Securing The Edge

by Avi Freedman, Akamai Technologies

Common wisdom has it that enterprises need firewalls to secure their networks. In fact, as enterprise network practitioners can attest, the “must-buy-firewall” mentality has pervaded the field.

Maybe you’re a believer too. But if you have any geeks working for you, do you realize they may have tunnels behind your firewall to their home machines? Oh, so you don’t allow ports other than 80? That’s okay—HTTPS-based VPNs, or users running ssh on port 80 to forward ports or run network extensions, are all growing in popularity.

How often does your network and service topology change? How often is your firewall updated? Are your home users running firewalls? What privileges do they have once they access the VPN from their home network? Do they run open 802.11 or have teenage sons or daughters with curious friends?

If you feel smug about your answers, what about other companies you know of?

I ask these questions to illuminate gaps in centralized security, which I see as an obsolete concept that probably never made much sense. It may be an over-used buzzword, but I believe “edge” security makes more sense. We should secure devices connected to the Internet, and ensure they only appropriately trust other devices on local and remote networks. Today, it’s too easy to drop an 802.11 bridge, or have rogue tunnels penetrate any corporate perimeter, so if your hosts and services aren’t secure, all perimeter security does is provide false confidence.

To give you some context for my assertions, let me tell you about my experience. In the late ‘80s and early ‘90s, I ran a university’s public access and departmental router networks, and Unix machines. In the last decade I founded, led, and am now peripherally involved with the first Philadelphia ISP (netaxs.com). My “day” job is with Akamai, where we run over 13,000 mostly Unix machines, providing web and application delivery—HTTP, HTTPS, DNS, and Streaming. These days I’m more on the network side of things, but I still do active sysadmin duty for netaxs.

Typically, regional ISPs or national network pro-

Centralized security is an obsolete concept AND WAS PROBABLY

NEVER A VERY GOOD IDEA. SECURING DEVICES CONNECTED TO THE INTERNET AND ENSURING THAT THEY ONLY APPROPRIATELY TRUST OTHER DEVICES ON THE LOCAL AND REMOTE NETWORKS—COMMONLY CALLED “EDGE SECURITY”—IS THE KEY.

viders need to run services such as TACACS+ (router authentication/logging), Radius (general authentication/logging), NNTP (Usenet News), SMTP, DNS, and HTTP for back-end functionality or to make users happy. Particularly ancient ISPs (such as netaxs, panix, std.com, and io.com) even run shell servers with access granted to anyone who ponies up \$20 a month.

This may surprise you, but few regional providers run firewalls! Some place filters on routers to help with DoS-type requests, but not to block services for host security. Still, most old-time regional ISPs have good security records—even those that supply shell account services. Many ISPs are loath to set up firewalls, and instead farm out this potentially profitable business because of liability concerns—primarily that network or service topology will change and the firewalls won’t be updated.

Some customers have implemented firewalls, only to have dialup access and hosts running RIP bridge their internal network to one or more networks or the Internet. Recent stories tend to involve users running tunnels in or out of corporate networks, or around insecure home or remote office networks with VPN access to the main network.

PORT 80 USE MAKES SECURITY TOUGHER

In the old days, SLIP or PPP over dialup lines were used to “extend” the “stub networks” that end users run to the Internet or to the core of corporate networks. Over the last few years, the mainstream approach has been to use PPTP or IPSec-based VPNs. But recently, the mainstream security industry has started selling HTTPS-based VPN systems, generally running on TCP port 80, and estimates

suggest these HTTPS-based VPNs will rapidly gain market share.

Also, many Unix-literate “advanced users” connect their personal machines to corporate or ISP networks with ssh tunnels, either doing forwarding of specific ports or actually running PPP inside of ssh, using ssh as the mechanism for both transport and encryption. Users often run these ssh sessions so they appear to be outbound connections to port 80, and thus are not typically blocked by firewalls or router filters that deny access to web applications. Similarly, HTTPS-based VPN technology can make outbound connections in any environment where users can access remote Internet applications using HTTPS. There are also VPN applications, such as vpnd, which I’ve run under Linux to run VPNs on arbitrary ports, initiated from inside a firewall.

Basically, either ssh or an HTTPS-based extension of the corporate network can allow connections originated from the outside to reach the inner core of the corporate network without passing through firewalls for proxying, packet filtering, content filtering, or monitoring! One possible configuration is a host configured to NAT connections originated from the outside. Another is a host configured to route a single IP address from the local LAN’s IP range over PPP via an ssh tunnel, with the corporate host also configured to proxy-arp for the remote IP address. Clearly enabling MAC address filtering on the switched infrastructure could block the second configuration, but that is rarely done, and still does not block a host from providing access from the outside via NAT.

In the old days, network administrators simply withheld analog phone lines from users to protect core corporate network from unplanned access. Protection becomes more difficult with VPNs on potentially any port. One can reset sessions after a few minutes, but that may break other applications. Or, one can search network flow statistics for lengthy communications with strange external hosts. But most perimeter security vendors today don’t have such features, and in a large network it is difficult to track down all persistent connections.

A draconian measure is to require all traffic be proxied, including DNS, as you can certainly build a tunnel via UDP or TCP on port 53. But such proxying blocks all access to HTTPS-enabled Internet websites. If you take this step, be prepared for the hassles involved with making major modifications to user behavior.

PROTECT YOURSELF

If you’d rather not cut off your internal network from the Internet, I strongly suggest you secure end devices before focusing on centralized (non-host-based) security. Please

note I’m not saying firewalls are useless. It’s just far too easy and commonplace to rely exclusively on them. This way lies breach! Firewalls can be useful for logging activity, and for intercepting worms, viruses, and undesirable content (although server-based blocking is often more effective). But first secure the hosts, and then layer on these additional features if you require them. Here are a few more specific suggestions:

- Workstations running Windows need personal firewalls. (I recommend ZoneAlarm or BlackIce, but several such packages exist.)
- Servers running Windows should have all inactive services/ports locked down, and should be port-scanned to ensure no goodies are left functioning by the OS vendor. Programs such as nmap from insecure.org are excellent for finding open ports on your network.
- Workstations and servers running Unix-type OSs should have all inactive ports turned off, and verified with port scanning. It’s slightly off subject to discuss hardening of public-access shell account servers, but it bears mentioning that most distributions of Unix OSs have far too many suid/sgid programs.
- Peer-to-peer networks are probably not a great idea. In fact, any potentially “back-doored” or just poorly written software that communicates outside of the network should be administratively restricted.
- Obviously, a regular regime of monitoring and patching of workstations and servers should be run.

These suggestions have been old news for 10 years—but people often ignore or de-prioritize them because they view their external and/or internal firewalls as a safety net. But there is no perfect safety net—even securing hosts won’t make you impregnable. Since any determined attacker (an employee’s teenager, for example) can probably find one user to “come in as” through sanctioned and supported VPN channels, it’s also important to ensure authenticated users don’t get the keys to the kingdom when inside your network.

You might be surprised at how many companies implement hacks to enable web access to monolithic back-end applications such as SAP and Siebel. I’ve seen multiple instances of these “light” client implementations authenticating to the back-end using a common username and password that allows access to all data at the application or database level!

Also realize that any NFS or SMB implementation for sharing user data effectively opens up all data to anyone who can obtain local root/administrator privilege. Even if mail services require separate authentication, many users

store mail on file servers, either as Unix mailboxes or as backups of laptop hard disks.

I've talked about the tunneled VPNs people can build through port 80, but many companies are moving to a web services architecture that places port 80 as the gateway for an increasing number of applications. While securing web services is the subject of another article, I'll remind developers that just being able to execute objects/functions/code demands the highest level of security—in the areas of parameter checking, access restrictions to object/function and parameter definitions, and of course granular access restrictions to the objects/functions.

THE BOTTOM LINE

I firmly believe that securing the hosts and services on a network is the only reasonable approach to ensuring security of corporate data. Obtaining network access shouldn't open up access to machines in the network, and obtaining access as one user shouldn't open up other user data without additional authentication.

I can hear the objections now—this kind of edge security is difficult!

Yes, securing all hosts, and then ensuring user data sets

are isolated from each other, is a thorny challenge, beyond state-of-the-art for many companies. But moving toward edge security is both necessary and worthwhile. Let's face it: most large networks are probably already penetrated or easily penetrated by a moderately determined attacker. Plan to replace NFS with AFS; ensure that users don't back up laptops to file servers with weak user isolation; and, of course, secure every internal host as if it was directly connected to the Internet. Good luck, and may the force be with you. ☺

AVI FREEDMAN (freedman@freedman.net) is currently Chief Network Scientist with Akamai Technologies; Chief Network Architect and Board member of FastNet; and is engaged with the startups, Money For Clue Consulting and Watch My Net. Avi has been working with Akamai for 3 years. Prior to that, he led Engineering at AboveNet. In 1992, Avi started Netaxs, the original ISP in the Philadelphia area. In 2002, Netaxs was acquired by FastNet. Avi's interests lie in routing, performance analysis, and security.

© 2003 ACM 1542-7730/03/0300 \$5.00