

Brought to you by:



Buy this book at:



! " # \$ % &
#!
' % # (
) %
* "
& % &
* # &
% # + &
, &
' # &
+ % " % &
- &
, (/ # * &

“private” public networks and the extremely “public” Internet as we think of it is anything but distinct; because these network service providers are frequently also Internet service providers, the confusion is further compounded.

If you’ve done a search on the World Wide Web for the keywords virtual private network or VPN, you’ve probably even discovered that for years companies like Pacific Bell have been marketing an extended telephone service using the same terms.

Thanks to all this confusion, it can be very hard to understand exactly what qualifies as a VPN, how a VPN can be implemented, and exactly what a VPN can and cannot do for you. That’s what we’ll explain for you here.

1.1 What Is a VPN?

Very simply put, a *virtual private network* uses a public network’s infrastructure to make the connections among geographically dispersed nodes, instead of using cables owned or leased exclusively for one single network’s use, as is typical for a wide area network (WAN). To the user, a VPN looks just like a private network, hence the *virtual* in its name, even though it is sharing a web of cables with the traffic of hundreds or thousands of other users at the same time. It has all the characteristics of a private network—limited access to only authorized users, for example—even though it is sharing the same public infrastructure with other users. Another way to describe it is that a VPN is a logical local area network (LAN) that connects an organization’s geographically dispersed sites in a way that makes them all appear to be part of one single network.

There are a variety of public networks that can be employed to make a VPN’s connections, but the most prominent and most public network available is, of course, the Internet. Because the Internet is everywhere and the Internet is where most of the VPN development is taking place, and because it is, as we’ll see, the most ubiquitous and cost-effective medium for a VPN, we’ll concentrate on VPNs running over it in this book. We will devote Chapter 4 to VPNs implemented through other networking services, and we will explain the differences in detail at that time, but since the Internet is the predominant medium and the technology is essentially the same regardless of the network being used, the Internet is where we will concentrate our discussion.

To illustrate how a VPN differs from a typical WAN, let’s look at a leased-line network, as shown in Figure 1-1, and then show how a VPN differs from it. For the sake of simplicity, this is only a three-node network, a company headquarters and two branch offices linked together with three leased lines.

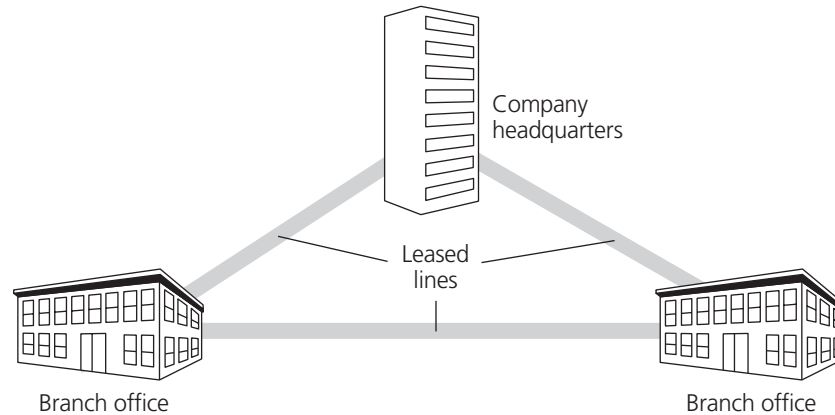


Figure 1-1 A typical leased-line wide area network.

Each office has to have a cable connecting it to each of the other facilities. Another arrangement would be to have the leased lines go through a hub, perhaps in the company headquarters.

Either way, with this type of network, the company actually owns the cable or pays a monthly fee for every mile of cable connecting its facilities, whether that cable is in use 10% or 100% of the time, whether it is being used to capacity or only a fraction of the capacity they're paying for. The costs escalate with every mile that separates the offices and with every node that is added to the network (requiring more strings of cable to connect it to the rest of the organization). Economies of scale are limited to what you can negotiate with the line provider, who is trying to recover from you all the costs for those cables.

Your message uses only those cables to get from point to point; there are no detours. Send a packet of data in one end of the cable and it travels right down that cable to the destination. It works much the same way the LAN connecting your office to the file server on your LAN does. This is a nice, secure connection, but it also means that if the cable is cut, perhaps by a backhoe operator putting in an irrigation line in an Iowa cornfield, that connection is down for the count. It will stay down until either the break is repaired or the traffic is rerouted manually around the break (if your agreement with the service provider offers that guarantee).

In a similarly simple three-node VPN, as Figure 1-2 shows, leased lines are dispensed with in favor of connecting each site to a public network. Instead of the hardwired pipeline between nodes of a standard wide area network using dedicated connections, the connections of a VPN are made through the web of cables, what is often described as the “cloud,” of a public network such as the




Figure 1-2 The same company using a VPN instead of leased lines.

Internet. Each office requires a single connection, most commonly a leased line and commonly referred to as a *local loop*, to the nearest public network *point of presence* (POP). That POP may be only a few feet away or it may be miles away. From that POP the data is carried by the web of connections—the cables, routers, and switches that make up the public network—to the POP serving the destination office, then through the local loop at that end on to their final destination.

As you can see, the connections—the local loops—between your company's offices and the public network are dramatically shortened. They can even be the "dry copper" connection provided by your local telephone service, perhaps an Integrated Services Digital Network (ISDN) connection. The costs for these short connections are correspondingly lower.

Within the cloud, however, as opposed to the leased-line design, there is no one single connection between point A and point B. Instead there is a web or matrix of cables connected by routers that the messages travel through. By using a public network—the Internet, for example—a network shared by hundreds or thousands or millions of other users, the cost of all those miles of cable is shared. We'll see further on in this chapter that this sharing can produce impressive savings.

It is also much more fail-safe than a single leased line, since a message can take any one of a number of different routes to its destination. It provides a redundancy, a safety net, that virtually guarantees that the traffic will continue to flow reasonably smoothly. If one cable is cut, the message will simply be routed automatically around the break to its destination, a service that is not generally available with a leased line.

1.2 What a VPN Is Good for and Why You Should Consider Building One

There are several uses for a VPN. It can be an extended intranet, connecting geographically distant facilities into a cohesive network. It can also be an extranet, linking, for example, customers and suppliers for increased efficiencies, such as electronic data interchange (EDI). Looked at this way, a VPN can do virtually anything that a more traditional leased-line WAN can do. In fact, so far it doesn't seem to offer services much different from any WAN.

But there is a third service that a VPN can offer that no leased-line WAN can offer, and that is in providing remote access services. A VPN lets road warriors with their laptops connect into the home office through an Internet service provider, riding through the public Internet to log on to the office network, rather than running up long-distance charges by dialing up to a remote access server thousands of miles away. As we'll see, that offers potentially impressive savings. While a VPN as an extranet or intranet offers some cost efficiencies over the typical WAN, the savings produced by using one for remote access are significant.

Hence the excitement that has developed over VPNs. Building a VPN would seem, at first glance, to be simple common sense. Why not take advantage of an existing infrastructure for the connections, instead of going to the expense of stringing your own cable or paying someone else to drag fiber through conduit to tie your facilities together? Or why go to the expense of leasing dedicated connections when they may only be used to a fraction of their capacity or for only a fraction of the time?

It does make sense, but as you'll see, there are downsides to VPNs. But before we take a look at the potential negative points to VPNs, let's see what the potential benefits are. As we said, there is a powerful logic to using an existing infrastructure to connect your facilities, rather than building your own.

The claims made for VPNs make them sound like the greatest invention since the electric light. The primary advantage cited is that a VPN is vastly less expensive than a network using leased lines. As we already mentioned, the VPN is also claimed to be more flexible and scalable, compared to a traditional WAN. Then, too, by using the international resources of the Internet, the vendors say that a VPN can offer connectivity virtually anywhere in the world. Finally, you'll hear that a VPN is an extremely cost-effective way to service a mobile workforce of telecommuters and road warriors.

To a degree, believe it or not, it's safe to say that most of these claims are true. Fortunately, some of them, such as actual dollar savings, are even measurable, while others are less tangible but no less real.

1.2.1 Economies of Sharing

It's a fact that a VPN escapes the cost of leasing the cables to connect your network. By using an existing public network for your VPN, you are sharing the cost of that public network with all the other customers. The cost of the public network is spread over a large customer base. You're not paying every month, by yourself, for every mile of each leased line, whether it is fully loaded 24 hours a day, 7 days a week or not. In most cases you're paying a flat, monthly fee that is a fraction of what you would pay for leased lines providing the same service.

Compare it to your personal telephone service, for example. While you pay a base charge for the local loop between your home and the telephone company's central office a block or two away, whether you are using it or not, you do not pay for every inch of cable between your home in Poughkeepsie, New York, and your daughter's dorm room at college in Palo Alto, California, whether you are using it or not, at least not directly. That cable is shared by thousands of callers, each paying perhaps a dime a minute for the time they are actually "online."

In the past, on long-distance telephone circuits, one call used one circuit, which was one pair of wires that could be traced from your home in Poughkeepsie to your daughter's dorm in Palo Alto. When you hung up on that call another took its place on the long-distance trunk, so at least you were only paying for time used. Today calls are multiplexed on that long-distance circuit, with the "silence" between words being filled with parts of other conversations. This spreads the cost of the wire over more than one customer, allowing each of them to enjoy the benefit of lower long-distance rates. From three dollars for 3 minutes the rates have dropped to three dimes for 3 minutes. But your conversation will still be carried over one circuit between Poughkeepsie and Palo Alto.

A packet-switched network such as the Internet allows even greater multiplexing, and thus greater efficiency, as each message is broken up into packets, and each packet is slotted in with others from other users and routed through a web of connections. No one circuit becomes overloaded, at least in theory, and every circuit, at any second, is efficiently utilized, carrying pieces of perhaps thousands of conversations. It also provides a safety net that a circuit-switched network or a leased line does not. If one link is overloaded or goes down, the traffic is automatically rerouted to its destination. (For a more complete description of how the Internet works, see Section 4.1.1.)

More importantly, the cost of all the fiber and copper and switches and routers is being spread over the millions of customers the Internet serves. You are leveraging to your advantage the investment in the hundreds of thousands of miles of cables and the uncounted routers and switches that go into making

up the Internet. Your major expenses are only the cost of that short loop that connects your office to the network access server (NAS) or POP of your Internet service provider (ISP) and your monthly Internet fee. The average price for a leased T1 (1.544 Mbps) connection is about \$1,800. A typical connection from a company's offices to the local ISP's POP costs \$400 to \$500 a month, because the chances are you'll actually use less than a full T1 line to your POP, perhaps even a 128 Kbps ISDN line or a digital subscriber line (DSL) of some sort at an even lower cost of \$50 to \$150 a month. If you're a small operation, your cost may be a monthly subscription for a dial-in connection to your ISP.

The savings can be considerable. According to a white paper by Infonetics Research, a study commissioned by Sun Microsystems estimated savings of from 20% to 47% by switching from leased lines to a VPN. In another analysis, Infonetics estimated savings of 20% to 40% for VPNs serving branch offices and 60% to 80% savings for a VPN serving remote access users. As we'll see later on in this chapter, when we look at the remote access aspect of VPNs, every analysis of VPNs produces similar savings estimates. The experiences of VPN users bolster those findings, as we'll see in Chapter 2.

Another source, *Data Communications* magazine, in their May 21, 1997, issue, ran their own numbers on a VPN, comparing leased lines, a frame relay service (see Chapter 4), and an Internet-based VPN solution (Table 1-1). The sample scenario was to connect three sites in the United States (Boston, Los Angeles, and Houston), plus one transatlantic link to London. All were connected at 64 Kbps. AT&T was the carrier and provided the charges, including local access circuits of 5 km to the nearest POP. Leased-line and frame relay figures were provided by Lynx Technologies, Inc. of Fairfield, New Jersey, a tariff tracking consultancy. Internet figures were based on average monthly ISP charges in the United States.

As the *Data Communications* analysis shows, the frame relay first-year cost is only about 17% lower than the cost for leased lines, but about twice the first-year cost for the Internet VPN. However, much of the frame relay

Table 1-1 *Data Communications* magazine VPN cost comparisons.

| | Leased line | Frame relay VPN | Internet VPN |
|-------------------------------|------------------|--------------------|-----------------|
| Annual charges | \$133,272 | \$89,998 | \$38,400 |
| Installation | \$2,700 | \$5,760 | |
| Four VPN encrypting devices | | \$16,000 | \$16,000 |
| Total cost, first year | \$135,972 | \$111,758 | \$54,400 |

first-year cost is the one-time charge for installation and the encryption devices. The annual charges (operating costs) are about two thirds of those for leased lines, though still more than double the annual charges for an Internet VPN.

By this analysis, the Internet is obviously the most economical choice for your VPN, but for the extra operating expense of the frame relay choice you do get added services that are not available on the Internet, as we'll discuss in Chapter 4. If you need those services, you'll see that, as economical as the Internet is, it is not the best choice for you.

Because of the way telephone charges are computed, the greater the distances and the larger your user base, the greater the savings you'll enjoy. Telephone charges are computed by the call, and the rates increase with the mileage covered. Distance means nothing to the Internet, and Internet service is usually billed at a flat rate, regardless of the number of times you use it or the amount of data transmitted. In Chapter 2 we'll look at some real-life VPNs and see how the savings can stack up in action.

1.2.2 Flexibility

It is true, also, that a VPN offers flexibility that is not available to a leased-line-based wide area network. To add a node to the latter requires leasing a new line, possibly more than one, perhaps even installing some cable. Leases have to be negotiated, perhaps rights-of-way arranged. Routers and switches have to be installed and configured.

Let's go back to our first three-node leased-line network. Your company, Giant Widgets, has long done business with Associated Grommets, a supplier of grommets for your widgets. Your company is flush with cash, the widgets market has been really strong lately, and you decide to buy the grommets factory. Once you've acquired it you want to put it on your existing network. Figure 1-3 shows what happens when you have to bring it into the loop of a leased-line network. Three new lines have to be leased and somehow integrated into your existing system.

Now let's take the same scenario if you're running an Internet-based VPN (Figure 1-4). You've purchased the grommets factory, and it just happens to already have a link to the Internet. They've been selling grommets through the Internet for years, after all. All that's needed is to slide into place the VPN system, generally a hardware box or some software, and they're on your network. Or suppose that instead of buying the Associated Grommets factory you just want to extend your VPN to it, turning your VPN from an intranet into an extranet. The scenario is essentially the same.

If you already have an Internet link on the facility you want to link to the VPN—for a Web site, for example, or email—getting your VPN up may be as




Figure 1-3 Adding one node to a leased-line WAN.

simple as installing a server and some software. For a small-scale VPN, if you're running Windows NT 4.0, adding a node to an existing VPN can be as simple as implementing the Point-to-Point Tunneling Protocol (PPTP) software that comes with the operating system and establishing a dial-up link to the Internet. A mobile user running Windows 98 can use PPTP to connect through the Internet for remote access to the home network from wherever he or she is. (But be forewarned: PPTP as implemented by Microsoft presents some serious security concerns, as we'll discuss in Chapter 7. Fortunately, it is scheduled to be replaced in Windows 2000.)

If speed is a requirement, you may need to call up your telephone provider and request a digital loop to their Internet service. If either ISDN or digital subscriber line (DSL) service is available, the existing copper that serves your telephone system may be put to use, saving the cost of stringing new cable, with only some terminators to be installed at your end. Another option worth exploring is the availability of Internet service through the cable TV network in your area, using a cable modem.

If your VPN is to serve as an intranet, connecting local area networks already up in distributed offices, it will require a router to connect you to the Internet and a firewall for protection from hackers. These may even be combined in one box that also integrates the VPN features along with its other functions. It still doesn't require leasing a thousand miles of dedicated cable.

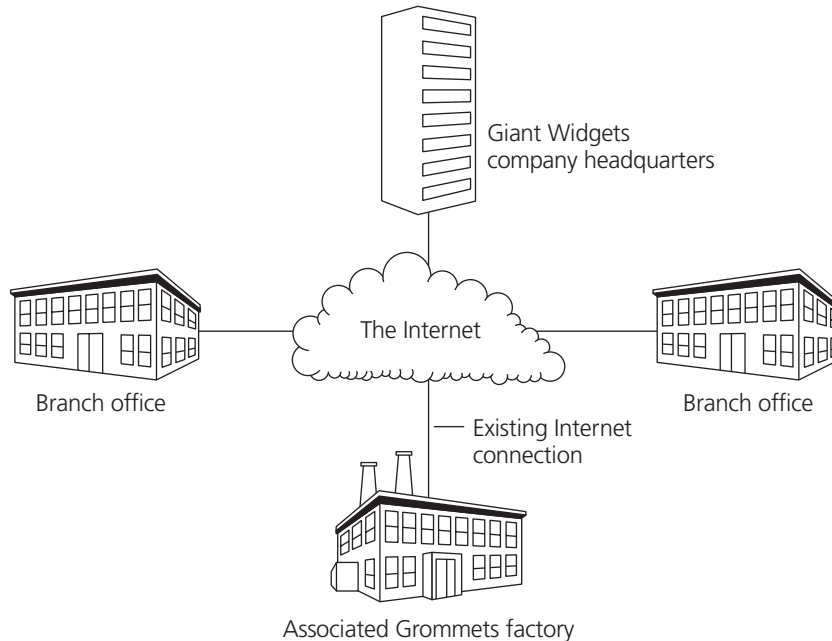


Figure 1-4 Adding the Associated Grommets factory to your existing VPN. With an existing Internet link, it is almost plug and play.

So, turning your VPN into an extranet by adding a supplier or customer can be accomplished by negotiating a compatible VPN connection between you and them. As we'll see, this will not be as simple as plugging in a toaster, but it will still be less expensive than negotiating a dedicated connection. As the interoperability of VPN products improves, the problem will ease further.

1.2.3 Worldwide Connectivity on a Budget

With a VPN you can have a network node virtually wherever there is an Internet POP, and today the Internet is virtually everywhere, even on Tristan da Cunha, the remotest inhabited island in the world. For the cost of an Internet connection, a modest company that has overseas sales representatives can have them directly on their network, a luxury previously limited to major corporations with very deep pockets.

Even for a large corporation, the ubiquity of the Internet may prove to be a significant advantage. Frame relay and ATM networks are not as widespread as the Internet and are more expensive. International leased lines are costly and may not be available in areas where the Internet has a presence, or such lines may be available only from a very expensive government monopoly.

Also, the Internet is, to a great extent, oblivious to national boundaries. Certainly there are those countries where Internet service is limited or access is restricted, but the same or greater problems are likely to be encountered using any other form of networking as well.

As we'll see, the Internet may not be the solution to your VPN needs. However, in terms of accessibility it is hard to beat.

1.2.4 The VPN and the Mobile Workforce

It is in remote access services (RAS) that VPNs show the greatest savings. In the usual scenario, a road warrior must dial into the public telephone network to reach his home base remote access server (Figure 1-5). From his motel room in Los Angeles he immediately begins running up long-distance charges to the home office in New York City, where he goes on the network, perhaps to check his email. Furthermore, instead of logging off to read and respond to his messages and then logging on again to upload his answers, he is likely to stay online, while the telephone company counts up steadily increasing profits at your expense.

If there are any advantages to this they are only that your cost per minute for an 800 number call may decline, perhaps to as low as \$0.07 per minute as



Figure 1-5 Accessing the network via dial-up service.

the time used goes up, and it may make it easier for the accounting department to keep track of the costs.

On the other hand, if your VPN supports remote access, that road warrior in Los Angeles can dial in to a local Internet POP there, perhaps through GlobalNet or WorldNet, perhaps using a local or regional ISP with whom he has an account. Once that connection is made, the Internet cloud connects him to the home office, avoiding all those miles at long-distance rates (Figure 1-6). Your cost is only the flat monthly fee he pays for his Internet service no matter how long he remains online, plus what the motel charges for a local call, not the per-minute charges for a long-distance connection. Even a series of 800 number calls at \$0.07 per minute are not likely to be lower than the cost of a \$20 per month Internet account with no time limits.

In terms of flexibility, a VPN can serve a mobile workforce or telecommuters in a way that no leased-line network can. A salesman in Dubuque can plug his laptop into the RJ 11 telephone jack in his motel room and dial the ISP that he has a subscription with, perhaps the local PSINet or WorldNet number. If he's staying in one of the upscale hotels that offer direct Internet connections, he can take advantage of that.




Figure 1-6 A remote user in Los Angeles connects to the home network in New York via the Internet through a secure VPN.

Once he's logged on to the Internet he can log on to his home network, turning his Toshiba laptop into a node on the network. He can check his email, upload new orders, review the status of last week's orders, download a new catalog and price lists, leave email to his assistant, check on his commissions, and log off. He can do virtually anything from that motel room that he could do from his workstation at the office.

The savings for such remote access can be even more impressive than for those in the more static intranet or extranet scenarios. VPN vendor TimeStep has created a cost analysis that is worth looking at because it puts some real numbers into the scenario (Table 1-2).

Notice that the capital costs for the Internet scenario are about six times that of the dial-in scenario. On the other hand, the monthly operating costs for the Internet scenario are only about one fourth those for the dial-up scenario. In less than 3 months, the difference in the capital costs has been recovered. The first-year return on investment is over 345%. Granted, these are numbers generated by a VPN vendor, but your own analysis is likely to return similar results.

More recently, Cisco Systems has analyzed the savings achieved by outsourcing a VPN-based remote access system using their Access VPN technology versus maintaining a dial-in remote access service in-house, using modem ports. As Table 1-3 shows, the estimated savings are 38% using this solution. Naturally, an in-house VPN solution is possible, too, with some trade-offs since you take on more management costs.

By their figures, a dial-in port for remote access costs about \$1,500 a year, including local access line charges and the cost of a modem. Assuming a conservative ratio of users to ports of 5 to 1, the cost per user for such a port is about \$300 per year. On the VPN side, a 1.024 Mbps dedicated line port and interface, including initial purchase, facility costs, and maintenance, runs about \$2,500 per year. However, such a port can serve 20 users, resulting in a per-user cost of \$125 per year, less than half the cost of the dial-up port.

If, instead of using the dial-in ports Cisco proposes in their chart, the remote users can connect to the VPN by making a local call to an ISP, the savings are even greater, since access charges are the minimal cost of the local call and a subscription with an ISP at \$19.95 per month.

In the real world, the potential savings can be impressive. Mazzio's Corporation runs a VPN to serve their restaurants that we'll discuss in more detail in Section 2.3.1. They discovered they had more than 20 company users dialing in and running up a phone bill of \$1,800 per month. By bringing them onto the VPN, that monthly cost will be eliminated.

Forrester Research set up a study comparing a leased-line network serving 2,000 remote users with a VPN serving the same number. As Table 1-4 shows, the 60% savings were about the same as predicted by other studies.

Table 1-2 TimeStep cost analysis of VPN savings with a mobile workforce.

| Direct dial-up scenario | | | | Internet scenario | | | |
|----------------------------|----------|-----------|----------------|--|----------|-----------|-----------------|
| Item | Quantity | Unit cost | Extended cost | Item | Quantity | Unit cost | Extended cost |
| Capital costs | | | | | | | |
| Terminal server* | 10 | \$550 | \$5,500 | T1 line startup** | 1 | \$3,000 | \$3,000 |
| | | | | Channel service unit/Data service unit** | 1 | \$995 | \$995 |
| | | | | Router** | 1 | \$1,950 | \$1,950 |
| | | | | PERMIT/Director*** | 1 | \$11,995 | \$11,995 |
| | | | | PERMIT/Gate 2520*** | 1 | \$3,995 | \$3,995 |
| | | | | PERMIT/Client*** | 75 | \$99 | \$7,425 |
| Total capital costs | | | \$5,500 | | | | \$29,360 |

| | | | | |
|--------------------------------------|-------|------|----------|---------------------------------------|
| Monthly operating costs | | | | |
| Long-distance charges**** | 1,500 | \$10 | \$15,000 | |
| | | | | ISP charges** |
| | | | | 75 |
| | | | | \$20 |
| | | | | \$1,895 |
| | | | | \$1,895 |
| | | | | \$700 |
| | | | | \$4,095 |
| Total monthly operating costs | | | | |
| | | | \$15,000 | |
| | | | | % operating cost savings |
| | | | | 73% |
| | | | | Payback period (months) |
| | | | | 2.70 |
| | | | | Return on investment (for first year) |
| | | | | 345.70% |

*Average industry cost ("Intranets on the Road"; Forrester Report; June 1997).

**Price from UUNet (August 1997).

***TimeStep's PERMIT/Director is a Windows-based software suite that offers VPN user management, authentication, key management, and administrative tools. PERMIT/Gate 2520 is a tamper-resistant gateway that secures data for transmission over the Internet using a suite of hardware-based encryption and authentication protocols. PERMIT/Client is a software suite available for Windows and Macintosh operating systems that offers VPN services including encryption and authentication protocols.

****Price from MCI (August 1997).

*****Average local loop cost (varies region to region).

Table 1-3 Cost comparisons of an in-house remote access system vs. an outsourced VPN system. Source: Cisco analysis, July 1998.

| | In-house | Outsourced | Savings |
|---|--------------------|--------------------|--------------------|
| Ports and toll-free access | \$957,000 | \$700,000 | \$257,000 |
| Network backbone | \$500,000 | \$450,000 | \$50,000 |
| Staffing | \$440,000 | \$0 | \$440,000 |
| Security | \$185,000 | \$100,000 | \$85,000 |
| 24 × 7 help desk | \$750,000 | \$550,000 | \$200,000 |
| Network management | \$75,000 | \$0 | \$75,000 |
| Totals | \$2,907,000 | \$1,800,000 | \$1,107,000 |
| Savings based on outsourced solution | | | 38% |

Table 1-4 Forrester Research found a 60% savings in a 2,000-user VPN.

| | Private network | VPN |
|-----------------------|--------------------|--------------------|
| T1 lines | \$48,000 | \$68,400 |
| Routers and servers | \$208,000 | \$44,800 |
| Phone and ISP charges | \$2,160,000 | \$1,080,000 |
| User support | \$600,000 | (included) |
| Total | \$3,016,000 | \$1,193,200 |

Notice, too, where the savings are made. The cost for T1 lines is higher, to allow the higher speeds needed for the consolidation to single-line access for the Internet, the extranet, and the intranet VPN. However, this is more than made up for by the consolidation of equipment (routers and servers), the phone and ISP charges, and the cost for user support. Again we see that the major savings come in operating expenses, not that the savings in capital costs from equipment consolidation are insignificant.

The consulting firm Gartner Group estimates VPN savings of at least 50% for remote access. For this reason they predict that by the year 2003, 70% of Fortune 500 companies will use VPNs for their road warriors. Telecommuters can benefit from a VPN as much as road warriors. If you have a VPN, telecommuters don't require a dedicated line between their home office and your headquarters. They can log on to your LAN through an account with their local Internet service provider. For speed, if the ISP offers it, they can use an ISDN or DSL local loop that offers digital speeds over copper telephone lines.

If that's not available, perhaps they will have a cable modem through their local cable TV service.

Similarly, with a VPN, managers who need to be on call at all hours can jump into a problem from home. When that late night call comes in from an irate customer in Tokyo, they can shuffle down the hall in their pajamas, turn on the computer in their study, and use their Internet account to log on to the company server to untangle the problem. A network administrator can log on and sort out a Paris-based user's account problem without making the 20-mile drive to the office at 3 in the morning.

1.3 Every Silver Lining Has a Cloud

All of this makes VPNs sound too good to be true. Putting in a VPN will boost your bottom line, satisfy your customers, turn your weary sales force into a jolly chorus line, and cure the common cold.

Well, not quite. There are challenges to creating a working VPN, especially one that uses the Internet. The Internet, as we all know, is an anarchic maze. It is rightly seen as a hostile environment, a tropical beach with appealing blue waters, which happen to be inhabited by sharks, jellyfish, moray eels, and stingrays. There are reefs and currents that must be negotiated and an occasional hurricane that may blow you on the rocks.

The primary concern on the Internet is, of course, security. For a VPN to work, the traffic it carries must get to its destination safely. It must be protected from sniffing and snooping, hijackers, denial of service, and from other hacker attacks (see Section 3.3 for a discussion of these terms).

As we'll see, VPNs, along with their Internet connections through firewalls and the like, are designed to cope with these hazards. But this also means that a VPN presents challenges that may not be encountered on a more secure internal local area network. There will probably be increased system management loads. These can be handled internally, or they can be outsourced, but either way they are going to result in some increased costs.

With regard to the Internet particularly, there are also issues of quality of service and reliability. The Internet can get bogged down. A fire in a manhole in Chicago fuses some fiber-optic cable into a useless lump and traffic all over the country slows to a crawl. Someone enters the wrong command and routers get confused and send everything off into a void until the problem is caught.

In Chapter 3, we'll explore some of the issues you need to consider before going the VPN route. We'll also go into more technical detail on how VPNs work in Chapters 5–7, but for now we'll take a brief look at how a VPN does what it does and some of the complexities it can deal with in the process.

1.4 How a VPN Works

How does a VPN make a private, secure connection through the very public Internet? What if you're running a LAN already that doesn't use the Internet's IP protocol, or that has IP addresses that are fine for your LAN but that would be banned from the Internet? Just as important, how do you keep unauthorized users from sneaking into the system? And how do you keep all those hackers out there from snooping into your data as it whizzes around the cloud?

To manage the first problem, that of making a secure connection and hiding your LAN's addresses from the Internet's routers, your existing network's headers are hidden using a process called *tunneling*. To keep unauthorized users out of the VPN, a system of user authentication is established, combined with security gateways such as firewalls that guard any Internet connection. To keep hackers from sneaking a peek at your company secrets, the data in transit is encrypted. To make sure someone hasn't tampered with the transmission en route, it is authenticated at the receiving end.

We'll go into the protocols that manage this and the details of how they do it in later chapters. For now, so you can understand some of the terminology, we'll take a quick look at the processes that are used to meet the challenges.

1.4.1 Tunneling

There's plenty of confusion in the terminology of VPNs, and a lot of it crops up right here. The term tunneling implies that, somehow, a pipe—a real, solid, direct route—is established through the Internet, a connection like you'd see in the telephone network, a circuit-switched network.

However, we know that the Internet is a packet-switched network. It simply doesn't work that way. In reality, a single message is broken up into packets, and each packet finds its own way to its destination, bounced from router to router in accordance with the address it carries, the route's routing tables, and the paths available. One packet may make its way from New York to San Francisco by way of Chicago, the next one by way of New Orleans. Once they arrive, the packets are reassembled in the right order to re-create the message.

So there is no such thing, really, as a "tunnel." However, there is a connection of sorts between sender and receiver in that the VPN equipment at each end has agreed on how to communicate (what protocols and security arrangements are to be used at each end). A control "circuit" may exist between them, a series of background messages that travel the Internet the same way as the message packets, taking care of administrative details of the communication, but it doesn't work like the solid link of a telephone connection.

Another term you'll hear in reference to the process is *encapsulation*. The terms tunneling and encapsulation seem to be used interchangeably. Probably the simplest way to differentiate between them is that *tunneling* is applied to the whole process of moving a message through the Internet for a VPN, while *encapsulation* refers to what is done to each individual packet that makes up the message.

In tunneling, each packet, including any existing header it has acquired from the LAN where it originated, is encapsulated—wrapped up, hidden—by a new envelope or capsule that carries the addresses of the source and destination VPN servers. In this encapsulation process, the VPN software, which we'll see may be running in a workstation, a network server, a firewall, or a router, appends a new header with new source and destination addresses to the packet before sending it out on the Internet. The original header becomes nothing more than part of the payload.

The Internet's routers and switches, which only know to look at that first header, are therefore oblivious to any invalid network addresses that may be buried in the packet, perhaps even one using a foreign network protocol, such as NetWare's IPX or Apple's AppleTalk. In this way, a company with NetWare or AppleTalk LANs in widely dispersed facilities may tie them together through the IP-based Internet with the Internet, and the local networks, being none the wiser.

1.4.2 Securing the Data

While tunneling allows non-IP data or data bearing illegal addresses to be moved through the network, it does not secure the data. While the network's routers may be unaware of the contents, anyone can intercept the packets and read and even tamper with their contents unless they are further secured. For security, a combination of encryption, verification, and authentication is required.

Encryption secures the data from anyone who does not have the key to decrypt it. A variety of encryption technologies are used, as we'll see in Chapter 5. Along with encryption there's the problem of making sure that only the right people can get access to the system and decrypt the data—there must be a way that users are *authenticated*, so that they are the only ones who get the key they need to unscramble the data. As we'll discover in Chapter 6, user authentication, and making sure users receive the keys they need to decrypt the data, is a vital part of a VPN, and a major challenge; the larger the VPN, the greater the challenge.

Finally there has to be some way of making sure that the data that is received has not been tampered with while in transit. For data *verification* yet another process is used.

1.4.3 Making the Combination Work

It is the combination of these elements—tunneling, authentication, encryption, and verification—that makes VPNs possible. To the user the process should be transparent. When Alice in her office at Giant Widgets logs on to the company network she should see the server in the former Associated Grommets factory in Malaysia the same way she sees the server down the hall from the cafeteria where she gets her morning coffee and bagel.

For this to be true, all the bits and pieces that make up a VPN's protocols have to work together. The Associated Grommets VPN server in Malaysia must communicate smoothly with the Giant Widgets server in New York. They have to agree on how every packet is to be encapsulated. Encryption must be agreed upon and somehow they must both have the right keys for the encryption/decryption process. They must be assured that some sneak hasn't gotten into the loop, that they are talking with authenticated users, and that the data has not been tampered with, and for all this to work they must be using the same methods.

It is in assembling all these pieces, and agreeing on them, that things can get a little sticky. Right now the VPN industry is like a teenage kid in the middle of a growth spurt. His arms and legs have suddenly grown out of his clothes. He bumps into chairs, drops the carton of milk off the table, his hormones are raging, his skin is bumpy, he suffers from selective deafness and blindness ("Take the garbage out? What garbage? Out where?"), and his mind is controlled from somewhere in outer space.

As you know, if you've done any reading about VPNs, there are a number of VPN protocols out there. Microsoft has offered up PPTP, Cisco Systems has developed L2F, and Digital Equipment has AltaVista; as of this writing, L2TP and IPSec are working their way through the Internet Engineering Task Force (IETF) development and approval process.

There are at least a half a dozen encryption algorithms to choose from, and different key lengths, some of which are legal in the United States but illegal to export anywhere outside the U.S. A number of different vendors are offering both hardware and software solutions to the VPN market. It seems that more than half the firewall vendors and router manufacturers are getting into the act, along with technology companies like Lucent. Every Internet security company that offers encryption has joined in the fray. Everyone has different ideas of how best it should be done, and few agree.

By means of protocols such as L2TP and IPSec, the IETF is trying to bring some order out of all this chaos. They have a number of different working groups studying these and other protocols in an effort to develop some standards. These working groups are made up mostly of volunteer software engineers, employed in academia and industry, who are often working for

competing vendors. The process is one of negotiation, discussion, and clarification and works amazingly well. In the end, standards emerge from these working groups that we hope will offer a reasonably high degree of interoperability between those competing vendors' products that use the same standards.

The final step in the process rests with the vendors. Once the standards are published, it is up to the vendors to implement them. Just *how* they implement them is as important as all the rest of the process. They may not implement them exactly the same way, but for the industry to thrive, they must work together. A VPNet VPN should be able to talk to one from Shiva, Ascend, or Cisco. Fortunately, there are industry programs that are testing products for interoperability.

1.5 Where We Go from Here

The purpose of this book is to sort out some of this confusion. So far we've seen what a VPN is, some of the benefits it offers, and how it works. We've also warned that there are some downsides to running a VPN. We'll probe each of these topics more deeply as we proceed. Chapter 2 examines how VPNs are being put to use in the real world so that you can get some idea of how you might put a VPN to use for yourself and the benefits you should see from it.

