Understanding Firewalls

Link your LAN to the Internet and you open it up not just to your legitimate users but to the whole world. A firewall can help to ensure that your corporate information is only available to those entitled to see it.

By David Morton

High on the list of things which the administrators of networks are being asked to do is the connection of the isolated LAN to the global Internet.

In many cases the initial requirement will only specify email: few companies perceive an immediate benefit in allowing their staff to surf the net in their desktop machines.

However, as most of us have discovered over the years, initial requirements have a habit of growing with time. What was seen initially as a simple SMTP and POP3 requirement, to collect incoming mail and deliver it to the right desktop, can soon grow, as the benefits of Web access for market research, and the potential of corporate Web sites for the dissemination of information to customers, become apparent to the management.

But before you connect your company's network to the Internet, there are a number of things which you must consider. Most of us are familiar with the benefits of leased lines against ISDN, or the advantages of large globally-recognised ISPs over their smaller competitors. But at least as important as the technological and managerial considerations is the concept of security. While your site is an island, disconnected from the rest of the Net, it is secure against attack from outside. Once connected, it's potentially as easy for your vital data to leak out as it is for your users to read a newsgroup or look at a Web page.

Risks

If you don't plan and implement appropriate security precautions from the outset, a direct connection from your company's network to the Internet is the equivalent of knocking a substantial hole through the wall of your stock room to the street outside. Just as that hole would make it easier for your mail-room staff to carry your product out to your customers' vehicles, it lets others walk in and remove valuable items. The only difference is that while you can see light-fingered individuals making off with your stock, you may never know that someone's made off with your list of suppliers, customers, or the plans for next year's world-beating product.

The Firewall

Between your network and the outside world, you must have a barrier which restricts access only to those whom you wish to enter, and which only allows those individuals to communicate in a manner which you define: you can let them read the Web pages on your HTTP server, but they must not be able to read the company accounts or send mail using your SMTP server and hence pretending to be from within your company.

Such a barrier is known as a firewall, after the type of installation used in the construction industry to stop the spread of a fire from one part of a building to another (despite the curious illustrations used in some magazines, it's nothing to do with a wall of fire).

In this sense it's something of a misnomer: a real firewall usually only delays the fire's spread for long enough for the occupants to leave and the fire fighters to arrive. It also doesn't assume that there's always a fire on one side of it, and something expensive and flammable on the other. By comparison an Internet firewall has to assume that there are all the fires of hell on the outside and that the contents of your network are slightly more combustible than a stick of dynamite soaked in petrol.

It's important to realise that, no matter how cleverly you design and implement your Internet firewall, and however finely its configuration is honed, it is of value only in conjunction with a correctly-specified security policy. Just because you have a firewall

"A direct connection from your company's network to the Internet is the equivalent of knocking a substantial hole through the wall of your stock room to the street outside."

product installed and configured, doesn't mean that you can dispense with such network niceties as secret passwords or sensible access controls from user accounts, any more than locking the street door lets you dispense with the safe in the cash office.

The Risk Of Connection

Remember that any connection to the Internet is a potential security hole and that an attacker doesn't necessarily need a permanent connection like a leased line to gain access. A leased line gives the attacker a more reliable target to aim at, but a dial-up connection using a modem or ISDN can be just as vulnerable.

In the case of ISDN, many ISPs operate a "dial back" system. When someone tries to access your site (be they welcome visitor or unwelcome assailant) the ISP's ISDN system dials your site and then drops the line, indicating to your connection that it should dial back and establish communication (it's done this way so that you pay the call costs, rather than the ISP). With this arrangement an ISDN connection can be available to the attacker virtually on demand. Unlike a modem connection there's no requirement for someone, or some automated process within the company, to have made the connection to the outside world before the attacker can strike. However even a modem connection can be held up by an attacker: the assailant sends a large email message, knowing that the next automatic mail connection will be up for a considerable period, during which time an attack can be mounted.

Possessions

Having decided you need protection, it's worth considering what you need to protect. Most of us imagine that an attack will be restricted to an attempt to access secret data, and we may consider that all the important data is adequately protected already. Perhaps accounts and personnel data is kept on a separate network, one which will not be connected to the outside world, or perhaps we believe that encrypting such files will be sufficient. But an attack can cost time and money even if it never goes near any top-secret files. The files which the intruder has accessed may not have contained any great secrets, but will it be inconvenient if the information in them has been changed very slightly? How important is it to you that you can rely on the information you store on your system? And if the information isn't important enough for you to want to protect, why are you wasting disk space storing it?

Resources

While the intruder is wandering round your system looking up your canteen's casserole recipe, they're consuming resources. Not a great many, if they're just listing files, but they're your resources, you paid for them, and you don't really want to hand them out to everyone who strolls by.

Reputation

Finally, and perhaps most importantly, you need to protect your reputation. If your system has an Internet connection, then it's highly likely that you send mail using the SMTP protocol, so one of your machines will be running an SMTP server. If you're connected to the net with no security barrier then it's quite possible for someone to connect to that SMTP server and send a mail message which appears - as far as the recipient can tell - to have come from your site. Such a message could be obscene, illegal or any combination of the two, and you may be completely unable to prove that the message wasn't sent by an employee. The damage done to a company's reputation in this way could be incalculable.

Also, it's possible for intruders to break into your system and use your spare hard disk capacity to store pirated software, pornography etc. If the computer press finds out about this, your public relations and marketing staff are going to have a hard time explaining that you really do take security seriously within your organisation. Your customers won't be too impressed either.

Denial Of Service

In addition to simple intrusion and actual data theft, there is a third type of attack to consider: this is known as denial of service, often abbreviated to DoS. This may take the form of mail-bombing, where vast quantities of email are sent to the victim's site until the mail system is overwhelmed. An even more violent intrusion may involve sending large numbers of network requests into the system under attack. Each request needs to be dealt with and resources assigned to handle the traffic. If enough network requests can be fired into the system from outside, then the valid requests generated internally will be swamped, and the entire network brought to its knees.

Such attacks don't just come from personal enemies or commercial rivals. It's easy to imagine that, since you have no antagonists and your competitors are all gentlemen, that there's no likelihood of such attacks being mounted. Sadly there are those who - rather like youths who steal cars for kicks - like to mount such attacks purely as a form of "sport", and who enjoy keeping a score of the networks they've visited, the embarrassing mail messages they've forged, and the servers they've crashed (or run the latest interactive game server on, without permission).

Don't assume for a moment that anyone needs detailed technical exper-

"It's quite possible for someone to connect to that SMTP server and send a mail message which appears - as far as the recipient can tell - to have come from your site."

Firewalls

tise in order to crash your system. Once a method is discovered, it rarely takes more than a couple of days before an "exploit script" is posted to the hacking-related Internet newsgroups. The most infamous exploit utility is WINNUKE - run it on your PC while dialled up to the Internet, type in the address of a Web server, and WINNUKE attempts to crash it. If that server is running under Windows NT, WINNUKE will probably succeed, too, because it exploits a recently-discovered bug in the way that NT handles certain incoming packets.

The Best Policy

Having decided that attacks can and most probably will happen, the next stage is to decide your security policy, which is an essential precursor to the design and implementation of your firewall. Security policies can be divided into four categories.

Do Nothing

You can choose to have no security policy at all. This is an approach which is only appropriate to those with nothing whatsoever to protect.

Obscurity

A second approach is known as security through obscurity: the assumption that because you're a small company or even a single-user machine and not a household name, no-one will ever find you or if they do find you, they won't be interested in your data.

This approach makes a fundamentally flawed assumption: it assumes that attacks on your data, resources or reputation come from rational people who only do such things for economic gain. Sad to say this isn't true, as some of the most significant threats come from those individuals who're doing it for the thrill of the chase, or who derive pleasure from sheer vandalism, making their day by deleting (if you're lucky) or minutely altering (if you're not) someone else's quarterly results.

The kind of attacker who will break into a system and wreak havoc simply to add another victim to their score, won't care that the victim is a company no one's ever heard of, it's simply one more on the score sheet. Indeed some such individuals are thought to target new arrivals on the Internet (by watching publicly available sources of registration information), which makes it doubly important to establish your security policy and secure the system before you establish your connection: don't connect and assume you can tidy up the details later, because the damage may already have been done.

Host Model

The third type of security model is the one which most network managers and administrators are familiar with: the host security model. In this case, each host is secured separately and every effort is made to eliminate security holes on a host-by-host basis. The problem with this is that it tends to assume that all those who are in a position to change the configuration of a host machine are sufficiently knowledgeable about the hazards and threats to that machine. While this is the case when there are defined roles for servers and workstations (as in the days when we all had DOS clients and NetWare servers) those boundaries are now much less defined. It's quite easy for a user to turn a Windows machine into a server "just to share the data for an urgent report with Fred down the corridor" without realising that - in some configurations - he's just shared that directory with the entire wired world. Indeed in some cases creating such an insecure share can expose not just that subdirectory, but the entire drive.

Network Model

When you're connecting your internal network to the Internet, what is required is known as a Network security model. In implementing a Network security model and controlling network access to all your hosts and services, rather than securing them individually with a host security approach, a single network firewall can protect hundreds or even thousands of machines.

With a correctly-configured firewall, even the most suicidal user can only breach the security of their machine within the confines of the network within the firewall, and not expose the system to outside interference no matter how hard he tries.

That's not to say that a firewall should be used as an excuse to abandon the basics of good host security. Just because you've bought the best firewall product available for your network configuration, and paid a consultant vast amounts of money to set it up, does not give your users *carte blanche* to abandon the basics of personal passwords which are kept secure (and not on a note stuck to the screen). A firewall should always be an adjunct to good host security, not as a substitute.

"With a correctly-configured firewall, even the most suicidal user can only breach the security of their machine within the confines of the network within the firewall."

Firewall Functions

Logically a network firewall is less of the impenetrable barrier suggested by the "firewall" metaphor and more closely akin to the moat around a medieval castle - complete with drawbridge - or perhaps an international border, with barbed wire and minefields perforated by controlled crossing points. It restricts all accesses from outside the network to a carefully controlled point, it prevents the attackers from even testing your other defences

(such as an individual host's security) and it restricts traffic leaving the network for the outside world to the same point.

Having made sure that all traffic entering or leaving your company's network passes through this one point - the firewall - the firewall software can then ensure that all traffic passing through it conforms to the security policy. This security definition depends on your system requirements, and is a matter for careful consideration when configuring the firewall software itself. However the firewall configuration is only a small part of the whole problem of defining how your site security should work: it's no good having the most rigorous firewall implementation in the world if your users are allowed to have modems in their workstations and dial up their own Internet accounts from the desktop to collect their personal mail.

Platform Choice

Physically the firewall is unlikely to be one identifiable lump of hardware with a large "firewall" label on it. It is much more likely to be a combination of hardware components - such as a router and one or more host machines - which may be dedicated to the task of running the firewall, or which may undertake other tasks in addition to this role. There's no security reason why the host machine running the firewall software should run the same operating system as the rest of your network - you don't necessarily need an NT firewall for an NT network or a NetWare version for a Novell site. The firewall is simply a black box, and what goes on inside it is of no relevance to the way you manage it.

Indeed in many installations the only machine in the building running a flavour of Unix may be the one implementing the firewall and the Internet connection. This is because the history of Unix and the Internet is such that Unix firewalls have been under development for much longer, and are much better understood by systems administrators than their younger siblings. Of course, there may be good management reasons for not introducing an entire new network operating system to your site simply to implement this one function. An NT firewall may be less developed, but if your network administrators know and understand NT, you're much more likely to get a solid and secure configuration from that, than if the same team are fumbling about on a Unix machine they don't fully understand.

However there are other crucial considerations when considering the choice of platform for your firewall, not least of these is cost. If you choose Linux (the free Unix clone) as the operating system for the machine which will run your firewall software, then all of the software you need is available free and much of it comes with the more comprehensive Linux distribution sets. On the other hand if the idea of getting security software for nothing bothers you, and you're unsure about using something as strange as Linux, then you may wish to consider a commercial alternative. Be prepared to dig deep, however, as some of these commercial firewall packages come expensive. Shiva's Raptor product for NT, for example, costs US$6,500 for 50 users rising to US$15,000 for the unlimited version. [*An article on setting up a Linux firewall is currently being prepared for future publication - Ed.*]

The detailed implementation of a firewall can be a hugely complex field, with variations on numbers of routers, the presence of one or more so-called "bastion hosts", ie, machines to which incoming attacks are deflected by the firewall.

Even when one of these many architectures has been chosen because of its suitability to the network you wish to protect, the configuration of the software is not something to be taken lightly. Even hardened firewall experts have been known to spend many expensive days getting things right. Whole weighty text books exist on the subject, some of which are only comprehensible to those with a degree in obscure languages and advanced obfuscatory techniques.

An Example

To take the simplest possible case as an example of a firewall in action, we might choose what is known as a "dual homed host" architecture. In this implementation we have a host computer which has two network interfaces. One of these is connected to our internal network and is probably physically represented by an Ethernet card. The other is connected to the outside world, or "hot" side of the firewall. This might be another Ethernet card connected to an ISDN or leased line router, an ISDN Terminal Adapter, or even a humble modem.

Normally such a dual-homed machine would route network packets from one network interface to the other. All data that needs to get from one network to the other would simply go in through one end of the firewall machine and pass straight out the other side.

But to serve as a firewall machine, this facility is disabled. IP packets from the Internet can reach this host, but can go no further, and in the same way packets from the internal network can reach it, but cannot leak out onto the Internet outside.

Of course such a configuration would have little utility. You could achieve the same level of service (ie none at all) by simply not connecting

"Physically the firewall is unlikely to be one identifiable lump of hardware with a large "firewall" label on it. It is much more likely to be a combination of hardware components."

Firewalls

to the Internet at all. What our firewall machine now has to do is act as a filter of the data passing through it: to permit some kinds of connection and data, and to forbid others.

Packet Filtering

This filter has to look at the data packets arriving at each side, and decide if it is appropriate to pass them through. An IP packet has a header which contains information on the source address, the destination address, the source and destination ports, what protocol the packet represents, and so on. When a conventional router handles an IP packet, it only makes one important decision: if it knows how to send the packet on its way to the correct destination address it does so; if it can't do this, it returns the packet to the sender with a message saying that the destination is unreachable.

But the router within our firewall is more cunning than this. It's what is known as a screening router. This looks not just at the source and destination addresses, but at the type of protocol and the port which the packet represents. The screening router doesn't only consider if it can pass the packet on to the destination address, but also if it would be a breach of the security policy if it were to do so.

In an email-only installation you might configure the screening router to block all incoming packets other than SMTP, and this is a common configuration when sites have only a hardware router to implement a simple firewall. In this case it's likely that the outgoing packets will not be filtered at all: everyone within the firewall has access to everything outside.

The next level of complexity might be to allow email packets to flow, but to block all packets to and from hosts which you know to be dangerous: a colleague has suggested that the local schools and educational establishments might come into this category, although that might appear a rather harsh judgement to some. You can make the packet filtering ever more complex: allow ftp, for example, but block telnet, rlogin or similarly hazardous operations.

Rules

It's important that the software you use to implement packet filtering has an understandable way of applying the filtering rules, and that it applies the rules in the order you specify them. Some packet filtering systems have been known to attempt to re-order the rules to make the filtering software more efficient and to improve its performance. The problem with this kind of improvement is that it can lead to unexpected results, and be difficult - if not impossible - to debug when you find that packets you would wish to have transmitted unimpaired are being dropped on the machine room floor by the screening router.

Equally important is that the router should log those packets which it rejects. This is essential, not only in the setup and debugging phase, but in the day to day operation of a screening router. If you don't log the dropped packets (and check the logs regularly), the screening router will leave you blissfully unaware of an attack in progress until the attacker finds a chink in your armour and wreaks havoc with your system. This is just as important as detecting failed login attempts on a host security system: a few dropped packets may not be significant, but a consistent pattern should be examined in greater detail: perhaps with a view to implementing more draconian filtering on packets from the host in question.

The problem with packet filtering on its own is that it is something of an all-or-nothing operation. If you choose to allow a service, then you allow all elements of that service. If you wish to permit some operations within a service but deny others, then a packet filtering system like a screening router may not be sufficient. Certainly a simple packet filtering scheme on a hardware router should not be regarded as sufficient protection on its own, in any but the simplest - and most severely filtered - implementations. While a screening router can provide a reasonable degree of protection from unsophisticated attacks, it can sometimes be overcome, and should be used in conjunction with other techniques.

The Proxy Server

One way of giving greater control to the system administrator, and greater flexibility to the firewall, is to use proxy servers. These reside on the dual-homed host already mentioned and effectively impersonate the Internet, to the user within your network. This is - as far as the user is concerned - a rather good example of smoke and mirrors: the users think they're directly connected to the Internet and that their commands are directly controlling the real server, however in reality their instructions - and packets - go no further than the machine running the proxy - all attempts to make a direct connection are filtered by the screening router.

If we consider the example of a proxy ftp server, this would run on the firewall machine, and the user's ftp client would communicate with it. The proxy server in turn would communicate with the real ftp server out on the Internet. Requests for file transfers from the client machine are not passed directly to the unknown server but are first screened by the proxy, which can implement rules according to the site's

"It's important that the software you use to implement packet filtering has an understandable way of applying the filtering rules, and that it applies the rules in the order you specify them."

security policy. For example, one might choose to block access to files from certain sites, or more realistically one might forbid outgoing file traffic from all but a chosen few machines and to a chosen few trusted recipients.

Such proxy servers are available for most of the common Internet protocols: Telnet, ftp, http etc and if the firewall machine is running Linux, they're available free. From the user's point of view the client software may need to be told that it's communicating via a proxy, but this is usually a minor configuration option. On Microsoft's Internet Explorer, for example, it is only necessary to tell the software that a proxy server is being used, and the address of the machine on which the server resides. Once the client software has been set up, so far as the user's concerned he has a conventional connection to the Internet, until they attempt to do something that the security policy doesn't permit.

It's important to realise that a proxy server is not - on its own - any form of effective firewall. In addition to providing a proxy, you have to implement a packet filtering screening router to completely block any direct access to the Internet from the users' machines. All too often one sees proxy servers described as security mechanisms, when on their own they are no more than an Internet cache: something which can be quite useful, but which should not be confused with a secure system.

Monitoring

An additional benefit of a proxy server and packet filtering system,is that it makes monitoring of Internet use quite straightforward. While most users might not be delighted to discover that the jpeg files they've been viewing at lunch-time are listed - and possibly cached - in the proxy server, from a system administrator's point of view letting this fact "leak out" to the users can have a dramatic - and highly beneficial - effect on the cost of a dial-up connection.

While it's tempting to add a few more useful services to the firewall machine - after all it's probably not doing very much, and it might be the only Unix host in the building, this temptation should be firmly resisted. It's a good idea to ensure that - from the point of view of the outside attacker - the firewall machine looks an unattractive and boring target. It should have no user accounts, nor anything resembling interesting data. Ideally it appears as just one IP address with no further hosts visible beyond it, and for all the attacker knows it might be a solitary home user's machine with no useful data on it worth stealing.

Conclusion

It's important to realise just what a firewall - any firewall, no matter how complex the implementation - can do for you, and what it can not. It can concentrate your security efforts to one single point, and enforce a well-defined security policy. It can limit your exposure to outside threats but only those you have anticipated. If a new threat arises which wasn't taken into consideration when the design and configuration of the firewall (and its associated security policy) was being planned, then it won't protect you.

Equally it won't protect against viruses (it has no way of knowing that a network packet is part of a virus) or against incompetent or malicious staff. If there are ways around the firewall (like other access ports to the Internet from workstations) then it's as useless as a real firewall with a large hole in it. However if things do start to go wrong, and you've controlled your access to and from the Internet with a well-configured firewall, it represents the best way to log and monitor what new threat has appeared or what unanticipated problem has come to light, so that steps can be taken to combat the new hazard.

But the most complicated firewall implementation imaginable won't protect your valuable data, if it's implemented as part of an ill thought out and implemented security policy or if you treat it as an afterthought to be bolted on to your Internet connection when you've got the bugs out and it's all working nicely. It must be planned carefully, before you expose your company data to the hazards of the Internet.

If it seems like too complex a problem to tackle, treat it as an opportunity to really understand how your valuable data is secured, what policies you're implementing, and how you're controlling them. Examine your managerial approach to data security, see what policies are defined, and which ones you have allowed to define themselves, or just grow and change at random. But above all do consider all this before you make that fateful connection, otherwise you may find you've closed the firewall only after your building's ablaze.

PCNA

The Author

David Morton (dmorton@cix.co.uk) is a networking consultant and freelance writer, who previously worked in an R&D capacity at the BBC.

New Reviews from Tech Support Alert

Anti-Trojan Software Reviews

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

Inkjet Printer Cartridge Suppliers

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

Windows Backup Software

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our “Editor’s Choice” award

The 46 Best Freeware Programs

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.