the ISP router is 210.1.3.10. What we effectively do is tell our router that its default gateway is 210.1.3.10, and that is where it will send all the packets it doesn't know how to deal with directly.

The right address

OK, so now our file server wants to communicate with the SMTP server 210.1.10.1 somewhere on the internet (in this case it is actually at our ISP, to keep things simple). It compares IP addresses and determines that the device is not on the local subnet, and so sends it to the default gateway address. Router A checks the address and notes that it doesn't correspond to any of the subnets attached to it directly. So if the router is stumped, what does it do? It sends it to its own default gateway (210.1.3.10), just as the file server did.

Now router B has the packet, and it takes a look at the IP address and sees that 210.1.10.1 is actually on the same subnet as its internal port — 210.1.10.10. It can thus ARP for the MAC address of the SMTP server and deliver the packet directly. Each time a packet traverses a router in this way it's known as a "hop", and the number of hops a packet has to make determines how quickly you receive your data.

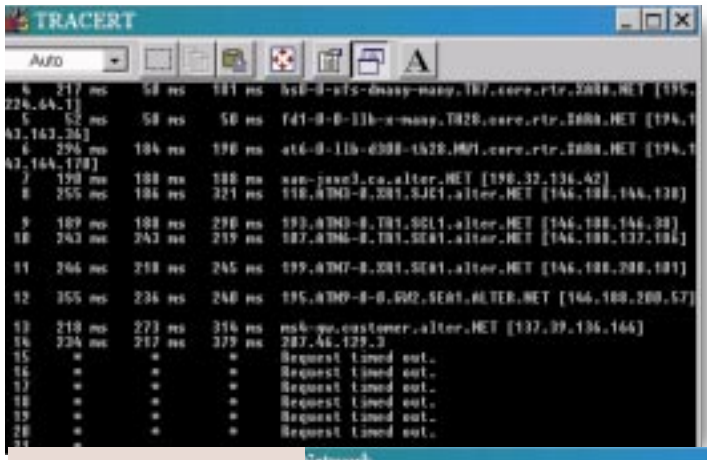

▲FIG 4
A SAMPLE TRACERT OUTPUT

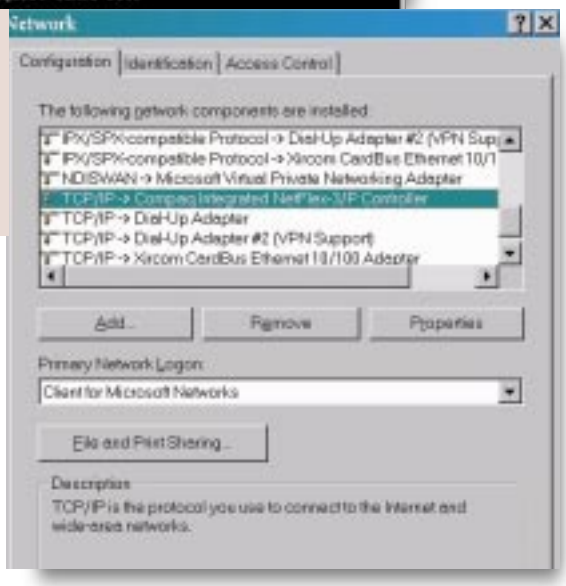

▶NETWORKING CONFIGURATION APPLET FROM THE CONTROL PANEL

Try dropping to a DOS box under Windows and typing TRACERT MICROSOFT.COM. You'll get a display similar to that shown here [Fig 4], with each hop represented by a new router. As you can see, it is sometimes a long, tortuous and slow route to Microsoft!

RIP it up

Of course, given the size of the internet, or even a large corporate network, routers need more than just a 'default gateway' to go on if packets are to find their way from point A to point B some time this week. Routers learn about the network automatically using something called RIP — Routing Information Protocol. This is used to allow routers to tell each other about the segments they are attached to and the addresses they know about. This is done by 'advertising' what they know across the network, whereupon each of them listens and updates its own routing tables accordingly. This is what makes the internet so resilient. Should any router fail for whatever reason, it will stop advertising and the routers that were communicating with it directly will start to find other ways around it. This dynamic change in network topology is completely automatic and is called 'convergence'. Finally, it's also possible to 'tell' a router about a specific route between two points, and these are called 'static routes'.

Hopefully, between this and the IP addressing/subnetting tutorial you have enough information to create your own IP network and get it attached to the internet. If you want to delve into TCP/IP a lot deeper than I have the space to here, you could do worse than check out the book *Windows NT TCP/IP* by Karanjit Siyan (ISBN 1-56205-887-8). Published by New Riders, it costs £26.95 from Computer Manuals (0121 706 6000).

PCW CONTACTS

Bob Walder can be contacted via the PCW editorial office(address, p10) or email networks@pcw.co.uk

The heat is on

Bob Walder tackles firewalls, a crucial weapon in the war against corporate hacking.

In previous columns I have covered the mystic world of IP addressing, WINS, DNS, routing and internet connectivity. Anyone who has followed this little lot and actually got as far as connecting their network (big or small) to the internet will no doubt be worried sick about all those nasty little hackers out there just waiting to break in, rifle their server's disks for juicy data, and wipe everything clean or install terrible viruses, before disappearing back into whatever adolescent world they inhabit. What do you mean, you hadn't even considered it?

One in five respondents to a recent survey admitted that intruders had broken into, or had tried to break into, their corporate networks, via the internet, during the preceding 12 months. This is even more worrying than it sounds, since most experts agree that the majority of break-ins go undetected. For example, attacks by the Defence Information Systems Agency (DISA) on 9,000 U.S. Department of Defence computer systems had an 88 percent success rate but were detected by less that one in 20 of the target organisations. Of those, only five percent actually reacted to the attack, says the National Centre of Supercomputing Applications (NCSA).

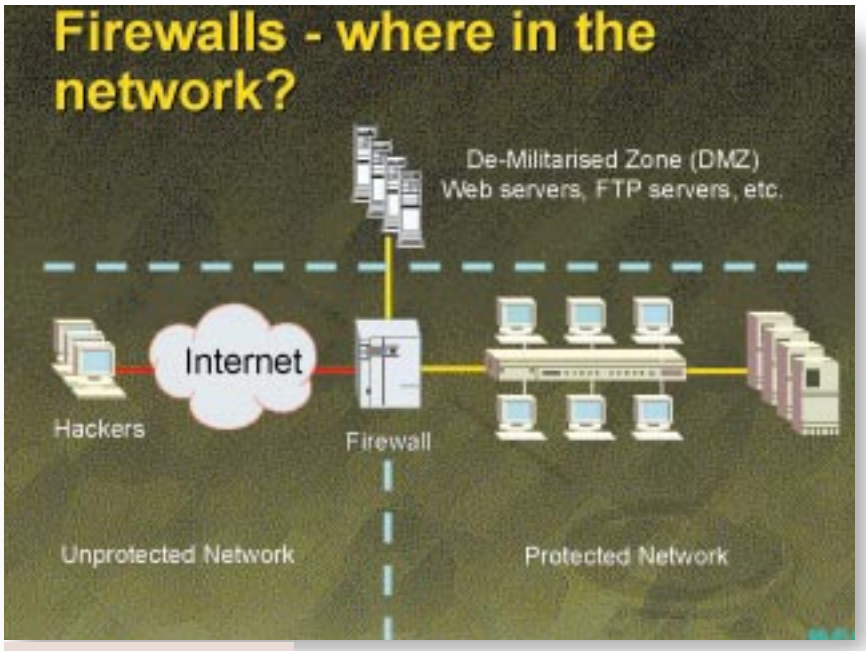

▲ KEEPING THE BAD GUYS OUT: A FIREWALL IS A NETWORK'S MEANS OF PROTECTION AGAINST DANGEROUS NET ELEMENTS
▶ ALL NET TRAFFIC CAN BE CONSTANTLY MONITORED USING A STATEFUL INSPECTION FIREWALL
▼ PROXY SERVERS ACT AS INTERMEDIARIES IN CLIENT/SERVER OPERATIONS

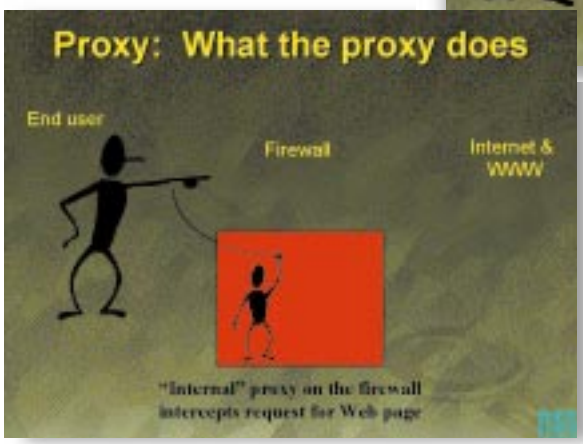

Mind the gap

So, is there anything you can do about it? Plenty, as it happens, so I thought I would set you on your way by covering the subject of firewalls: how they work, and how you use them.

Essentially, a firewall should be thought of as a gap between two networks — in our case, an internal network and the internet — occupied by a mechanism that lets only a few selected forms of traffic through. The important thing to remember is that there are three main firewall architectures currently in use.

Static Packet Filtering Working at the Network Layer (i.e. very low down) in the OSI stack, packet filters make simple deny or permit choices depending on the network address of the packet and a number of rules defined by the administrator. Packet filtering is fast, transparent (no changes are required at the client), flexible and cheap: most routers will provide packet filtering capabilities, and pure packet filter firewalls do not require powerful hardware.

Dynamic Packet Filtering/Stateful Inspection Some vendors are touting this as the "third generation" of firewall architectures, but it's really just an extension of the basic packet filtering architecture employed by most routers. Stateful Inspection occurs low down in

BEHIND THE MASQUE

Another important feature of today's firewall is Network Address Translation (NAT). This is also occasionally referred to as "IP Masquerading". Here's how NAT works. When one of the machines on your internal LAN wants to communicate with one elsewhere on the internet, it sends a packet to the firewall using its normal IP address. On its way through the firewall, the packet is altered — its return address is replaced by the IP address belonging to the firewall itself. All responses to those packets therefore come back to the firewall. The firewall then reverses the address-swapping process and passes the reply straight back to the target computer .

Why is this useful? First, the outside world never gets a glimpse of the addressing structure of your protected network. Everything looks as though it is coming from a single address — that of the firewall. This then gives rise to a useful side effect. If the internal address is never seen by the outside world, there is no need for any of the computers behind the firewall to have legitimate IP addresses. This provides complete freedom for your internal IP numbering system. There are several groups of IP addresses which are considered to be "reserved" and are never actually used on the internet, making them ideal for use in an internal numbering system protected by NAT. They are: 10.0.0.0 to 10.255.255.255; 172.16.0.0 to 172.31.255.2551; and 192.168.0.0 to 192.168.255.255

The use of NAT also means that the firewall will allow any number of machines on the LAN to share a single, legitimate IP address, thus making your internet connectivity costs that much lower.

the network stack, at the MAC or Network Layer, thus making it fast and preventing suspect packets from travelling up the protocol stack to the operating system above (never a good thing!). Unlike static packet filtering, however, Stateful Inspection makes its decisions based on all the data in the packet, corresponding to all the levels of the OSI stack.

The state of the connection is monitored at all times (hence Stateful Inspection), allowing the actions of the firewall to vary based on the administrator-defined rules and the state of previous conversations. In effect, the firewall is capable of remembering the state of each ongoing conversation across it, thus allowing it to screen all packets for unauthorised access while maintaining high security, even with connectionless protocols such as UDP.

Inevitably there are penalties to be paid for this level of security

Proxy Servers Working very high up at the Application Layer of the OSI stack, a Proxy Server firewall acts as an intermediary for user requests, setting up a second connection to the desired resource either at the application layer (an application level gateway) or at the session or transport layer (a circuit-level gateway).

Proxy code is effectively split into two halves, and actually "stands in" for both client and server operations, relaying valid requests between the trusted and untrusted networks via the proxies. For a simple client request to retrieve a web page, the proxy server fools the client into thinking that it (the proxy) is the required web server. It then passes the request to its "external half", which pretends to be the client making the request. As far as the outside world is concerned, the protected internal network does not exist — all that is visible is the external portion of the proxy. The web server passes the page back to the proxy, which transfers it to the internal proxy, which finally passes it on to the user's web browser. Unlike Packet Filter and Stateful Inspection firewalls, a direct connection is never allowed between the two networks. The penalties paid for this level of security, however, are performance — Proxy Server firewalls have large processor and memory requirements in order to support many simultaneous users, and flexibility — the introduction of new internet apps and protocols can often involve significant delays while new proxies are developed to support them.

While static packet filtering alone is usually confined to the router these days and not considered strong enough for enterprise-class firewall devices, the differences between the remaining two architectures are negligible. True proxy servers are the safest, but impose a severe overhead in heavily loaded networks. Dynamic packet filtering is definitely faster, though most high-end firewalls are hybrids these days, using elements from all three architectures.

Most of the currently available products are available as dual or tri-homed gateways: they have two or three separate network interfaces for the internal, external and "De-Militarised Zone" (DMZ) networks. The DMZ adds extra protection for the internal network, providing a secure subnet which allows internal web, FTP and mail servers to be accessed from both the trusted and untrusted sides of the firewall without compromising security. Even if an attacker on the external segment manages to compromise machines on the DMZ, everything on the inside remains guarded by the firewall.

A DMZ is the only safe way of allowing external users access to some of the servers on your site.

PCW CONTACTS

Bob Walder *can be contacted via the usual PCW editorial office (address, p10) or email* networks@pcw.vnu.co.uk.

Domain and simple

Companies often end up with a proliferation of domains, and trying to rationalise the structure can be a daunting task. Bob Walder has ways of making it easier.

Much has been written about the lack of a true directory service in NT and the problems inherent in managing multiple domains in an NT network. But if domains are that complex to administer, why do organisations so often finish up with tens or hundreds of the things spread about the place?

History lesson

There are a couple of historic reasons for this. One is that no matter what people tell you, size does matter. In Windows NT Server 3.1, domain controllers were limited to storing 10,000 objects in the security accounts manager (SAM) database. Many larger companies found this to be insufficient as the network grew and were forced into using multiple domains. With the release of Windows NT Server 4.0, the limit was increased to 40,000 users and the maximum recommended size of the database was 40Mb. But for some it was too late since the domain structure was already fixed, and for others even this number remained too small.

There are other reasons, too, of course. Sometimes a network's communications infrastructure dictates the domain structure to minimise replication across slow WAN links. Large organisations might also want to delegate administrative tasks to a number of people and the only true security boundary in an NT network is the domain. Finally, there is simple growth, whether organically within a company or via acquisition of others. Either way, it is possible to finish up with numerous domains which would be better merged together. The problem is that as things change, it is often necessary to collapse the domain structure to something simpler. When two companies are merged, for instance, rationalisation of the domain structure is always desirable. Microsoft recommends collapsing domains in preparation for the move to NT5 and Active Directory. Unfortunately, domain reconfiguration is not that straightforward because of the unique security identifier (SID) associated with each user object.

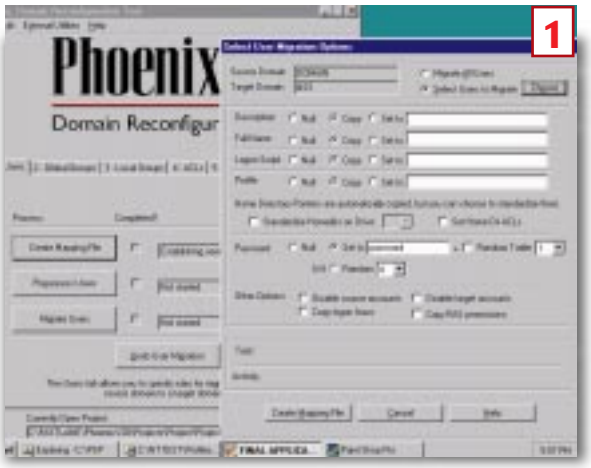

1

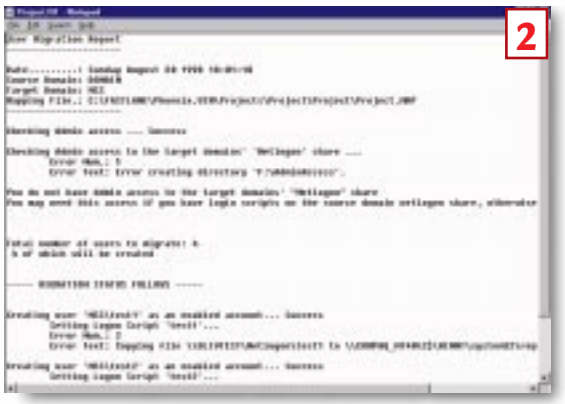
2

It is an unfortunate fact of life in the NT world that SIDs are not portable across domains. And security is completely dependent on the SID rather than the user name. It is the SID that is a member of a Local Group and ACL and User Rights, *not* the user name, which is there for display purposes and ease-of-management only. To merge two domains into one, therefore, would require manual creation of users and groups in the target domain, and manual adjustment of the various shares and access rights to ensure that these users could continue to access resources in the source domain until the migration was completed.

Rationalisation of the domain structure is always desirable

Domain reconfiguration in a large enterprise cannot be accomplished overnight, so you have to ensure that as users are moved to the new domain, they still have access to all their original shares, mailboxes, printers and other resources until the migration is finished. I recently came across a product called Phoenix, from FastLane Technologies, which helps with the whole reconfiguration process. Going through the various stages step by step I will highlight the problems and show where Phoenix may be useful. In a large network with many domains, these stages represent a huge manual effort which is not only tedious but prone to errors. Tools like Phoenix, which should probably have been included as part of the base operating system, certainly take much of the pain out of the process.

1 **Users** [Fig 1]. Assuming we already have NT Server installed on a primary domain controller in our new domain, the first task is to create the user objects that will be members of the new domain. Where we are consolidating a number of domains into one, this can be a huge task. Phoenix allows the administrator to perform this creation automatically, selecting the users from a list of those

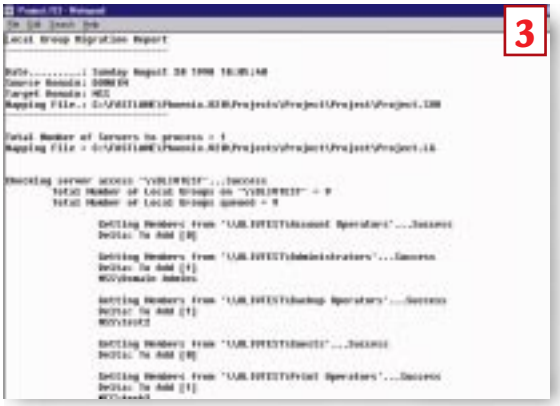

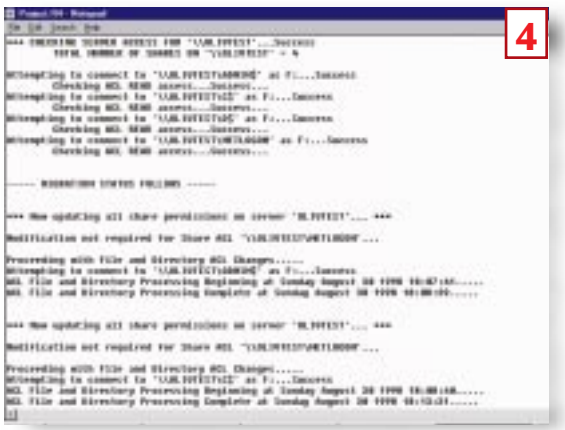

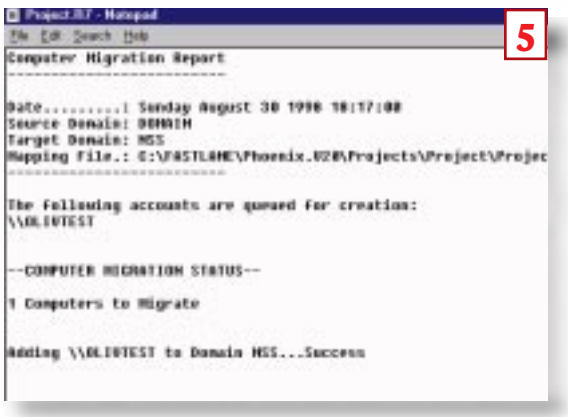

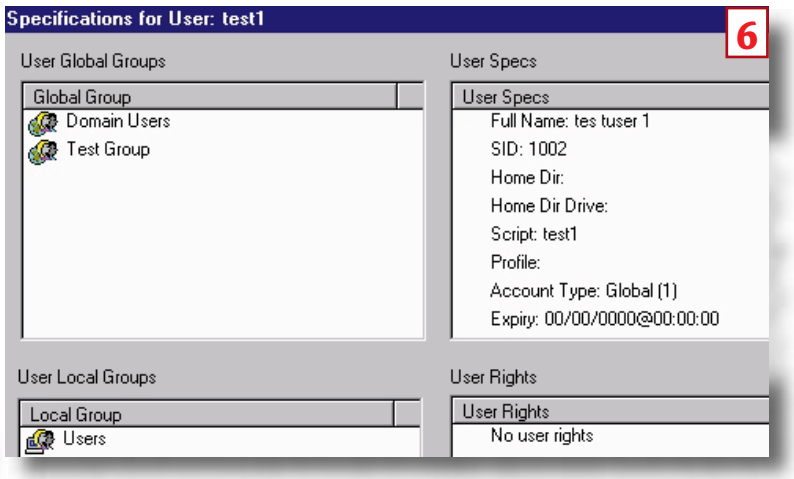

already on the existing domains. User names to be migrated can be chosen individually from a pick list, or can all be migrated in one go with the appropriate user objects duplicated (with new SIDs) in the target domain. The remaining stages focus on finding an objects source SID and inserting the SID of the destination object so that the new users maintain the same access as the original account.

2 Global Groups [Fig 2]. Like users, global groups cannot be copied across domains so new objects need to be created in the target domain to mirror the source groups, and the appropriate users must be added to the new groups. Once again, this is a tedious manual process that can be automated using Phoenix. If multiple domains are merged to one target, groups with the same name can be merged or created as separate groups. For example, you probably wouldn't want to merge all the "Domain Admin" groups into a single object in the target domain. As the new global group objects are created, they are populated with the same members that were in the source groups, using the new user objects created in the target domain in stage 1.

3 Local Groups [Fig 3]. In order that a newly-created user object has the same access to resources as the original account, Phoenix searches all local groups for the original users' SIDs and appends the SID of the target user account. Global groups in the source domain are processed in the same way, since local groups can include both users *and* global groups. As this process can be resource intensive, you can run the external application locally on each server throughout the enterprise. This can be done using the DR Distributor, which distributes a secure scheduler and an automated updater down to local computer level. The administrator can then push the updating of local group migration, ACLs and user rights to the computer itself to minimise the load on the network. The local computer spawns the update application and updates its own data locally, as a central console maintains and reports all updates.

If an organisation is implementing a staggered migration, it is possible to schedule the automated updater on the remote computers to run every day, which allows the network to quickly adapt to massive changes in stages.

4 ACLs [Fig 4]. Like local groups, all ACLs must be searched for SIDs of the original account. As each is located, the SID of the target account is appended, therefore ensuring that the target user object retains the same access rights as the original.

5 Rights [Fig 5]. Although the previous stages give the destination user access to the same global groups, local groups, shares, files and directories, the user still may be unable to log on. Such computer-specific access is determined by user rights and advanced user rights, and this stage ensures that once again destination SIDs are appended wherever source SIDs are found, to ensure that the new account has equal access to the same physical computers as the old account.

6 Computers [Fig 6]. The final stage of domain reconfiguration creates computer accounts on the target domain for all computers in the source domain. Other tools in the Phoenix package make light work of moving Domain Controllers and Exchange mailboxes between domains, too. Once the initial migration phase has been completed, Phoenix can also be used as a day-to-day domain management tool.

PCW CONTACTS

Bob Walder *can be contacted via the usual PCW editorial office (address, p10) or email* networks@pcw.vnu.co.uk

FastLane's Phoenix *is available from Peapod Distribution on 0181 606 9990*

What's in a name?

A good naming scheme, with a DNS server, is the most efficient way to provide automatic addressing and naming services on a Windows network. Bob Walder shows you how to do it.

In a recent issue, we went into some detail on IP addressing and the use of WINS and DHCP to provide automatic addressing and naming services on a Windows network. Since then, I have received several emails asking me to do a follow-up piece on DNS in Microsoft environments — so here goes.

Domain Name System is the cornerstone of the internet, since it provides the means to turn all those long-winded IP numbers into equally long winded-names — but at least they are easy to remember. For instance, how many of you know the IP address of the Novell web server off the top of your head? Try dropping to a DOS prompt and type ping.www.novell .com. After a pause, you should see a reply from 137.65.2.11. If you had typed PING 137.65.2.11 in the first place, you would have got exactly the same result, but marginally quicker. That is because the first thing that happens when you try to PING a domain name is that it must be resolved by a DNS server.

The simplest way to provide name resolution in small networks is to use the HOSTS file. This is a text-based file which can be found on most TCP/IP systems, and which contains a simple list of IP addresses and the names that relate to them [Fig 1]. This file can name common systems both inside and outside an organisation and each address can have several names, usually a "formal" name followed by a number of less formal "nicknames" or aliases. Hence, in our sample HOSTS file, the marketing server can be referred to by its IP address of 10.1.1.2, its full name of dilbert.marketing.acme.co.uk, or by its shortened aliases of Dilbert or Marketing. While the use of HOSTS files is possible in smaller networks, they can have serious drawbacks in larger ones.

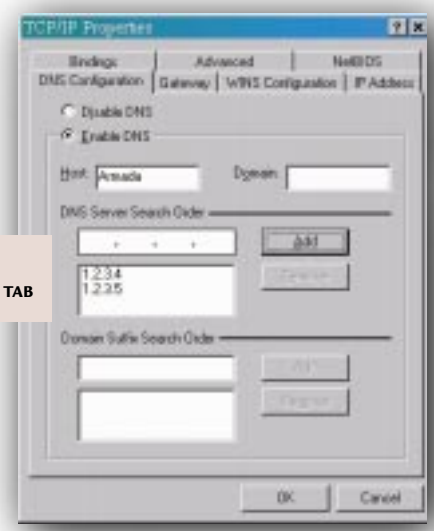

▶ **FIG 2 DNS CONFIGURATION TAB**

The main problem is that a copy of the file must exist on each and every TCP/IP client which intends to refer to resources by name rather than IP address. This approach is obviously not very scaleable

DNS provides for central control over names and IP addresses

and presents systems administrators with a potential nightmare in a network with hundreds or thousands of clients. Ensuring that each and every HOSTS file is always up to date as network changes are made is bad enough, but there is also the temptation for users to create their own files with customised naming conventions, making it difficult for "hot desking" colleagues. Of course, it is possible to manage HOSTS files by keeping a master version on one of the central servers, and downloading it to clients automatically on a regular basis, but this approach, too, can have its problems in large distributed networks.

Obviously, such solutions will not scale in large organisations and certainly will not scale to the internet. The problem of internet naming has, to date, been largely satisfied by DNS which allows a computer that is registered to the internet to be uniquely identified by that name wherever it may be located. Because computers work with numbers rather than names, the second major function of DNS is to translate the unique host name into the appropriate IP address required in order to establish communications.

The use of DNS is certainly mandatory in some form if you wish to participate in the internet. However, a good naming scheme coupled with the implementation of DNS servers can also make life considerably easier for users of a large corporate TCP/IP network, even if it is not connected to the internet. DNS provides for central control over names and IP addresses and removes the need for individual HOSTS files — although these can co-exist quite happily where required. It allows control to be applied both in a distributed fashion and where it can be most effective, in local sites and divisions. The domain name protocol is quite complex syntactically, although its operation is straightforward. A

[FIG 1]

```
; Hosts
;
; IP address    name                          alias
;
127.0.0.0       loopback                      Bob
10.1.1.2        dilbert.marketing.acme.co.uk  Dilbert Marketing
10.1.1.3        dogbert.sales.acme.co.uk      Dogbert Sales
10.1.1.5        ratbert.accounts.acme.co.uk   Ratbert Accounts
10.1.1.10       bwalder.acme.co.uk            Bob
192.168.1.52    bgates.microsoft.com          Bill
```

host, given a name, asks the server for a name-to-address translation. If the name server does not possess the means to perform that translation directly, it will pass the request on to a server with a higher authority than itself. This process can be repeated until the request is satisfied, which will always happen unless the requested address was incorrectly specified or there is some unforeseen problem, such as a DNS server being temporarily unavailable. Those of you with internet connections will have one or more entries in the DNS tab in your Network configuration [Fig 2, p299]. This will provide you with everything you need to resolve all those addresses out there on the internet.

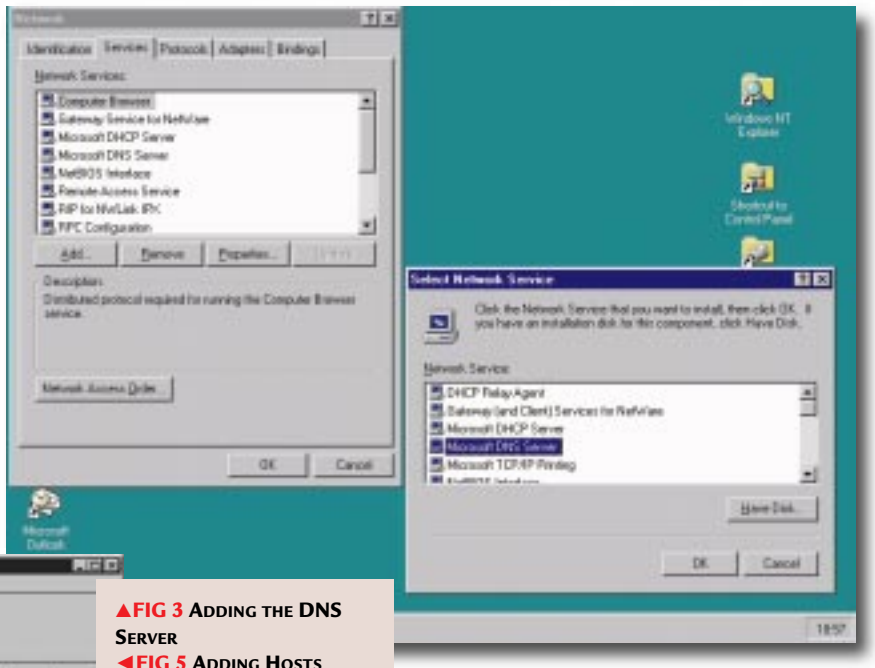

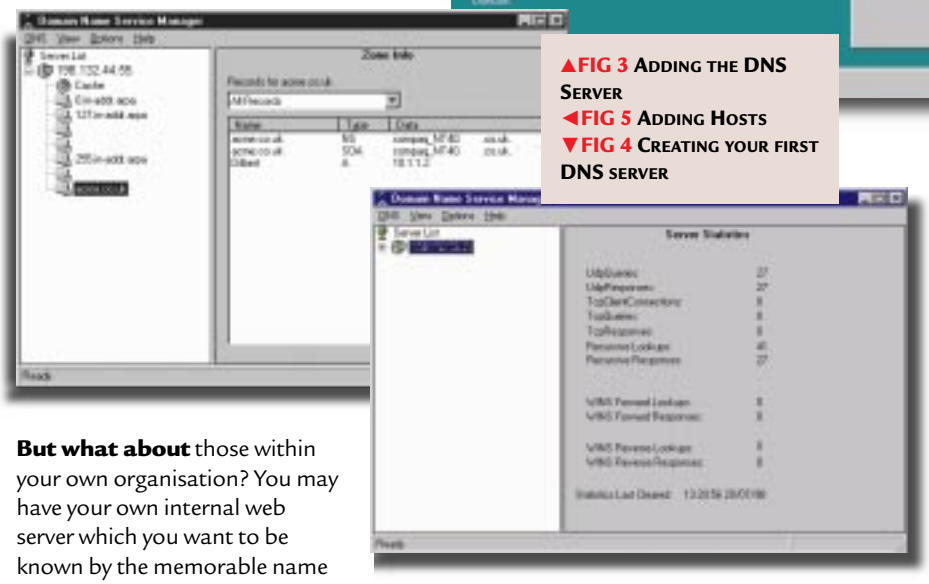

▲FIG 3 ADDING THE DNS SERVER
◀FIG 5 ADDING HOSTS
▼FIG 4 CREATING YOUR FIRST DNS SERVER

But what about those within your own organisation? You may have your own internal web server which you want to be known by the memorable name of DILBERT rather than the slightly less memorable address 10.1.1.2.

Luckily, it is easy to add your own DNS server in a Microsoft site, since the appropriate software is included with NT Server 4.0 (prior to this, you had to rely on a rather lacklustre offering in the Resource Kit for 3.51). One nice feature of the Microsoft product is that it hooks neatly into WINS if you already have WINS servers configured. DNS configuration can be quite complex, unfortunately, so I am going to concentrate on the absolute basics to provide internal naming services.

- **Call up** the Control Panel on the server and double-click on the Network icon.
- **Click Services,** Add, and select the Microsoft DNS Server [Fig 3].
- **Supply** the path to the installation files and reboot your server when finished.
- **Once the server** has rebooted, go to Start, Programs, Administrative Tools (Common) and the DNS Manager.
- **Click on DNS,** New Server and provide the IP address of your new DNS server. The database files are initialised and this address then appears in the Server List in the left-hand pane [Fig 4].

The use of DNS is mandatory in some form if you wish to participate in the net

A number of zones are automatically created and can be ignored for now. Your first job is to create a zone to represent your internal network.

- **Highlight** the DNS server you have just created and click on DNS and New Zone.
- **Click** on Primary and then the Next button.
- **Give** the Zone a name (e.g. ACME.CO.UK), tab to the Zone file field and click on the Next button to accept the default file name.
- **Click** on the Finish button. The new Zone is created with some default entries.

All that is left is to enter the host names you wish to resolve. Taking our DILBERT example in the Domain ACME.CO.UK, we would simply highlight the Zone ACME.CO.UK and click on New Host. Enter the host name (DILBERT) and the address (10.1.1.2), and click on Add Host. Your Domain should now look like the right-hand pane [Fig 5]. Now, include the address of your DNS server in the DNS configuration tab of all your clients (first in the list) and you're ready to go.

Refer to my advice on the use of DHCP in an earlier issue to see how you can make that change across all your clients in a matter of seconds. Now if you drop to the DOS prompt and type PING DILBERT.ACME.CO.UK you should receive a reply from device 10.1.1.2 — and away you go.

PCW CONTACTS

Bob Walder *can be contacted via the PCW editorial office (address, p10) or email* networks@ pcw.co.uk

Ghost story

No, not things that go bump in the night, but a utility that automates network installation. Bob Walder investigates.

There are still occasions when you may need to have a set of network client disks for your PC in order to attach to the network without booting into Windows 95 first. Leaving aside the possibilities of corrupt Windows installations which might force such drastic moves, I can cite one practical example of my own.

When I create new machines for the test lab, I use a nifty utility called Ghost <www.ghostsoft.com>, which is designed to automate the process of installing Windows 95, NT and OS/2. The idea is that you create a "standard installation" on one PC and then create an image of the hard drive in a special "Ghost file" that can then be written to any number of other PCs at the click of a mouse.

The spirit of DOS

The catch is that you need to Ghost from partition to partition, or from disk to disk, and when creating a client you need to run Ghost from DOS rather than Windows. The easiest place for me to store the standard Ghost files is therefore on my network drives, but this necessitates being able to boot to DOS and attach to the network in order to access those files.

This is done using the Network Client Administrator, a utility which is installed via the standard NT Set-up routine. Once installed, you must log in as Administrator (or equivalent) and run through the following procedure:

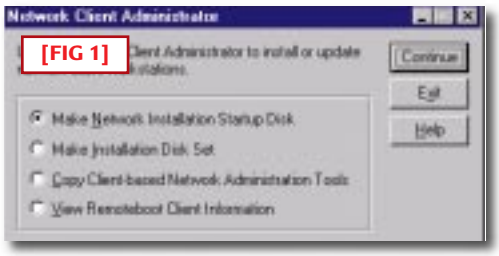

1 Click Start, Program, Administrative Tools (Common), Network Client Administrator to fire up the utility. You are presented with the window in [Fig 1].

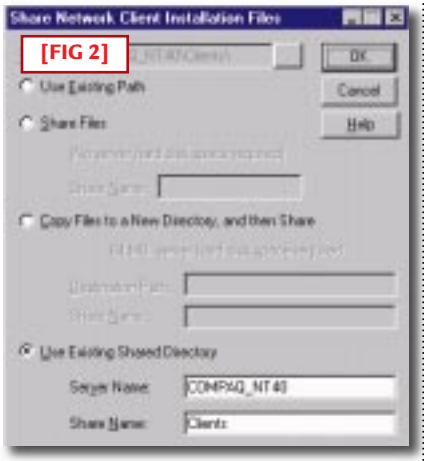

2 Making a full installation disk set is a pain, and the available client options for that are limited, so we will make a Network Installation Start-up Disk. Click on the appropriate option and then Continue. The difference with this option is that the resulting disk merely boots and attaches to the network, following which it automatically runs a full client installation from the server to the local hard drive.

3 The next window looks like [Fig 2]. If this is the first server to host the utility, you can allow it to copy the client installation files to the local server hard drive and create a share called "Clients", simply by entering the source path (the installation CD) and selecting "Copy files to new directory, and then share". Alternatively you can leave them where they are on the CD-ROM by selecting the "Share Files" option, although the CD must then be available each time you perform a client installation.

4 After the files have been copied to a server, you will simply enter the server name and the share name (usually "Clients") where the files were copied to, and select the "Use existing shared directory" [Fig 2].

5 Select the floppy drive type, client and network card from the next window. We are creating disks containing the Network Client for MS-DOS and Windows but you might find you have a problem with the network card. There are not many to choose from in this menu, which is strange given the range of devices supported within NT.

If your network card does happen to be among those listed, congratulations, because you've got an easy life and the rest of the process is plain sailing. For most of us, however, we will have something else installed in our client machine and so we have a bit of jiggery-pokery to perform. To begin with, just select any old card — I chose the 3Com Etherlink III — and we will fix the configuration files later.

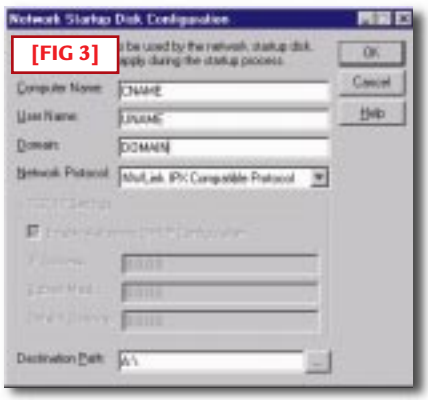

6 The next window [Fig 3] is where we specify the unique computer name we want to assign the client machine, any user name which has access to the NT Server where the client files are located (best to make this one Administrator equivalent), the name of the domain to authenticate to, and the network protocol to use. You may only have one available, so the TCP/IP settings may or may not be applicable. If you select TCP/IP you can then specify whether or not to get IP info from a DHCP server. If not, you can enter the IP address, subnet mask and default gateway yourself. Either take a look at my previous column on IP addressing, or stick to IPX/SPX to keep things simple.

The final bit to enter here is the destination path, which is usually your floppy drive.

[FIG 4]

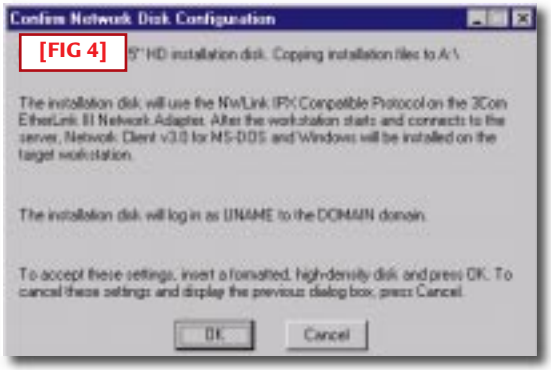

7 The confirmation screen is at [Fig 4]. Insert a floppy disk which has been formatted as a system disk (i.e. format a: /s) but is otherwise blank. Clicking OK will start the file-copying procedure

Once all the files have been copied, you have a disk that contains the files COMMAND.COM, CONFIG.SYS, AUTOEXEC.BAT and a directory called NET. The important stuff is all in the NET subdirectory, which contains the basic client software, configuration files and network card driver.

If you selected the correct card in step 5, then you are home free at this point: just stick the floppy disk in your client PC, reboot and away you go. For the rest of us, it's time for that jiggery-pokery I mentioned earlier.

The first step is to put the correct driver on the floppy disk. This is difficult to specify exactly, since every vendor constructs its driver floppies differently. In essence, however, you are looking for a subdirectory called NDIS, or perhaps DOS, which will be somewhere on the driver disk that came with your network card. In that directory will be a file with the .DOS suffix. In my case, the driver for the 3Com Fast Ethernet XL 10/100 PCI network card is called EL90X.DOS, so this file should be copied to the NET subdirectory. Just to be on the safe side, why not delete the driver that is there currently, which is called ELNK3.DOS.

OK, we now have the correct driver on the disk, so the next job is to amend the configuration files to point to it. Look in PROTOCOL.INI [Fig 5] for the line which says DRIVERNAME=ELNK3$ and change it to reflect the name of your new driver. Leave out the .DOS suffix, but make sure you leave the $ on the end. In my case the line will now read DRIVERNAME=EL90X$.

A similar operation must also be performed with the SYSTEM.INI file [Fig 6]. Here you are looking for the [network drivers] section, for a line which reads

```
NETCARD=ELNK3.DOS.
```

This time, you are going to replace the entire filename, so in my case the new line reads

```
NETCARD=EL90X.DOS.
```

It's not difficult to do, but it's just not that obvious if you're not used to it. All you need to do now is insert the floppy disk into the client workstation and reboot. OK, I lied. Because what actually happens, if you take a look at the AUTOEXEC.BAT file [Fig 7], is that the client software is loaded, the PC attaches to the network and then installs the full client from scratch on your local hard drive (that's what the z:\msclient\netsetup\setup.exe /$ does).

[FIG 5] PROTOCOL.INI

```
[network.setup]
version=0x3110
netcard=ms$elnk3,1,MS$ELNK3,1
transport=ms$ndishlp,MS$NDISHLP
transport=ms$nwlink,MS$NWLINK
lana0=ms$elnk3,1,ms$nwlink
lana1=ms$elnk3,1,ms$ndishlp

[ms$elnk3]
DRIVERNAME=ELNK3$
; IOADDRESS=0x300
; SLOT=1
; MAXTRANSMITS=6

[protman]
drivername=PROTMAN$
PRIORITY=MS$NDISHLP

[MS$NDISHLP]
drivername=ndishlp$
BINDINGS=ms$elnk3

[ms$nwlink]
drivername=nwlink$
FRAME=Ethernet_802.2
BINDINGS=ms$elnk3
LANABASE=0
```

[FIG 6] SYSTEM.INI

```
[network]
filesharing=no
printsharing=no
autologon=yes
computername=COMPAQ_NT40
lanroot=A:\NET
username=newuser
workgroup=domain
reconnect=no
dospophotkey=N
lmlogon=0
logondomain=domain
preferredredir=full
autostart=full
maxconnections=8

[network drivers]
netcard=elnk3.dos
transport=ndishlp.sys
devdir=A:\NET
LoadRMDrivers=yes

[Password Lists]
```

[FIG 7] AUTOEXEC.BAT

```
path=a:\net
a:\net\net initialize
a:\net\nwlink
a:\net\net start
net use z:
\\COMPAQ_NT40\Clients
echo Running Set-up...
z:\msclient\netsetup\
setup.exe /$
```

Of course, that means you have to go back and redo the changes you have just performed on the new installation on the hard drive. To be honest, I usually remove that last line of the AUTOEXEC.BAT and use the floppy disk itself to boot and attach to the network, replacing the NET USE Z: mapping to point to the shares I actually want to use. For instance, if you share the C: drive of your server as DRIVEC, you can replace the above command with:

```
net use z:
\\COMPAQ_NT40\DRIVEC
```

PCW CONTACTS

Bob Walder *can be contacted via the PCW editorial office (address, p10) or email* networks@pcw.co.uk

A new flame

Bob Walder reports on his vist to the Networld+Interop show in Las Vegas where security products were all the rage, with network security do-it-alls and new versions of firewalls.

I have just returned from the Networld+Interop show in Las Vegas, where one of the most common recurring themes was security. Every other stand at the show seemed to sport some kind of security product, aimed either at keeping intruders out, or telling you all about what an intruder is up to once he has got in to your system.

Today, even the smallest company is connecting one or more PCs to the internet for browsing and email, and so needs to be aware of the implications of opening up its sensitive business data to the unwelcome attentions of strangers. So, the abundance of security-related products is hardly surprising, given the popularity of the net, and security is a subject which I intend to cover in more detail in future columns.

Flame throwers

A number of companies announced new versions of their firewall products, including the likes of Checkpoint and WatchGuard. Other vendors more renowned for their hardware expertise, such as Matrox and Bay ,were extolling the security features in their latest products.

Some companies, like Abirnet, are already saying that firewalls are redundant. Abirnet produces SessionWall, a network monitoring and blocking tool providing the ability to monitor your TCP/IP traffic across the full range of protocols: HTTP, FTP, SMTP, POP3, UDP etc. You can selectively ignore, report, log, alert or even block traffic, depending on combinations of the source and destination addresses, protocol and contents of the data packet.

According to Abirnet's Kurt Ziegler, firewalls only provide a defence at your "front door", doing nothing to protect you from internal attacks (which make up 80 percent of all hacking incidents). Tools like SessionWall, though, provide a constant vigil against internal *and* external attacks, and thus a more effective overall defence.

Although I take Kurt's point to a certain extent, I do not foresee the end of the firewall just yet. After all, take house security: it is one thing having security cameras at home so you can identify a burglar; but surely it is a wise precaution to lock the front door, too? In effect, I think most companies will want to use firewalls *in conjunction* with security monitoring tools such as SessionWall and ISS RealSecure to provide a true belt-and-braces approach to network security.

As firewalls move from the category of

"propeller-head" to "commodity", the smart vendors are already recognising the need to further differentiate their products. This is achieved by layering additional services on top of the firewall, so that critical processes which are best performed at the corporate gateway to the internet, such as virus scanning or bandwidth control, can be executed by the firewall box. According to Checkpoint, the customer wants tight integration of all security and traffic control components, together with policy-based management for all business.

Abirnet claims that firewalls are redundant and that SessionWall provides a constant vigil against internal and external attacks

Two of the most crucial services being developed by savvy firewall vendors include directory services support and bandwidth control. Checkpoint's FireWall-1 is one of the first firewalls to offer integrated LDAP support with the release of version 4.0, announced at the show. This simplifies user management within security policies, as it allows FireWall-1 to share user and group repositories already established within an organisation.

General management and ease of configuration is also improved by virtue of the new Java-based GUI for defining and managing user-level security information and flashy new Security Wizards to guide the inexperienced administrator through the process of defining a security policy which does not incorporate glaring holes.

Checkpoint Charlie

Taking a step towards the day when we can use a single console to manage all our security devices, Checkpoint's Open Security Manager now allows the administrator to define, manage and distribute security policies for a range of third-party security devices such as Cisco PIX, Microsoft Steelhead, and routers from Bay Networks, Cisco and 3Com.

For those wanting to move beyond simply connecting their network to the internet and into the realms of providing secure inter-site links across an otherwise insecure network, then FireWall-1 4.0 also provides some smart new VPN (Virtual Private Network) facilities. Both the firewall itself and the standalone remote client (SecuRemote 4.0) support a Public Key Infrastructure (PKI) as part of an overall enterprise security implementation. This makes use of X.509 digital certificates and Certificate Authority (CA) technology from

Entrust to simplify and automate critical VPN functions like adding and deleting users, managing encryption keys and providing encryption key backup/recovery.

Rising from the Chrysalis

That other great bugbear of encryption applications, performance (or lack thereof) has been addressed through a partnership with Chrysalis-ITS. The Chrysalis plug-and-play PCI card offers hardware-based DES and Triple-DES encryption, allowing firms to deploy VPNs at Fast Ethernet speeds without performance degradation. These cards generate and store private encryption keys on the card itself, guaranteeing the privacy of all communications to and from the FireWall-1 gateway.

Be on your WatchGuard

WatchGuard Technologies launched its new baby at the show. Called Firebox II, this is a complete turnkey hardware and software solution which WatchGuard has christened a "network security appliance". It is an attempt to place the Firebox alongside the fridge or microwave in terms of plug-and-play and in truth it is not far off that ideal, at least in terms of installation.

As ever, defining an effective security policy is anything but trivial in terms of the amount of thought required, but at least the

p283 ➢

WatchGuard Technologies' new baby, Firebox II, is an attempt to place the Firebox alongside the fridge or microwave in terms of plug-and-play

means of getting that security policy into the Firebox is made as straightforward as possible. And once it is there, WatchGuard provides an excellent range of graphical tools to monitor and report on activity through the firewall.

For those organisations which have multiple sites, or which provide managed services across the internet (such as ISPs), the Firebox provides a Global Security Manager through which an administrator can initialise, configure, manage and update multiple Fireboxes. Dual-image flash memory allows an administrator to download a new policy, monitor it to test its effectiveness and immediately revert to the former policy should anything prove unsatisfactory. For those wanting to provide secure links between sites, Firebox also includes IPSec VPN as standard.

Hitting the iSwitch

On the hardware front, both Matrox and Bay had something new to offer. Matrox was announcing its Fast iSwitch product, designed to provide high-performance switching, internet access, firewall, DMZ (De-Militarised Zone) and management software in a single, diminutive unit ideal for the smaller office environment.

This device features two, four or eight 10/100 auto-negotiated ethernet ports allowing it to be connected to the internal network, either directly to power users or servers, or to existing hubs or switches. External communications are provided by two enhanced serial ports for modems or ISDN terminal adapters, which are then available to all network devices attached to the iSwitch. The Fast iSwitch also performs bandwidth augmentation between the two serial devices, providing double the bandwidth when required.

Firewall is pretty basic, but provides a DMZ facility to host web and FTP servers internally, and Network Address Translation (NAT) to hide your internal network structure from prying eyes. The beauty of the iSwitch is that it allows a low-cost means of connecting an existing LAN, or group of standalone PCs, to the internet without the need for a dedicated server or router.

The Bay area

Bay's Extranet Switch takes a slightly different tack but is still aimed at small to medium-sized enterprises wanting to create secure extranets and personal VPNs (Virtual Private Networks) across the net. The model launched at the show was the Extranet Switch 1000, aimed at businesses with up to 50 simultaneous remote users and priced at $7,000 (UK price to follow).

Since Bay considers that the traditional remote access infrastructure is fraught with problems (long-distance phone calls for users, managing high fan-out of WAN connections at head office, etc.) the Extranet is designed to provide secure remote access across the net for up to 50 roving employees (up to 2,000 users with the Extranet Switch 4000) using a single box.

Directory enquiries

Directory Enabled Networking facilities allow authentication, performance levels, type of access and level of security to be defined on a per-user basis, with user information gleaned from existing corporate directories such as LDAP, RADIUS or NDS. The 1000 is designed to process VPN tunnels and perform encryption, and puts a filter on all 50 users simultaneously. Internal and external bandwidth management mechanisms ensure that the highest-priority users receive the best performance.

Tunnelling standards supported include IPSEC, PPTP, L2F and the new L2TP standard, while authentication can be performed against a variety of external servers including LDAP, RADIUS, Windows NT (with Active Directory to come), NetWare NDS, Security Dynamics or AXENT.

p284 ➢

Encryption support includes RC4, DES and Triple-DES.

The Networld+Interop show attracted around 60,000 visitors over three days, so while it is not in the same league as Comdex or Cebit it is well worth attending if you want to get an indication of how the networking market is moving a few months before it all starts to happen here.

More answers than questions: readers write about splitter cards

This month's Q&A section is all answers! I have had several responses from readers about the splitter cards mentioned by Dr Evans (*PCW June*). He was asking if a splitter card would be a viable option to purchasing and networking an additional PC to provide access to programs and data for two employees.

David Waller writes: "*I have used such a device, called Sharedware. It consists of an ISA card and a module which allows a second keyboard, mouse and monitor to be connected to a computer running Windows 95. It makes use of 95's multitasking capabilities to allow both employees to use the computer independently, and the user profiles allow the various people to have individual desktops, as if using a network.*

"*I first tested it on an older and slower 100MHz Pentium with a dial-up modem connection to our internet provider, as I felt that this would be one of the main uses of Sharedware in a school situation. With one account, one modem but two independent users, the cost of purchasing Sharedware was recouped with just one application!*

"*It worked perfectly. The users of the host and client machines could independently surf the net. It was difficult to make any quantitative conclusions on the speed deterioration because of the many variables involved, but the users did not seem to be waiting much longer than usual.*

"*I then tested it on a networked computer with an AMD K6 processor and 32Mb of RAM. The two users could independently log in to our Novell 3.12 network and access the software on the local machine, and the files in their home directories on the file server.*

"*The speed deterioration of the processing was as you would expect with any multitasking operations in Windows 95. If both users were using Word or Excel, there was no noticeable difference; but with graphics applications and DTP, there was (as expected). This is again a subjective conclusion, as no timings were taken.*

"*The price of the package is £200 but a workstation bundle is now available for £300 which includes Sharedware, monitor, keyboard and mouse, with a 30-day money-back guarantee. In conclusion, I can only say that Sharedware delivers what it promises, allowing two users to access the same computer with minimum hassle.*"

A slightly different viewpoint came from Andrew Pickup: "*I am head of IT in a preparatory school and we recently became connected, via a modem, to the internet. By chance I was sent a brochure about one of these splitters and so we bought it: £300 for card, mouse, keyboard and monitor.*

"*I believe the idea behind it is that a computer is only usually called on to use a fraction of its processing power. Theoretically, then, this power can be shared effectively to run two different 'computers' from one machine.*

"*There were some problems at first: the keyboard and mouse were poor quality, so both were replaced. The card kept trashing the computer and one day the hard disk was reformatted three times in an attempt to provide a 'clean' installation. Eventually one of the card's designers visited and declared that it was incompatible with the sound-card driver. Since then, with no sound card, it has worked fine.*

"*We are glad we bought it. It has provided double the internet use for the same phone charge, with no perceptible loss of speed, although there have been a few little quirks regarding the order of connection, the sending of emails, etc. It certainly runs the database effectively and two concurrent searches are just as quick as when only one person is using the database. A pupil using CorelDraw on one monitor happily co-exists with another who is browsing the net, and there is no problem with printing. There are crashes, perhaps once a day, but the machines are used intensively so it's not surprising. By the way, the computer is a K6 with 64Mb memory, but it did run with 32Mb.*

"*My advice to Dr Evans, though, would be similar to yours: buy another PC. If Access works OK on the machine he has at present, he could easily network the two using a crossover cable. A new computer and network cards would only cost twice as much as a splitter and monitor.*

"*Bearing in mind the importance of the information, I would not risk a splitter. After all, it doesn't really matter if* my *machine crashes, but we would not use one for our office administration. Dr Evans could get a nice setup for another £800.*"

So, anyone looking for a definitive answer will be disappointed yet again (isn't that just like a consultant?). However, as you can see, the splitter device is clearly suitable in certain circumstances, while for business-critical situations I would still recommend purchasing and networking a second PC if the budget will stretch to it.

My thanks to both David and Andrew for taking the trouble to write in about this.

Idiot proof: NetWare Connect

How big an idiot am I (answers on a postcard, please…)? In the July column, I published a request for information about modem-sharing products for NetWare. I was racking my brains so hard trying to think of third-party products that would do the job, that I completely overlooked Novell's own offering, NetWare Connect.

Karl Dyson, among others, wrote in to point out the error of my ways: "*There is a product available from Novell, called NetWare Connect, which allows you to share modems via a 'client' that effectively redirects any of your local COM ports to server-shared modems. It uses the NetWare security to allow different users different access times and levels, and can support both dial-in and dial-out. There are 32-bit clients for Windows 95, and 16-bit clients available for both DOS and Windows 3.1x, so this should answer Peter's* [Peter Williams] *problem.*

"*NetWare Connect also supports as many modems as you can fit in the server, within reason. You can plug modems into the server's standard COM ports, or purchase something like a DigiBoard — supporting four or eight modems per card, and taking the COM-port processing away from the server CPU. It runs as an NLM and is therefore configurable from a remote console as well.*

"*There should be something about it on the Novell site. Try* www.novell.com/catalog, *there is an entry for NetWare Connect 2.*"

PCW Contacts

Bob Walder can be contacted via the *PCW* editorial office (address, p10) or by email at networks@pcw.co.uk.

Abirnet www.abirnet.com
Bay Networks www.baynetworks.com
Checkpoint www.checkpoint.com
Matrox www.matrox.com
Sharedware www.sharedware.com
WatchGuard Technologies www.watchguard.com

Suits you, sir

How to achieve perfect connectivity from PDA to office network and back. Bob Walder has at last found a solution which almost fits the bill, and you don't need the special trousers.

Having decided to use Microsoft Exchange Server in-house for all messaging and groupware needs, I decided to go the whole hog and use Outlook 98 as the main email client.

I must digress for a moment here to explain that I have no great liking for Microsoft products in particular, neither do I have any great dislike for them. As with any software on the market today there are a number of features which I love, some which drive me to distraction, and even a few bugs to spice up my mundane daily routine.

Berating Microsoft-bashing

Microsoft comes in for a lot of stick from all quarters. Some of it is justified, but a lot of it is not. The company has its faults, as does its software, but "Microsoft bashing" seems to be in fashion at the moment and I don't believe that all the coverage you see is entirely objective.

Anyhow, whether or not you like the monopolistic hold that Microsoft has on the desktop market, the one thing the company provides is a wide range of software that you can be pretty sure will work together (most of the time). Hence there is a great deal of logic in a network manager deciding that with NT at the server and Windows 98 at the desktop, a sensible choice would be the equivalent Microsoft messaging products. This should never be the only criteria you use to select software, but when all other things are equal, it cannot harm to go for a one-stop-shop approach, can it?

My own messaging needs are not particularly demanding, either at the server or the client, so Exchange and Outlook fit the bill nicely. All my calendar and contact information resides on the network, allowing colleagues to schedule meetings, check my diary for appointments and make tentative appointments on my behalf.

Connectivity conundrum

The problem comes when I return to the office with my PDA, on which I have made some changes to contact records and entered a few appointments. I need a means whereby I can quickly and easily synchronise my calendar and contact data bi-directionally, so both the network and my PDA reflect the latest state of play.

Regular readers of this column will know of my fondness for the Psion machines with their excellent keyboards, superb screens and intuitive software. But to this day, what Psion has never got right is the connectivity software, so I eagerly awaited the arrival of PsiWin 2.1 since this promised me, at last, synchronisation with Outlook both for calendar and contacts data.

Unfortunately, it just doesn't work. The documentation is terrible — there are no clear instructions on how to utilise anything other than the default Outlook folders. Even when I gave in and used the default folders, it still did not provide true bi-directional synchronisation.

Compare this with Windows CE 2.0, another example of the logic behind selecting a single software vendor. CE 2.0's Active Sync technology looked like providing me with everything I needed. I acquired one of the new Casio Cassiopeia machines, plugged it in and watched as my Outlook data was synchronised: calendar, contacts, tasks, the lot. From then on, any changes I made either to the Cassiopeia or to my PC-based data was instantly synchronised, meaning the machine was always ready to go. Unfortunately, each time I came back, I noticed that the sync operation was too damn slow.

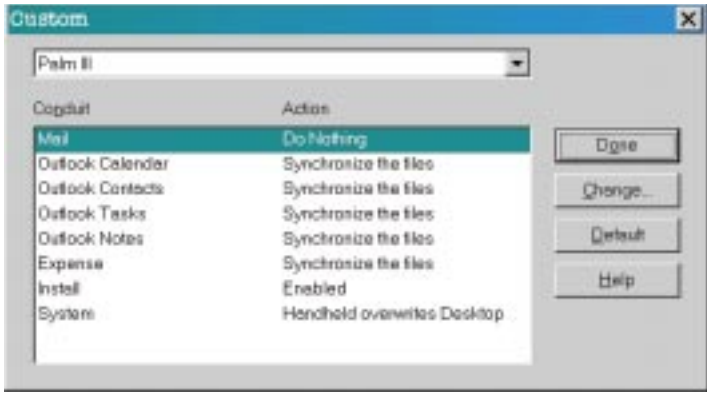

Palm III HotSync options

What a carry on! Nurse, the trousers...

I quickly tired of carrying the Cassiopeia around. It is one of the more portable of the new CE 2.0 machines but is still more like a small laptop than a true PDA. In the words of the old Monty Python sketch: "…of course it's pocketable — you just need the special trousers!"

So, a bit of lateral thinking was in order and what I came up with was the new 3Com Palm III, the latest generation of the

p295 ➢

Setting Date Book options in Desktop To Go. Note the various synchronisation options

PalmPilot PDA. Of course, it has no keyboard, but after a couple of hours' practise I found that I could write almost as fast on the Palm III as I could with pen on paper, and with 98 percent accuracy.

Perfection to go

The on-board software is adequate (although all PDA vendors should really try to emulate the Psion calendar) and you can add any number of interesting applications that can be downloaded from the web. The piece of magic which made it perfect for me as a PDA, though, was some third-party software called Desktop To Go, from DataViz <www.dataviz.com>.

This integrates completely with the Palm III HotSync utility and fully supports both Outlook 97 and 98. With this, you can choose whether or not to synchronise mail, calendar, contacts, notes and tasks, and you are not restricted to using the default Outlook folders. You can choose in which direction the updates take place, and when you opt for bi-directional, the merging is complete and accurate. If you change different fields on the same contact record in the PC and the Palm III, for instance, the two changed records are merged into a single new one on both devices.

The HotSync ring of confidence

I have never yet managed to upset it and I now have complete confidence in the process. And, this thing is *really* fast. I can press the HotSync button on the Palm III cradle and see almost 800 contacts, 200 diary entries and a hundred-or-so notes. Tasks are completely synchronised in less than a couple of minutes; far quicker than CE 2.0 can manage the same operation.

At the end of the day though, it is still a compromise. If I want to do a lot of writing, then I cannot do it on the Palm III. I need something with a keyboard. My preference would still be the Psion Series 5, although if I needed my diary information I would be forced to carry the Palm III as well. Or perhaps I could synchronise with a CE 2.0 machine instead and just carry that?

But if I need cellular communications, my choice is further complicated. I can connect my Nokia 8110 phone to the Palm III, the Psion or the Cassiopeia, but at some cost to both bulk and battery life. No, if I need cellular data and fax I switch to yet another device, the Nokia 9000i. But here again, the synchronisation options are abysmal.

What I want (what I *really, really* want) is something the size and weight of the Palm III with the on-board software of the Psion, the synchronisation capabilities of DataViz and the cellular communications facilities of the Nokia 9000i. I want the backlit mono screen of the Psion 3C, the keyboard of the Psion Series 5 and the whole package should run for days on a couple of penlight alkaline batteries. Is that too much to ask? Answers on a postcard, please.

Which Windows 95 is installed?

I have received another suggestion about how to determine which version of Windows 95 is installed on your system. This one that came from Roy Venkatesh.

"I missed your question in February's issue of *PCW* regarding the Pan-European version of Windows 95. Anyway, I am going to try and guess that you wanted to know how to determine which version of Windows 95 you had on your system.

"There is a useful page on the Microsoft web site at http://support.microsoft.com/support/kb/articles/q158/2/38.asp. In essence, it reads: 'This article describes how to identify the following information about your Windows 95 installation:

- Which version of Windows 95 you are running.
- Whether your installation of Windows 95 is

p296 ➢

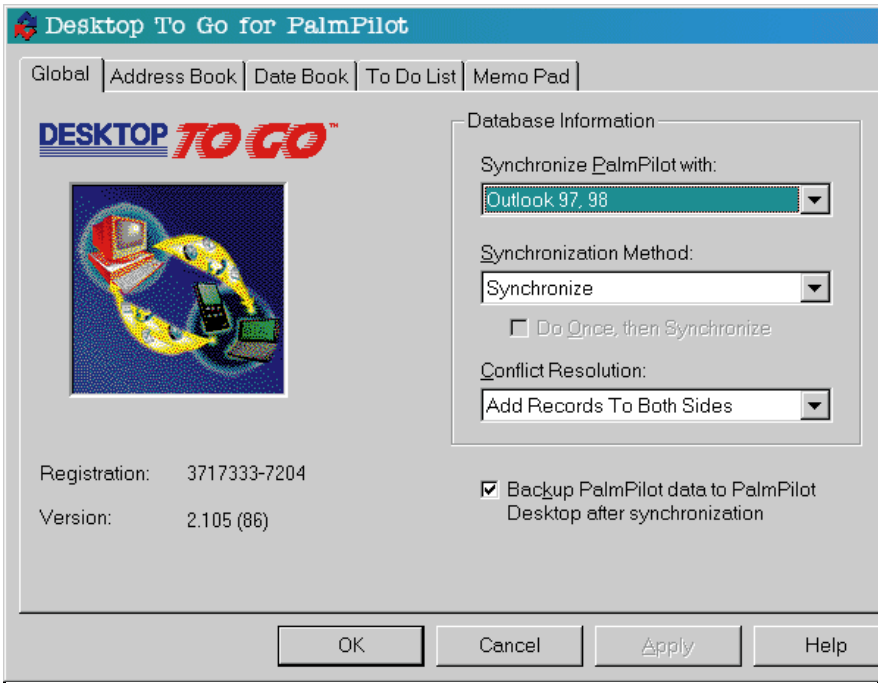

The Desktop To Go options screen

an OEM installation.

• Whether your installation of Windows 95 is an international version.

"'To determine which version of Windows 95 you are running, follow these steps:

1. In Control Panel, double-click System.
2. Click the General tab.
3. Locate the version number under the System heading and then see the following table [Fig 1].

"'*Notes*: note that if you are running OEM Service Release version 2.1, you see the version number 4.00.950B (the same as OSR2) when you follow the steps above. To determine whether you are running OSR 2.1, check for USB Supplement to OSR2 in the list of installed programs in the Add/Remove Programs tool in Control Panel and check for version 4.03.1212 of the Ntkern.vxd file in the Windows\System\

Fig 1: Find the version

Version number	Version of Windows 95
4.00.950	Windows 95
4.00.950A	Windows 95 plus the Service Pack 1 Update, or OEM Service Release 1
4.00.950B	OEM Service Release 2 (OSR2)
4.00.950C	OEM Service Release 2.5 (OSR2.5)

Vmm32 folder. If you are running OSR2.5 and you uninstall the USB Supplement using the Add/Remove Programs tool in Control Panel, the version number changes to 4.00.950b on the General tab in System properties.

"'Windows 95 may have been preinstalled on your computer. These installations are referred to as OEM (Original Equipment Manufacturer) installations. To decide if you have an OEM installation of Windows 95, follow these steps:

1. In Control Panel, double-click System.
2. Click the General tab.
3. Locate the Product ID number under the Registered To heading. This number typically contains 20 digits. If digits 7, 8 and 9 contain the letters OEM, you have an OEM installation of Windows 95. For example, the following sample Product ID number indicates an OEM installation: 12345-OEM-6789098-76543.

"'If you are using an OEM installation of Windows 95, you should contact your computer's manufacturer for Windows 95 support. To determine the language version of Windows 95, follow these steps:

1. Click the Start button, point to Find, and then click Files Or Folders.
2. In the Named box, type 'winver.exe' (without quotation marks) and then click Find Now [Fig 2].

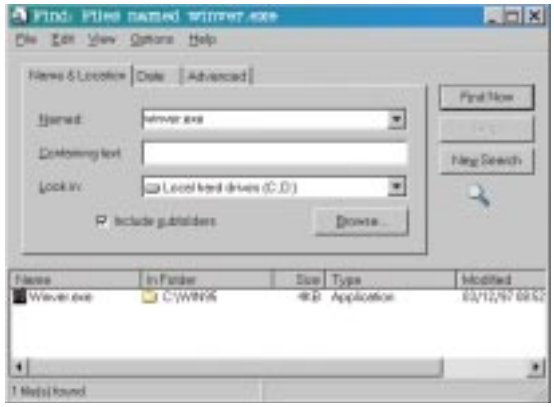

Fig 2 Result of the Find File operation on WINVER.EXE

Question & Answers: home-based web server

Q I have a small home-based business for which I would like to set up a web server and email server. We have a very small budget (almost none) and old computer systems (100MHz 486) to work with. Is a home setup a reasonable undertaking, or should I simply pay for these services to be hosted by another company?

James England

A This pretty much depends on what sort of services you want to offer and the amount of bandwidth you are going to need. As far as the software is concerned, first of all take a look at my book review in last month's column and get yourself a copy of *Red Hat Linux Unleashed*. This includes a copy of the Red Hat Linux operating system, which will run quite happily on your older 486 machines with as little as 4Mb of RAM. The CD also includes all the software you need to host DNS, mail and web services on your Linux server, so that should get you started for minimal cost.

The bandwidth is a different matter, though. Running your own web and email server requires 24-hour internet access, so unless your business already requires such access for other reasons, you will need to approach a solid, reliable ISP to get yourself a permanent 128Kbps link. This is the minimum you will require and it doesn't come cheap.

The alternative, therefore, is to look to the same ISP to register your domain name, host your web site for you and handle all your email. The things to watch out for here are that you have enough email accounts available, that there is enough disk space allocated for your web site and that it is possible to run CGI scripts on the web server if necessary, since all of these will cost extra if you don't get them included from the start.

Review: Server Resource Kit 4.0, Supplement 2

■ **Microsoft Windows NT Server Resource Kit Version 4.0, Supplement Two**
Price £46.99
Author Microsoft
Publisher Microsoft Press

This is something of a cross between a book review and a product review. I am looking at the latest supplement to the Microsoft Windows NT Server Resource Kit, to which I shall refer hereafter as "Supplement Two", given the length of its official title.

There is actually no book to be found in this package, although you can buy it from a bookshop. Supplement Two is in fact a two-CD set, expanding on the information in the original Resource Kit and its first Supplement.

For those who are unfamiliar with the Microsoft Resource Kits, they are invaluable tools for power users and network administrators. Occupying a niche between product documentation and software development kits (SDKs), this particular Resource Kit (RK) has extensive documentation about the NT operating system and utilities that fill some of NT tools' operational and administrative gaps.

Supplement Two's CD-ROMs contain all the information from the original RK, Supplement One, and the NT Workstation Resource Kit. Of course, the original RK also included three printed books, but if you are comfortable using online documents instead of hardcopy manuals you can use Supplement Two as your comprehensive NT resource guide.

Supplement Two's documentation is in Windows Help format which makes it convenient for searching and browsing, and the new stuff includes detailed information about the changes in Service Pack 3, the use of Point-to-Point Tunnelling Protocol (PPTP) with NT's Remote Access Service (RAS) to create Virtual Private Networks, and the use of X.25 with RAS.

In addition to the extensive documentation, however, the RK and its supplements come with a large collection of useful applications, utilities and other programs. Once again, Supplement Two contains all the software from the previous resource kit releases (including updates of many of the programs) as well as an assortment of new utilities. Some of the software is available for download from the web, but some is exclusive to the RK.

The trivial stuff includes new desktop themes and animated cursor editors while much of it is devoted to the serious task of network management. There are a number of command-line alternatives to NT's GUI utilities which benefit anyone who is firmly rooted in the DOS world. Particularly, it offers a way to include advanced NT commands in batch files and scripts. With these tools you can create scripts that automate complex systems and network administration tasks using Registry modifications, event log data and OLE automation.

Finally, Supplement Two contains a sample edition of TechNet, Microsoft's technical resource which is normally only available as a monthly subscription.

All in all, Supplement Two is an extremely valuable resource containing almost 1Gb of software and documentation at a reasonable price.

● *Thanks to Computer Manuals (0121 706 6000) for supplying me with a review copy.*

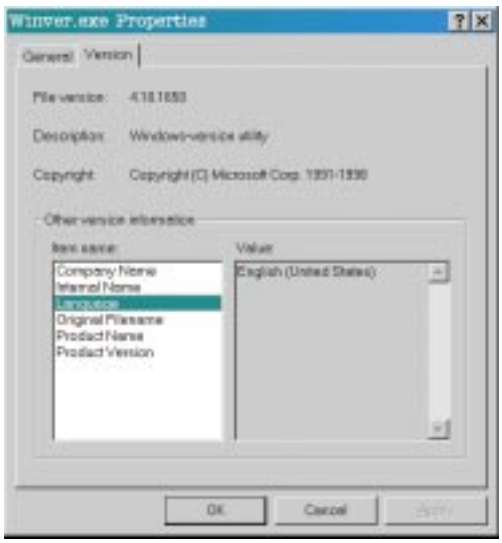

Fig 3 Viewing the properties of WINVER.EXE allows us to see the language setting

3. When the file is located, use the right mouse button to click the file, and then click Properties on the menu that appears.
4. Click the Version tab.
5. In the Item Name box, click Language. The language version is then displayed in the Value box [Fig 3]' ".

Roy goes on to say that you can get the Windows 95 Service Pack 1 for this version from www.eu.microsoft.com/windows/software/localize/pan.htm, or for other language versions www.eu.microsoft.com/windows/ software/localize/localize.htm. He points out that he has used the European Mirror site www.eu.microsoft.com as it is much faster! Thanks very much, Roy, for your assistance.

PCW Contact

Bob Walder can be contacted via the *PCW* editorial office (address, p10) or by email at networks@pcw.co.uk.

Outta space?

Bob Walder goes into orbit about capacity planning with NT. Whether or not you needed to upgrade your hardware was once a matter of guesswork but here is an NT tool to help you.

As I write this column roughly three months in advance of it appearing in print, you can understand why, on reading of Cisco's decision to refuse to loan its equipment to any European test labs or magazines for evaluation, I immediately thought "April Fool".

Well, there's definitely a fool here, and I don't think it is me. This bizarre decision comes hot on the heels of Microsoft's attempts to quash independent tests of its products when it does not care for the outcome of those tests. Apparently, it can do this because of a clause in the licensing agreement (by which we are all bound as soon as we open that envelope containing the CD-ROM) which states, in essence, that the software cannot be used for benchmarking purposes without Microsoft's express permission.

In other words, independent test labs cannot simply get hold of NT Server, test it against NetWare and tell you which is the fastest NOS. Nor can we test Exchange Server against GroupWise to say which is the most scalable messaging system — not without Microsoft's permission. By submitting its products for review, Microsoft is, in fact, giving its permission for them to be tested in this way but it leaves the opportunity for all sorts of censorship.

At least Microsoft is being evenly heavy-handed around the globe. Cisco's decision, however, seems to be more of a slur against European networking professionals in particular. Does it not think we are capable of providing a competent review of its products on this side of the Atlantic?

Or is it, perhaps, that it does not trust its own European staff to provide adequate support to labs over here during such tests? Either way, the only option for Europeans interested in lab tests of Cisco kit is to rely on US magazine reviews or to send their own staff to Cisco's labs to work with its technicians using tests designed by Cisco itself.

The latter option doesn't sound too attractive, does it? As for the former, many European readers would rather see the kit tested by home-grown magazines and organisations. That is why so many US-owned publications spend so much money producing their own UK-specific copy rather than simply rewriting the work already done for the US parent magazine.

I make my living from testing network hardware and software. My company runs an independent test lab and I produce a fair amount of material for UK networking publications. Most of the major publishing houses over here have their own in-house labs, as well, but if the likes of Cisco have their own way, there will not be much for us to test in the future and that cannot be good for us poor "second class citizen" Europeans.

Do you not find it just a little bit annoying, this American assumption that they are the centre of the computing universe?

Tip of the month — have you ERD?

One of the worst things that can happen on an NT system is for the Registry to somehow become corrupted, since for many people the only way back from such a disaster is a complete system re-install. There are a number of precautions you can take to avoid this, though.

The first is to back up your system properly. Yes, I know I go on about this time and time again and I am sure I am teaching my granny to suck eggs here, but it is often those of us who ought to know better who end up trying to restore a system from six-month old backup tapes (I am admitting to nothing here!) Even if you back up your data on a daily basis, please make sure that your chosen backup software is also making sound copies of your NT (or Windows 95) Registry database.

Another good tip is to install two copies of NT (Server or Workstation) on your machine. This is not in contravention of any licensing agreements, since you are only ever going to use one copy at a time. The idea is to install a full copy as normal, with all the bells and whistles you require: this will go in the WINNT directory. Then install a second copy in another directory, say WINREPAIR, which is a minimal installation with just enough to get the machine up and running. If you ever damage any part of your primary NT installation you can always boot quickly and easily using the copy stored in WINREPAIR in order to get at your data on the NTFS partition.

Make sure you keep your Emergency Repair Disk up to date. It is tempting when first installing NT to skip right past the bit that asks: "Do you want to create an Emergency Repair Disk?", since it gives you the option to do it later, but how often have we installed NT and never gone back to create the ERD? The thing is, this disk contains a complete copy of your Registry along with a few other critical system files, and it is used during the repair process. This is where you boot into the install process, but instead of installing NT you choose to repair an installation.

Even if you created an ERD when you first installed NT, how many new pieces of software, new drivers, or new Service Packs have you installed since? Every update to your system potentially changes the critical system files stored on the ERD, so each time you make a major change to NT you should create a new ERD.

Go on, now… find that original ERD (if you created one at all), blow the dust off it, drop to the MS-DOS prompt and type RDISK –S. Go on… we can wait.

Capacity planning with NT

"How can I accurately predict how much load my NT Server can take? I am currently running file and print services together with MS Mail for about 50 users. I would like to upgrade to Exchange Server and put the remaining 40 users on the system but will my existing hardware be up to the task, or should I invest in a new machine?"

So writes P Hurley of Bristol, with a question which is echoed almost daily by various consultancy clients of mine. The other favourite is: "Can I run both Exchange Server and SQL Server on the same box?" Questions such as these are asked daily by network administrators up and down the country. Unfortunately, capacity planning is one of those areas where technology consistently falls foul of business requirements.

You know the problem. Your users expect sub-second response time for all their file retrieval and database query operations. They may get it when the system is first installed and there are relatively few users to thrash it. But time goes by and that sub-second response time creeps up to two, five, or even ten or more seconds as application functionality is increased, and more and more users compete for limited resources.

Determining if and when you need to upgrade your hardware is often a matter of pure guesswork and it is not always possible to be sure exactly which components require upgrading. Perhaps it

Book review

■ **Red Hat Linux Unleashed**
Author David Pitts, et al
Publisher Sams
Price £37.50

Have you ever wondered what to do with those spare 386 and 486 machines with 8Mb or 16Mb of RAM that are lying around your office, no longer capable of running the latest and greatest Microsoft operating system?

Why not get yourself a copy of *Red Hat Linux Unleashed* and turn them into web servers, firewalls, mail servers or FTP servers? For those of you unfamiliar with Linux, it is a freeware Unix-like operating system that can run in as little as 150Mb of disk space and 2Mb of RAM! Red Hat Linux is simply a commercial distribution of the Linux OS, which considerably simplifies its installation and configuration.

For those of you completely unfamiliar with Linux (or Unix in general), however, this is not quite the book for you. It assumes a basic level of Unix knowledge (how to log on, move around the system, perform the more basic system administration tasks, and so on) and does not attempt to take you from scratch at any point. This ensures that Unix devotees will get the maximum from this volume (although I must confess to being a complete Unix novice, yet even I managed to install Linux and configure some services with few problems).

If you know a little Unix, however, this book covers Linux system administration and management and how to handle file systems and printers and all that good stuff. In other words, it contains all the information you will need to make a relatively painless transition from other brands of Unix to Red Hat Linux.

It even starts off with a chapter covering Linux installation, taking you through the whole procedure step by step. And in case you were wondering where to get the software, there is a CD-ROM included with version 4.2 of the operating system, together with a complete developments environment and all the OS source code.

Also on the CD is software for network aliasing (virtual hosting), Perl, Python, Tcl/Tk, LISP, a high-performance web and FTP server, Sendmail SMTP server and client, and an X-Windows system. The majority of the book sets about telling you how to install, configure and manage these in a no-nonsense, easy-to-read kind of way. Buy it now, and make the most of all that hardware you were just about to throw into the skip!

p313 ➢

Questions & Answers — can *you* help?

Q I read your column regularly and have seen articles on SAPS for modem sharing on an NT network. But I'd like to share a modem on a network running NetWare 3.12. There are 30 PCs running Windows 3.11 and 95, and five of the Win95 PCs have modems with dedicated lines. But as the number of users wanting access to a modem for email increases, so the cost of dedicated lines becomes prohibitive. Is there a version of SAPS or another product which would allow sharing a modem and a line?

Peter Williams

A Sorry Peter, but I don't know of a product that will do this. I am sure, though, that there must be something out there that will do the job so I will throw it open to readers of this column. If anyone knows of a modem sharing solution for NetWare servers, drop me a line at the usual address and I will endeavour to review the software in a future issue.

is more memory you need, or a faster disk, or a more powerful processor, or perhaps you simply need an additional server?

Today's distributed network model is a far cry from the predictable centralised model of the mainframe era. Then, if you wanted to know how many users the box would support, or how much additional memory you would need to support a specific load, the IBM salesman would arrive with his big blue book and his calculator and tell you exactly how much money you would need to spend. Given the cost of hardware, capacity planning was an absolute requirement and the job title of Capacity Planner existed in every organisation that used IT.

The random adoption of PC hardware coupled with lower costs made the position of Capacity Planner seem superfluous in most companies. But today, many of the mission-critical systems of yesteryear have migrated down to the PC-based LAN and capacity planning is once again becoming a necessity. But where are the tools?

Oakland-based Bluecurve has developed a capacity and reliability planning tool for NT based on a technique which it calls "Active Measurement". Rather than rely on mathematical models which are difficult to create and maintain in a distributed environment, Dynameasure performs a controlled stress test on the actual NT system components. Current modules available include file and print services, SQL Server, Exchange Server and Oracle, with a web server version to follow.

Although any NT server (and even a NetWare one) can be the target of a simple file and print test, in order to run the others you will need to have the appropriate software (SQL Server, Exchange Server or Oracle) installed. Dynameasure works by creating test "data-sets" on the target server using the chosen application. For file and print tests a whole load of different files are spawned: for SQL Server tests, a dummy schema for a sales order processing system (complete with test data) is generated; and for Exchange Server tests, a bunch of different mailboxes and dummy mail messages (with various types of attachments) are created.

The tests themselves then manipulate these data-sets in various ways in order to simulate the activities of real users. Each transaction performs a typical activity that might be performed by a user in a distributed system, such as: update a customer record in an SQL database, copy a file to a server, or send a number of mail messages with huge binary attachments.

The architecture of Dynameasure itself is distributed in nature, allowing the various components to reside on a number of different machines for maximum flexibility and scalability. There is a single Control Server responsible for the underlying functions and communication between components. This must reside on an NT Server and is configured by an intuitive GUI utility, from which tests can be created and controlled and the results analysed.

Below this in the hierarchy is the Resource Client, which provides the communications channel between Test Clients and the Control Server, as well as providing an intermediate repository for statistics collected by the Test Clients. A single Resource Client can handle a number of Test Clients, and the software can reside on a dedicated machine or can co-exist with either the Control Server or a Test Client.

At the bottom of the tree are the Test Clients, which can be either Win95 or NT machines and each client can run a number of "Motors" depending on processor speed

p314 ➢

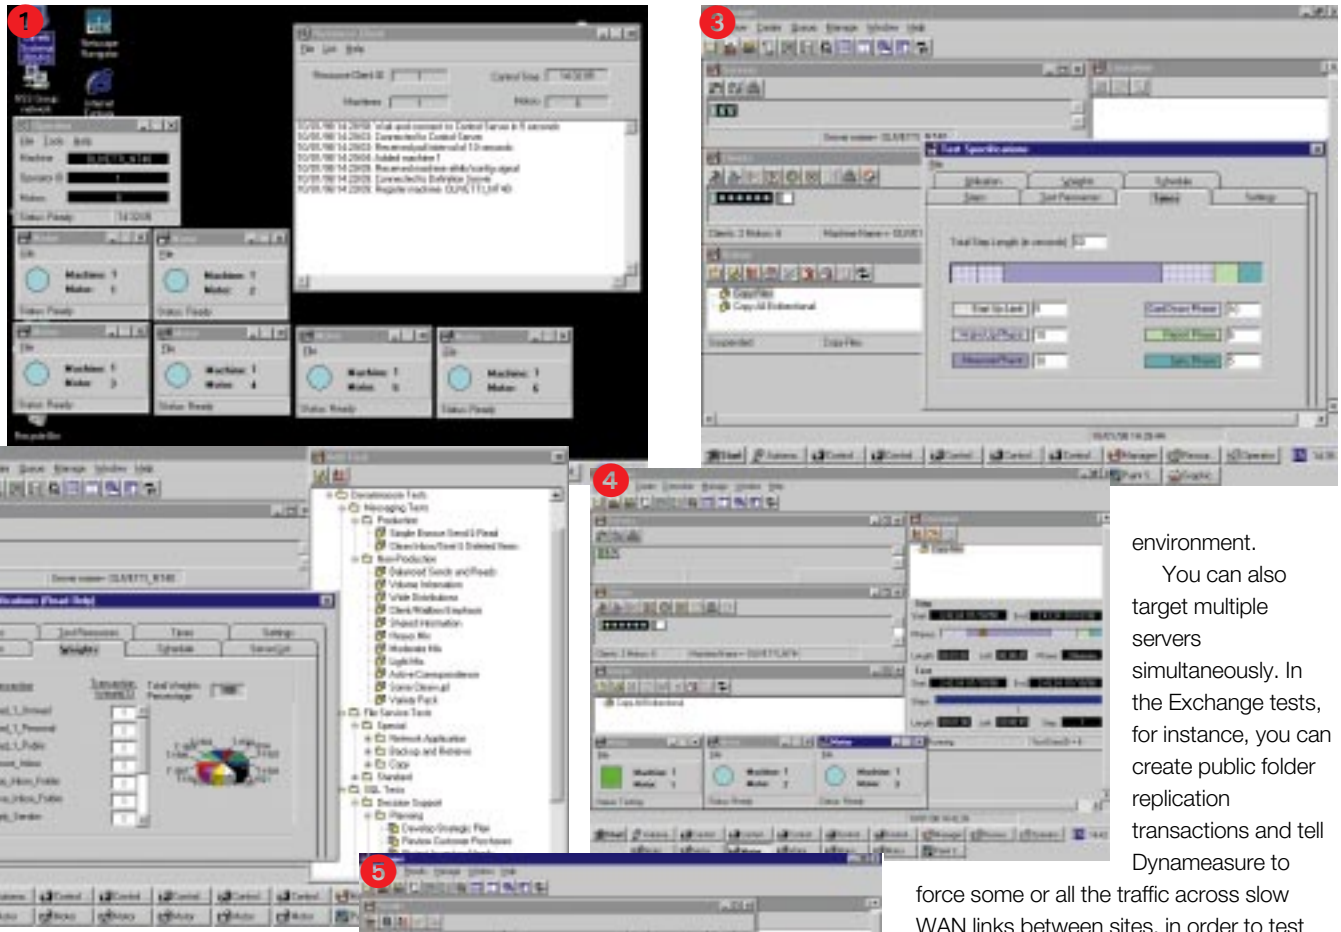

Fig 1 **A number of Motors running on an NT Server**
Figs 2 & 3 **Modifying test parameters**
Fig 4 **Running a test**
Fig 5 **Comparing results of two tests**

and memory configuration. Each Motor is a separate process (representing a separate client) and it is the Motors that execute the transactions against the target servers. Thus, with reasonably-specified clients and a high-speed LAN (100Mbps or more), it is possible to run tests involving hundreds of clients while only using a fraction of that number of PCs.

Each test consists of a number of steps. At each step you can add more Motors, thus allowing a controlled increase in the load, from a single user up to hundreds. At each stage in the test, the Motors collect statistics regarding data throughput, user response time and CPU utilisation and these are collated by the Resource Client and eventually passed to the Control Server. Once the tests are finished, the Manager Utility can be used to produce comprehensive text and graphical reports.

Although a huge number of standard tests are included as part of the package, one of the key features of the Enterprise version is the ability to define your own schemas, data-sets and transactions. This is particularly useful when you need to do some capacity planning for an existing system. For instance, how many users can you add to your sales order processing system before it falls over?

In this situation, you could define your existing SOP data-sets and transactions to Dynameasure and then perform the appropriate load testing against a real-life system to obtain your answer. The product could even be used to aid the database design process by running such tests against a system design which has been defined only in Dynameasure, and which doesn't even exist yet in a production environment.

You can also target multiple servers simultaneously. In the Exchange tests, for instance, you can create public folder replication transactions and tell Dynameasure to force some or all the traffic across slow WAN links between sites, in order to test your WAN infrastructure.

Dynameasure can be used to perform system health checks as well as for capacity planning. Baseline reports can be generated from a normal system and then further tests can be run periodically in order to detect changes in the capacity of an infrastructure over time.

This approach has parallels in other disciplines. Doctors, for example, use stress tests to measure cardiovascular fitness. A treadmill test can detect severe problems immediately and repeated stress testing over time can provide comparisons which can identify potential problems before they become dangerous.

If you need to get a handle on your NT Server capacity and reliability planning, then Dynameasure is worth a look. Contact the UK distributor, the Peapod Group on 0181 606 9990, for more information.

PCW Contacts

Contact Computer Manuals on 0121 706 6000 for Red Hat Linux Unleashed

Bob Walder is a journalist and networking consultant based in Bedfordshire. He can be contacted via e-mail at the usual address networks@pcw.co.uk