Attacks on firewalls

ONCE UPON A TIME IN THE WILD WEST

FRANK BERNARD

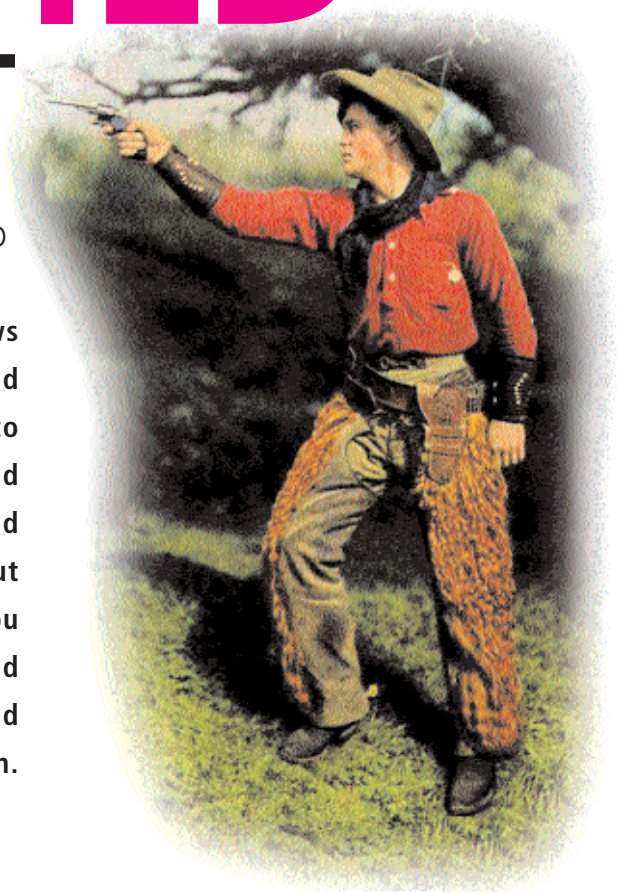

To defend against the digital bows and arrows of the Internet, you need a firewall. But when you've gone to the trouble of installing one and find nothing happens, you may be tempted to think that all that stuff about hackers we've been warning you about is all a load of rubbish, and you've wasted a lot of time and effort. You'd be wrong though.

As the new, involuntarily Open Source (following the recent high-profile attack on its systems) Microsoft illustrates (see *https://www.linux-community.de/News/story?storyid=531*), hackers are a threat that must be taken seriously. But the attack on Microsoft merely exposed the loopholes that were already there, and out of idleness, or to avoid interrupting workflow, were not closed. The fact is that Microsoft, and most other sensible companies, have firewalls installed, but this fact alone is obviously not adequate for Internet security. The combination of several (in themselves harmless) individual security defects creates a security loophole. In the Microsoft case, this involved a Trojan horse on an employee's home PC that allowed a hacker access to Microsoft's internal network by stealing passwords.

Several things could have stopped this attack, or at least reduced its impact:

- Had the Trojan horse been discovered by an up-to-date virus scanner, the passwords necessary for the break-in would not have been stolen in the first place.
- Had the firewall prevented a direct access to the computer, the passwords that were passed on would have been of no use.
- Had the analysis of existing system logs been better, the damage could have discovered earlier.

But what lessons can be drawn from this and what does it have to do with our topic?

- Security costs money. Nobody likes to leave large amounts of money in a desk drawer, so they lock it in a safe. Data is the most valuable asset a company can have nowadays. So computers containing this data must be protected.
- Security costs money on a permanent basis: Even a safe with lots of money may not be guarded at night, and a safe built in 1870 may be heavy and look solid, but probably isn't as secure as a more modern one. Security therefore needs to be constant, and must be adapted to thwart new break-in techniques.
- Security should overlap itself: double locks rather than single ones, infra red and microwave motion detectors – virus protection plus a firewall.
- Security and openness are mutually exclusive: This is the central tenet of an Internet security solution. The more information an attacker gets, the more weak points he can discover and exploit.

Restriction of information – divide and rule

Many attacks on a company network are planned in advance. The first thing the enemy needs is knowledge about the structure and possible weaknesses of the network. If this is made very difficult for the attacker to start with, it may be that he will lose interest in breaking in, or try the company next door.

For a LAN with Internet connection, true high security means:

- An external DNS server should not administer any internal names or addresses.
- Information, for example on the operating system (which the attacker finds out on log-in), version status and email system in use (by the SMTP greeting) should be suppressed if possible. An attacker could exploit these pieces of information for targeted attacks on known loopholes.
- Not even the DNS name *firewall.company.com* should be announced externally – even if it is obvious. Select a neutral name, for example *mail.company.com* if the firewall is also a mail exchanger.
- Also, if the firewall is scanned, no information should be revealed which could be useful for an attack. See the "Portscan" box for more on this.

Mastery lies in (access) restriction

Frequently, due to idleness, lack of staff and money, or simple pressure of time, compromises are accepted in Internet security. It can turn out to be almost impossible to implement changes later because of the structures that have been created. A more secure protection of the total network is usually possible only as the result of fundamental restructuring.

Portscan – right or wrong

This is how an nmap output ought to look if your system is reasonably secure (LinuxWall V2, Linux-Kernel 2.4.0-test10).

```
Starting nmap V. 2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Interesting ports on ([Internet address deleted]):
(The 1520 ports scanned but not shown below are in state: filtered)
Port       State      Service
22/tcp     open       ssh
25/tcp     open       smtp
113/tcp    open       auth

TCP Sequence Prediction: Class=random positive increments
                    Difficulty=1580537 (Good luck!)
No OS matches for host (If you know what OS is running on it,
see http://www.insecure.org/cgi-bin/nmap-submit.cgi).
```

An example of system wide open to attack:

```
Starting nmap V. 2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Interesting ports on [Name deleted] ([Internet-Address deleted]):
(The 1511 ports scanned but not shown below are in state: closed)
Port       State      Service
7/tcp      open       echo
9/tcp      open       discard
13/tcp     open       daytime
17/tcp     open       qotd
19/tcp     open       chargen
21/tcp     filtered    ftp
135/tcp    open       loc-srv
139/tcp    open       netbios-ssn
158/tcp    open       pcmail-srv
427/tcp    open       svrloc
5631/tcp   open       pcanywheredata
65301/tcp  open       pcanywhere
TCP Sequence Prediction: Class=trivial time dependency
                    Difficulty=16 (Easy)
Remote operating system guess: Windows NT4 / Win95 / Win98
Nmap run completed – 1 IP address (1 host up) scanned in 19 seconds
```

- There should be only one transition point between internal and external network, ISDN-cards or modems should be banned.
- A firewall should only open the services (ports) to the computers that are absolutely necessary.
- If possible, accesses should be permitted only from the LAN into the Internet, as these can be controlled more easily.
- Computers that are to be accessible from the Internet (or from an ISDN- or analogue line) should be placed in separate segments (so-called demilitarised zones, DMZs).
- The bandwidth for a specified service should be set to a minimum value, which limits the possibility of Distributed-Denial-of-Service attacks.

A clear network topology also creates the pre-condition for a transparent firewall policy. "Security by obscurity" just does not exist.

What is an Internet attack?

The answer to this simple question is extremely difficult. In certain cases this cannot be answered in a general way, and can even vary from system to system. Linux 2.4 with *netfilter* has received some

crucial improvements which both increase security, and, in cases of doubt, at least allow a classification. Almost everyone would agree with me if I were to claim that a ping (ICMP *echo request*) was harmless from the point of view of technical security. A ping on several IP addresses in succession could, however, be a host-scan (to determine which hosts are in fact accessible from the Internet, but are not registered in the DNS). A ping on a broadcast address (255.255.255.255) is a so-called Smurf attack, which can unleash veritable storms of packets.

Many Windows computers are pre-configured so that when online they try to make contact and exchange information with all computers that reply to the target ports 137-139. Windows computers also have 137-139 as source port. But there are implementations that use other, usually unprivileged source ports (for example Samba). So, if a cowboy computer rides up on the Internet and tries to access the SMP ports, the source port of the spy computer is then, too, often in the unprivileged domain. This is obviously not a configuration error, but a spy mission.

One question of particular interest involves computers that don't even "exist" yet. One of my customers was assigned a Class-C block of network addresses, with only the firewall and the router using any of them in the first few days, yet all addresses were probed during that time

Many packets use security loopholes which have been known and dealt with for years, such as so-called XMAS-packets (special TCP-packets, in which all flags are set) or packets with a prohibited destination TCP port (for example Port 0).

These are generally used in "Denial-of-service" (Dos) attacks, as they make the firewall (temporarily) unusable, therefore allowing no traffic at all to get into the network. But much more common are the attacks that collect information. Once this has taken place, the information is then used for targeted attacks.

Spotting an attack – the needle in the haystack

Many administrators shy away from implementing a paranoid security policy, especially one that reports every suspect packet, since this will, of course, make their logs a great deal larger (at least at first). However, this is the only option open to you if you want to tell if a presumed attack is real or not, and learn from the experience.

Being careful can mean that it can be easier to spot and fix configuration errors (such as incorrect network masks) on PCs. What's more, administrators will know in far greater detail what's really going on on their networks. Any change to the LAN infrastructure will be clearly visible.

So, after a few days of problem solving, all that should remain are packets that could launch potential attacks. Indeed, a *nmap* scan, as shown in the two tables, gives rise to about 1 MB of log. The size of the log file alone ought to sound alarm bells. Naturally *netfilter* has some advantages over its predecessor *ipchains*, in that such packets can be reported but only for a certain time. Remember, a hacker will want to be as inconspicuous as possible, which is why it is bad policy to suppress the packets completely.

Cops and robbers – Prosecuting hackers

Let's reverse the tables now, turning the hunter into the hunted. Using his IP address (which we will already have tracked) we can identify the attacker. This is because the hacker's IP address isn't normally false, as he'll want to get a response to his attacks or probes. In theory then, all we need is a quick *nslookup* and we'll be able to determine who his ISP is in order to send them e-mail about the antics of their customer.

Unfortunately this is only successful in very rare cases. The big providers are certainly able to say when and who was active with which Dial-Up addresses and can therefore identify the customer. This costs money though, so it doesn't always get done. Computers with fixed (and consequently easily traceable IP addresses) on the other hand, are often not the originators. Instead, they are more often victims (usually of Trojans) themselves. In the end, then, due to the difficulty in tracking an attack, and the large number of attacks that generally occur, it is usually pointless to try to find and prosecute attackers.

Outlook – much better with 2.4!

Quite apart from improved packet filtering, Kernel 2.4 offers additional options and countermeasures against SYN flooding and IP spoofing. Through Traffic Shaping (bandwidth restriction), for example, it is possible to guarantee that even if there is a "Denial-of-service" Dos attack, business will carry on as usual. This is something to look forward to, but even so, you must remember that no matter what security measures you implement, hackers will find a way around eventually. Indeed, no firewall is completely secure, and it is only through a secure configuration *and* permanent monitoring that attacks can be spotted and potential new (or old) security loopholes can be closed. This is where an analysis of the situation demands very good knowledge of IP protocols. And one of the best ways to get started on the right road is to conduct an attack on your own systems using one of the many tools available.

In all, then, yes, Internet security costs money, more money than many IT managers want to spend. But by ignoring the ever present, ever increasing problem of hackers, in the long term your company stands to lose much more money than a good security policy could ever cost. ■