13

Network Layout

intranet (in**'** trə nĕt'), *n.* Any collection of networks owned by a single entity that is too large to be controlled by that entity.

—

Corporations and other large entities often imagine that their networks are contained within a secure perimeter. While this may have been true when there were only few hundred hosts involved, large companies now have intranets with tens or even hundreds of thousands of hosts.

These nets typically have several firewalls, numerous connections to business partners (called *extranets*), VPNs to remote offices, provisions for telecommuting, insecure links to other countries, numerous cheap wireless base stations, and innumerable fax and data modems.

The control and management of such a large collection of networks is an open research problem. Why? By design, there is little centralization in IP technology, which improves the robustness of the network. But it also makes it hard to control from a central point, which is pretty much the CIO's job description. The internal domain name service may be centrally controlled, and the address allocations on corporate routers should come from a central authorization source. But it is easy for a rogue manager to purchase an Internet link, and modems are very cheap. A modem link to an ISP is an easy and cheap end-run around corporate network access policies.

Traditionally, network managers have lacked tools to explore their networks beyond the known bounds. It is easy to run network management tools on routers you know (providing that you have the community string), but it is harder to find new or unknown connections.

Intranets are constantly changing. Mergers and acquisitions bring new network connections—the board does not usually consult with the network people on the compatibility of the merging networks and the pending unification of their access policies. Business partners are connected, and sometimes disconnected.

Technical people tend to change jobs frequently. One of us consulted with the IT staff of a major company in 1996. When we revisited them in 2001, not a single person we had met still worked for the company. In fact, most of the 2001 crowd were recent college graduates. This

is typical: The technical people (and the CIOs!) tend to move on, and the networks they leave behind never match whatever documentation they happened to create. Connections are forgotten, as are the reasons for those connections in the first place.

The job of managing security is made harder by uncooperative employees. We know of one Silicon Valley company that tried to control incoming modem access by forbidding modem lines. The employees, who liked to dial directly into their computers from home, responded by installing "fax" lines. At the end of the day, the fax modem lines were reconfigured for remote access.

How does a company control this? Some perform war dialing on their own exchanges. Others have switched to digital telephony in their business—a standard modem doesn't work on an ISDN line. Should telephone companies supply reports of digital usage on corporate exchanges? The telephone switches could detect and note incoming and outgoing digital usage—both fax and computer modem—and summaries could be reported on the monthly bill.

A company can have better control over its firewalls, which are usually highly visible, and over interconnections to business partners. But the latter can be numerous and haphazard, and are often installed quickly (time-to-market concerns) and with little thought given to security issues. We once ran an authorized *ping* scan of Lucent's intranet, and got an irate call from Southwestern Bell. Investigation showed that the packets ran through our link to AT&T, and through AT&T's intranet to their extranet connection to Southwestern Bell. (These links were an artifact of the AT&T/Lucent corporate split. This particular problem was fixed.) Does your security policy include trust of your business partner's business partners?

Our point is that a large intranet is probably not as secure as you think it is. Large companies have many employees—a larger barrel is likely to have more bad apples. A large number of hacking attacks are made by insiders.

The security of an intranet bears on the security policy of the corporate firewalls. If Bad Guys can get in relatively easily, or are already there, then we don't need to implement quite as robust a firewall. We can concentrate a bit more on the convenience of our users, and a little less on the security grade of the firewall. This leads to greater performance and ease of use, while still keeping the casual attacker out of our intranet.

Given that most companies do not strip-search their employees when they leave the building, we are freer to provide commensurate security through the Internet link.

13.1 Intranet Explorations

The cartography of the Internet has been studied and explored in a number of ways since its inception. A summary of recent projects may be found on Martin Dodge's Web pages.[1] These tools can also be used to explore intranets by companies with access to these nets.

Maps of these networks can reveal a number of pathologies. Figure 13.1 shows a few unknown network pieces in one well-run network. The map in Figure 13.2 shows a routing leak: A dual-homed host routing company traffic to some external points that should not have been reachable. Such maps can indicate intranet connections that should have been severed in previ-

1. See `http://www.cybergeography.com/`.

ous divestitures, or connections through business partners or acquisitions that should have been controlled.

How tight are company intranets? The results vary widely, with the sorts of companies you might expect generally, but not always, doing a better job. Some interesting statistics are shown in Table 13.1.

13.2 Intranet Routing Tricks

If a host can't be reached, it is much harder to hack it. The hacker must run through a third party, utilizing transitive trust, and this can complicate things. We can play tricks with packet routing that can be easy and quite effective at hiding hosts.

One trick is to use unrouted or misrouted network addresses. Companies that have avoided direct IP connectivity with the Internet have been doing this for years, sometime to excess. If there is no direct IP connectivity—they use application- and circuit-level gateways only—they can run their own Internet, complete with root DNS servers and their own address allocations. We know of one company that assigned a separate /8 network for each state in the U.S. where they do business. It made allocation easy, though rather sparse.

We don't recommend this approach on such a large scale, because the company will eventually merge with some other company, and addressing excesses will become a major IP renumbering problem. Futhermore, they may have to rely solely on network address translation should they ever choose to use an IP-transparent gateway or set up a joint-venture DMZ.

But for small networks, it might make sense to misuse a little address space. One of us has a static /28-sized network at home, and needs some private address space for non-Internet hosts, like a printer or doorbell. The correct solution is to use some RFC 1918 address space, but in this case, the home network was doubled to /27. The extra 16 IP addresses are in use by someone else in the same ISP, so we *black-holed* some of their address space, but it is extremely unlikely that we would ever want to connect to those particular hosts.

Black holing can become a serious problem, and we know many companies that had to fix these problems when they went to IP-transparent gateways. The /8 networks that had been chosen and used nearly at random in the old days had to be completely renumbered.

Collisions can be a problem even if a company has faithfully used the RFC 1918 address space in this way. When companies merge, their address spaces are likely to collide, again requiring renumbering. It would be nice to pick RFC 1918 address space that is unlikely to be in use by future merger partners. Figure 13.3 offers some data that may be of some statistical help.

We can also use encrypted tunnels to allow outside users onto parts of our internal network. The tunnels can direct telecommuters and business partners to particular hosts, without giving them the run of our intranet. Check these carefully, though: It is easy to misconfigure a VPN tunnel. And this can cause the same problem of address-space collision: Whose 10.1.2.3 do you want to visit today? Life should improve with IPv6, when it will be easy to get unrouted but globally unique address space.

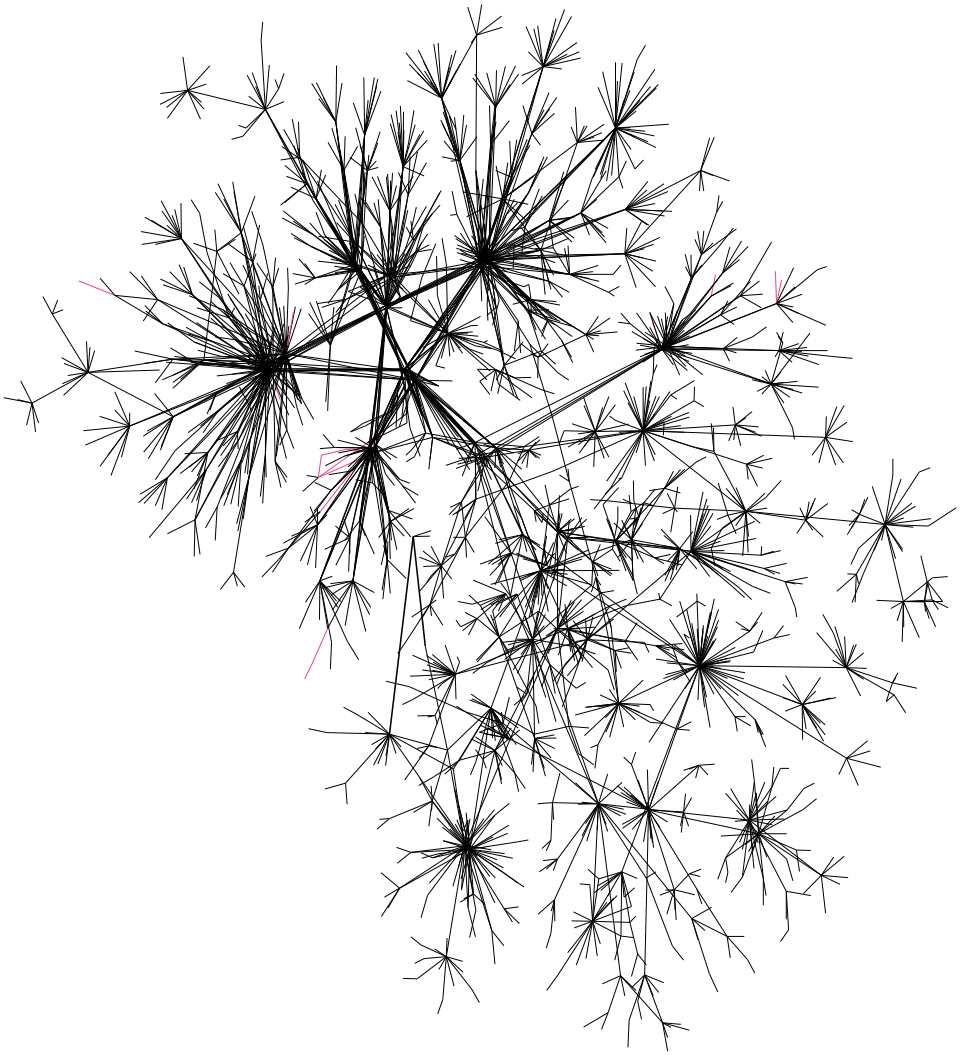

Figure 13.1: Most companies have an official list of the networks in their intranet. This list is almost always incomplete. In this especially well-run network, only a couple of links, shown in bold, were unknown.

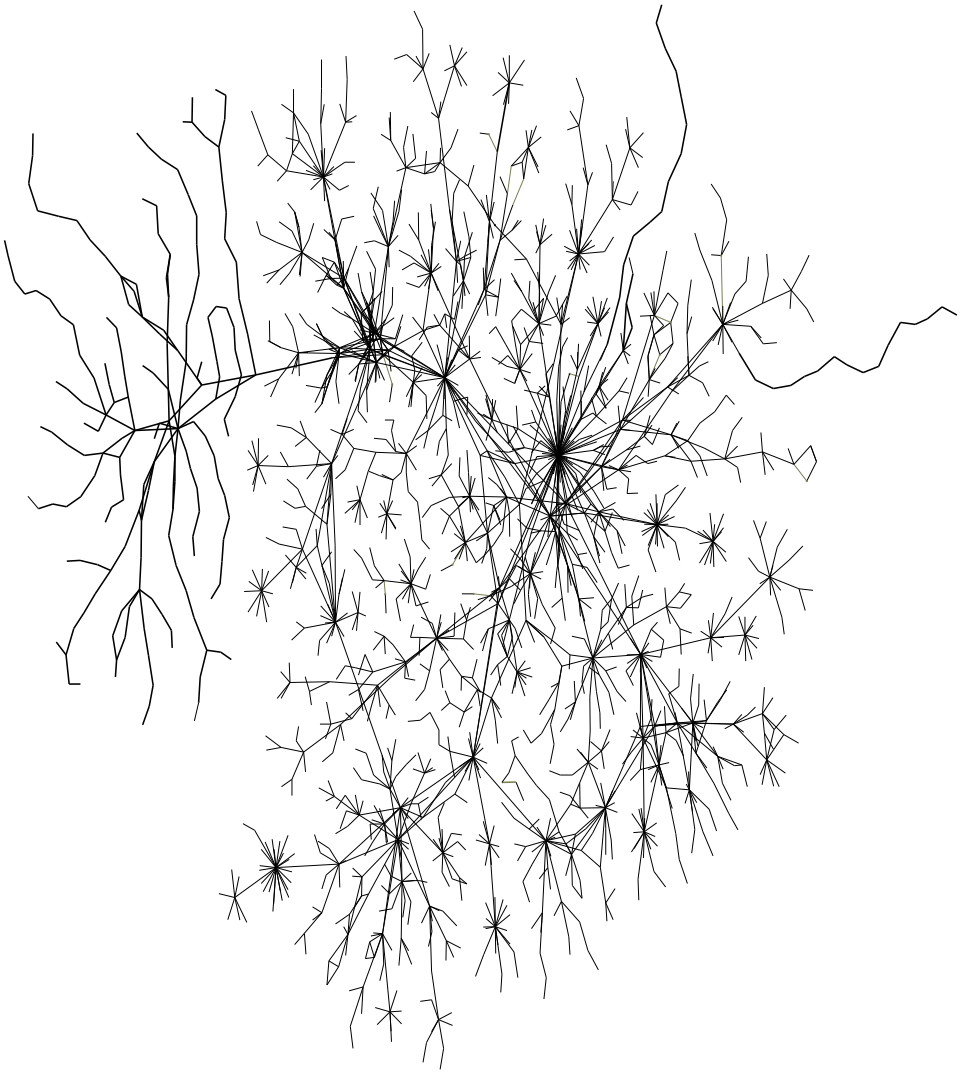

Figure 13.2: This intranet has several *routing leaks*, hosts that announce external routes into the intranet. The sections in bold lines are paths to "intranet" destinations that traverse the Internet, *i.e.,* are outside the intranet. These leaks are not very common and are generally easy to fix.

Table 13.1: Some interesting intranet statistics. This data was summarized from (authorized) scans of a number of Lumeta customers' networks.

Measurement	Range
Number of IP addresses found on the intranet	7,936 –364,171
Potential number of hosts defined by the list of "known" intranet CIDR blocks. Some companies allocate their space more frugally than others, which can ease network management and future network mergers.	81,340 –745,014,656
Percent of all the routers discovered on the intranet that responded to SNMP community string `public`. Most companies want this value to be 0%	0.14 %–78.57 %
Percent of all the routers discovered on the intranet that responded to common SNMP community strings other than `public`.	0.00 %–31.59 %
Number of hosts in the intranet that appear to have uncontrolled outbound access to the Internet. Some companies have policies prohibiting this	0–176,981
Number of hosts that accept UDP packets from the Internet (*host leaks*,) and also have access to the intranet. This violates nearly all corporate security policies. Such hosts are often home computers with tunnels to corporate intranets. They may also be running personal Web sites. Some have been gateways for hackers into corporate networks	0–5,867
Percent of hosts running Windows software. This is a rough statistic based on crude TTL fingerprinting.	36.45 %–83.84 %

13.3 In Host We Trust

We need firewalls when the hosts cannot protect themselves from attack. We also use them to provide an extra layer of protection around hosts and network regions that are supposed to be secure.

Traditionally, firewalls have been used to protect organizations from attacks from the Internet. The corporate gateway required the first firewall, and that remains an important location for the security checks that a firewall provides. The central location provides a focal point for implementing security policies efficiently.

Alas, this approach doesn't work very well anymore. The "internal" community has generally grown vast. In many companies, it can span many continents and administrative domains. Holes in the perimeter abound, from rogue employees, business partners, misconfiguration, tunnels, and legacy connections beyond the memory of network management staff.

Firewalls are used in more locations now. We find them in individual clients, between administrative boundaries, and between business partners. Though they can be inconvenient, firewalls can make an organization's network more robust in the face of successful attack. Firewall bulkheads can protect various corporate communities from security failures elsewhere. This is a lesson learned from the design of naval ships.

Most companies limit the use of internal bulkhead firewalls. A very common location is between the main corporate network and its research arm; these two groups often have different security policies, and sometimes mistrust each other.

Even in small companies, firewalls sometimes separate different tiny divisions. In some small companies, the developers might have a small collection of UNIX-based hosts with strong host security, but the sales and management teams may insist on using more convenient and more popular— but less secure—operating systems. (In one company we know of, the e-mail service for the UNIX hosts improved during the several days when the Melissa worm took out the production corporate e-mail service.)

With really strong host security, you may be able to skip the firewall altogether for a very small community of trusted hosts. But beware—the community may still fall if the trusted network services contain holes.

Ideally, a community behind a firewall shouldn't include more than about 40 hosts. Put another way, it's hard for a single firewall to protect a domain larger than that controlled by a single system administrator. Beyond that, it becomes easier for connections and security problems to escape the notice of the administrator. We realize that 40 is quite a small number, but we do see trends heading this way. Some banks now have hundreds of discrete firewalls, with a correspondingly large administrative management load. Conversely, we think that this extra overhead has purchased a great deal of extra security. A number of companies now offer mechanisms for administering a large number of firewalls. These attempts are promising, but be careful to protect the central administration site, and be careful not to install the union of all desired firewall openings.

From a security point of view, we see three levels of host-based security:

1. A small core of trusted hosts are rigorously locked down. They contain the master password or other authentication files, master binaries, and possibly console-only access. They have a trusted time source, and may serve as a drop safe for important log files. They may offer *ssh*

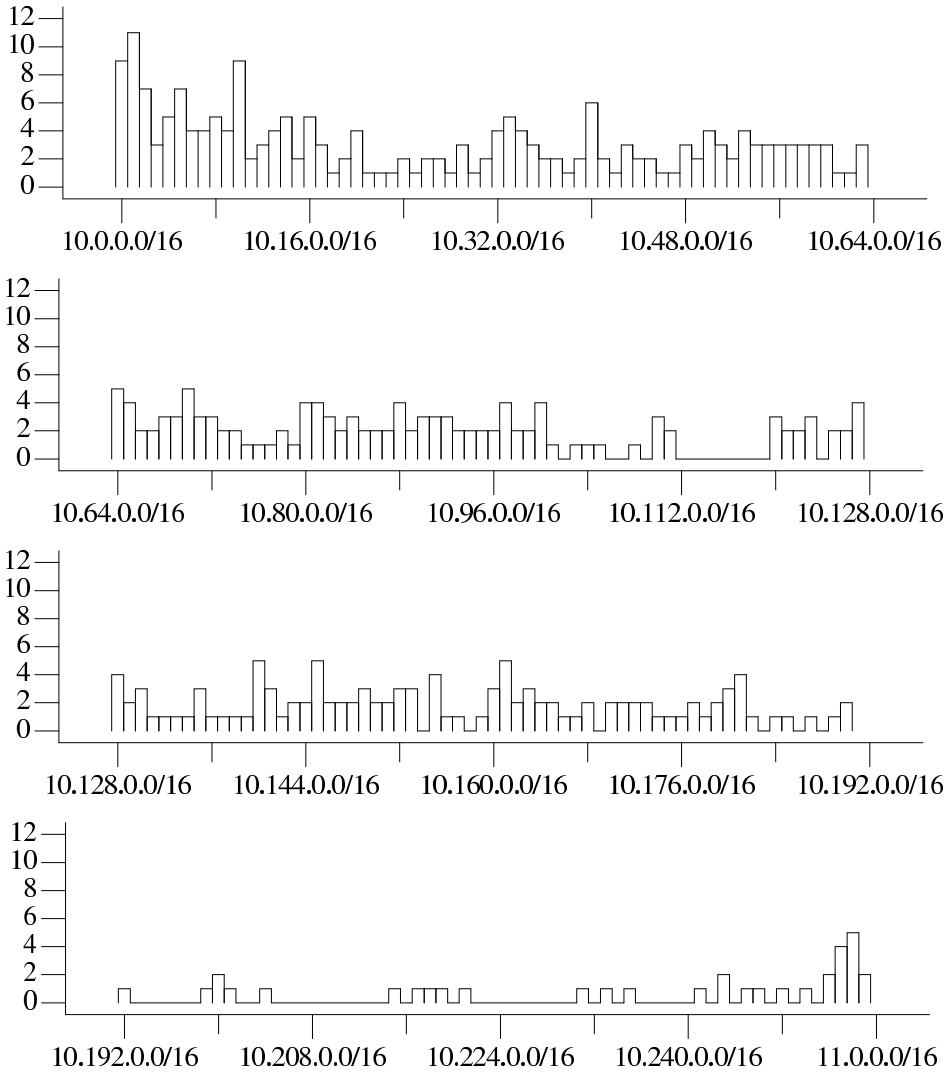

Figure 13.3: RFC 1918 address usage on over a dozen large corporate intranets, at the /16 level. If one chooses an unpopular RFC 1918 address, there is less likelihood of a collision in the case of a corporate merger.

service for a few administrators, but perhaps shouldn't. They may also offer dial-up access with strong authentication (but see the sidebar on page 256). If one of these machines is compromised, the game is over. (There is a trade–off here between emergency availability and security. Yes, these machines should be secure, but if 24x7 availability by skilled personnel is needed, you need to weigh the risks of *ssh* against the risks of whatever ad hoc mechanism will be installed at 3:00 A.M. on a winter day when the Miami site needs be repaired by a snowed-in administrator in Buffalo.)

2. The second level of host security uses hacker-resistant systems that are not keystones of the entire network. These hosts provide services that are important, even vital, but their compromise doesn't jeopardize the entire network. These hosts may run POP3 or IMAP servers, Apache, Samba, SSH, and NTP. Ideally, these services are jailed and/or relegated to a DMZ so that a server weakness won't compromise the other services.

3. Untrusted hosts comprise the third group. These hosts run software that we have little confidence in. They reside at the convenience end of the convenience/security spectrum. They often run out-of-the-box commercial software installed by unsophisticated users. If one or more of these hosts are corrupted, our gateway and basic services remain uncorrupted.

To date, Windows hosts fall into the third category, in our opinion. We do not know how to secure them, or even if it is possible. Some claim that Microsoft servers can be secured to higher levels by applying a long list of configuration changes, moving the host from convenient toward secure. We think the market would welcome machines that are configured for tighter out-of-the-box security.

Microsoft is not alone in this: Most UNIX hosts traditionally came with a lot of dangerous services turned on by default. A number of distributors in the Linux and BSD-UNIX fields have addressed this in a useful way: *no* services are turned on by default.

13.4 Belt and Suspenders

A paranoid configuration, for an application or circuit gateway, is shown in Figure 13.4. This is the kind of network layout you can use to protect the crown jewels, perhaps your payroll systems. In this scheme, which we call *belt-and-suspenders*, the gateway machine sits on two different networks, between the two filtering routers. It is an ordinary gateway, except in one respect: It *must* be configured not to forward packets, either implicitly or via IP source routing. This can be harder than it seems; some kernels, though configured not to forward packets, will still do so if source routing is used. If you have access to kernel source, we suggest that you rip out the packet-forwarding code. The outside router should be configured to allow access only to desired services on the gateway host; additionally, it should reject any packet whose apparent source address belongs to an inside machine. In turn, the gateway machine should use its own address filtering to protect restricted services, such as application or circuit relays. The inside filter should permit access only to the hosts and ports that the gateway is allowed to contact.

The theory behind this configuration is simple: The attacker must penetrate not just the packet filters on the router, but also the gateway machine itself. Furthermore, even if that should occur, the second filter will protect most inside machines from the now subverted gateway.

Should You Trust a Private Dial-up Line?

We admonish people not to rely solely on in-band administration of important computers. In-band signaling has obvious problems—for example, how do you fix a router over a network if the network is down because the router needs reconfiguration? In-band signaling used to be a security problem in the telephone system, allowing people to whistle notes that could give them free telephone calls.

Out-of-band access to a network element like a router usually implies a telephone link to it, using a modem. If the network is down, the phone system is probably still working (though this assumption should be checked for extremely vital equipment.) Can we trust the telephone system?

Certainly the router must be minimally protected by a password. Modems are easily discovered by "war dialing" or information leaks. One cannot rely on the secrecy of the telephone number.

Cleartext passwords on the Internet are subject to simple eavesdropping. Is this a threat on a telephone system? The technology is different, and the expertise is less common, but eavesdropping is possible on phone connections, and it doesn't require a man in a van with alligator clips outside your home. Governments have this sort of access, as do telephone company workers, and there are known cases of such abuse. And modern phone switches can implement a seamless phone tap easily, given administrative access to the phone switch. Hackers have obtained this kind of access to switches for over two decades.

These attacks are certainly less common than the typical Internet attacks described in this book, and the expertise is less widespread.

Therefore, as usual, the answer depends on your threat model. Who are you afraid of? How motivated are they to break your security? What will it cost you if they do? Challenge/response authentication can raise the barrier, but the highest security is still strong physical security and on-site, console-only access.

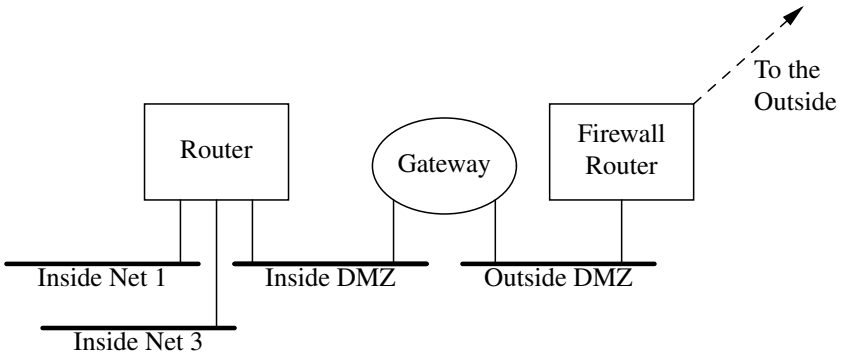

Figure 13.4: A "belt-and-suspenders" firewall

13.5 Placement Classes

In this section, we discuss four different "placement classes" of firewalls. Different organizational situations demand different locations and types of firewalls.

The first placement class corresponds to a large corporation. These are large installations whose firewalls utilize all of the bells and whistles. Typically, these will have a fancy GUI, a hot spare, a DMZ, and other expensive attributes. More than one DMZ might be used for different groups of semi-trusted machines. One of them might house Web servers, while another could be used for experimental machines. The goal is to isolate them from each other. After all, these machines are more exposed, and you want some way to protect them from each other.

This is the scenario in which you're most likely to want a "traditional" firewall. This firewall will likely be your best-administered one; however, it often has to be too permissive, as it has to allow in everything that anyone wants. Do your best to resist temptation here; when you do punch holes in the firewall, limit the legal destinations, and *document* everything, including the person and organization who requested the hole. Make sure the holes expire after not more than a year; six months is better. Renewal should require more than a *pro forma* request.

A second placement class is the departmental firewall. Large organizations have complex topologies on the inside, and different departments have different security needs and varying connectivity requirements. A good departmental firewall should block, for example, NetBIOS and NFS. These protocols are needed within a department, so that employees can share work more easily, but there is rarely much need for these protocols to cross departmental boundaries. If such is needed, an internal VPN is a better idea. Generally, router-based packet filters will suffice as departmental firewalls; it is reasonable to make compromises here toward connectivity for the sake of simplicity. DNS, for example, should probably be allowed between departments. Again, documentation and rule expiration are good ideas.

If your corporate security group has sufficient resources, it should build (and test) some sample rulesets. As we've noted, coming up with a set of rules that is actually correct is a nontrivial exercise.

There are also cost considerations. Most organizations probably can't afford full-fledged firewalls for each of their departments. If a packet filter won't do, a spare PC running Linux or one

of the open source BSDs is almost certainly sufficient, though many departments do not have the system administration cycles to spare.

Past that, individual hosts should be armored. The details of what to block are discussed in Chapter 11; what is of interest here are the criteria for deciding what to block. Different machines require different types of filters. A PC in an office environment should not block Windows file sharing and printer sharing, if they are needed to get the job done. Conversely, given the experience of Code Red, where people did not even know they were running Web servers on their machines, a default of blocking incoming port 80 on users' desktop machines seems like a good idea. As with all firewalls, at the host level it is a good idea to filter out services that are not used. This is even more important for machines that sometimes live on semi-trusted networks, especially road warriors' laptops. Armoring the host is sometimes not necessary for a general corporate machine. However, if a home machine is used for telecommuting, and the kids have another machine on the home LAN, it's a good idea to turn on the host-level firewall to guard against the Things that have infested the kids' machine. (If your kids are deliberately trying to hack your machines, you have other problems, which are well outside the scope of this book.)

The final placement class is what we call a "point firewall." This is generally a packet filter, and is part of a large and complex collections of networks and hosts that operate within a large framework.

Consider a large e-commerce site as an example. Many different pieces have to communicate, and there is a wide range of policies among them. The Web server needs to communicate with the inventory, order-taking, customer care, credit card verification, and billing machines, and probably many others, but the nature of this communication is very restricted. The order tracking system may need to do database queries to the inventory system, and it may need to generate e-mail to customers; however, there is no need for anyone to log in between these machines. E-mail retrieval is even less likely.

All of the different pieces can be laid out in a large, complex diagram, and the relationships among them defined. In each case, a firewall should be placed between the entities, with carefully tuned holes that allow only the minimum necessary traffic. If the Web server itself is outsourced, the hosting company handles other sites, some of which might even be your competitors. It is important to allow access only to the Web server, even if the requests are coming from the same LAN. Similarly, there may be a small and select group of people on the corporate network who need to access the sensitive database used by the Web servers, but others should not be able to.

Sometimes, as in the case of the content supplier, the best way to set up a firewall is to create a packet filter that allows in only VPN traffic. A second packet filter should be created *after* the VPN termination, to restrict what services even authorized users can reach. This way, you can ensure that only a few people come, and that they only talk certain specific protocols, and only to a particular group of machines.

Designs of this sort tend to be highly specific to the project in question. Space prohibits a detailed treatment here; it is a subject for a book unto itself. But one point should be stressed: In many such setups, by far the most dangerous link is a small, obscure one in the corner—the one that connects this massive production system to your general corporate intranet. *That* link needs to be guarded by a very strict authentication system.