Because compulsory tunnels are created without the user's consent, they may be transparent to the end user. The client-side endpoint of a compulsory tunnel typically resides on a remote access server. All traffic originating from the end user's computer is forwarded over the PPTP tunnel by the RAS. Access to other services outside the intranet would be controlled by the network administrators. PPTP enables multiple connections to be carried over a single tunnel.

Because a compulsory tunnel has predetermined endpoints and the user cannot access other parts of the Internet, these tunnels offer better access control than voluntary tunnels. If it's corporate policy that employees cannot access the public Internet, for example, a compulsory tunnel would keep them out of the public Internet while still allowing them to use the Internet to access your VPN.

Another advantage to a compulsory tunnel is that multiple connections can be carried over a single tunnel. This feature reduces the network bandwidth required for transmitting multiple sessions, because the control overhead for a single compulsory tunnel carrying multiple sessions is less that that for multiple voluntary tunnels, each carrying traffic for a single session. One disadvantage of compulsory tunnels is that the initial link of the connection (i.e., the PPP link between the end user's computer and the RAS) is outside the tunnel and, therefore, is more vulnerable to attack.

Static compulsory tunnels typically require either dedicated equipment or manual configuration. These dedicated, or automatic, tunnels might require the user to call a special telephone number to make the connection. On the other hand, in realm-based, or manual, tunneling schemes, the RAS examines a portion of the user's name, called a *realm*, to decide where to tunnel the traffic associated with that user.

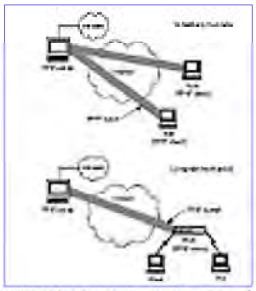

FIGURE 6.5 Voluntary and compulsory tunnels.

However, setup and maintenance of static tunnels increases the demands on network management. A more flexible approach would be to dynamically choose the tunnel destination on a per-user basis when the user connects to the RAS. These dynamic tunnels can be set up in PPTP by linking the system to a RADIUS server to obtain session configuration data on the fly.

Static tunneling requires the dedication of a *network access server* (NAS) to the purpose. In the case of

an ISP, this restriction would be undesirable because it requires the ISP to dedicate an NAS to tunneling service for a given corporate customer, rather than enabling them to use existing network access servers deployed in the field. As a result, static tunneling is likely to be costly for deployment of a global service.

Realm-based tunneling assumes that all users within a given realm want to be treated the same way, limiting a corporation's flexibility in managing the account rights of their users. For example, MegaGlobal Corp. may desire to provide Jim with an account that allows access to both the Internet and the intranet, with Jim's intranet access provided by a tunnel server located in the engineering department. However, MegaGlobal Corp. may want to provide Sam with an account that provides only access to the intranet, with Sam's intranet access provided by a tunnel network server located in the sales department. Situations like these cannot be accommodated with realm-based tunneling.

Using RADIUS to provision compulsory tunnels has several advantages. For instance, tunnels can be defined and audited on the basis of authenticated users, authentication and accounting can be based on telephone numbers; and other authentication methods, such as tokens or smart cards, can be accommodated. When deployed in concert with roaming, user-based tunneling offers corporations the capability to provide their users with access to the corporate intranet on a global basis.

RADIUS

The RADIUS client/server model uses a network access server to manage user connections. Although the NAS functions as a server for providing network access, it also functions as a client for RADIUS. The NAS is responsible for accepting user connection requests, getting user ID and password information, and passing the information securely to the RADIUS server. The RADIUS server returns authentication status, i.e., approved or denied, as well as any configuration data required for the NAS to provide services to the end user.

Roaming

Various ISPs have started to form strategic alliances—for example, the Stentor Alliance between MCI, British Telecom, and Bell Canada—that allow the partners to tunnel traffic across one another's networks. These agreements make it easier for your mobile workers to tunnel traffic to your corporate sites regardless of their location. If their work takes them to areas not serviced by your ISP, then they can call one of the partner ISPs in the area to use the VPN.

RADIUS creates a single, centrally located database of users and available services, a feature particularly important for networks that include large modem banks and more than one remote communications server. With RADIUS, the user information is kept in one location, the RADIUS server, which manages the authentication of the user and access to services from one location. Because any device that supports RADIUS can be a RADIUS client (see Figure 6.6), a remote user will gain access to the same services from any communications server communicating with the RADIUS server.

RADIUS supports the use of proxy servers, which store user information for authentication purposes and can be used for accounting and authorization, but they do not allow the user data (passwords and so on) to be changed. A proxy server depends on periodic updates of the user database from a master RADIUS server (see Figure 6.6). When corporations are looking to outsource their VPN to an ISP, they probably will arrange to have an ISP authenticate users of its PPTP server based on corporate-defined user data. In such cases, the corporation would maintain a RADIUS server and set user information on it, and the ISP

would have a proxy RADIUS server that receives updates from the corporate server.

For RADIUS to control the setup of a tunnel, it has to store certain attributes about the tunnel. These attributes include the tunnel protocol to be used (i.e., PPTP or L2TP), the address of the desired tunnel server, and the tunnel transport medium to be used. In order to take further advantage of RADIUS' capabilities—namely, its capability to track network usage—a few more items are needed—the address of the tunnel client (the NAS) and a unique identifier for the tunneled connection.

Previous | Table of Contents | Next

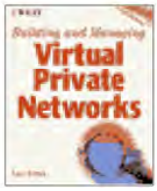

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

When combining dynamic tunneling with RADIUS, at least three possible options are available for user authentication and authorization:

1. Authenticate and receive authorization once, at the RAS end of the tunnel.
2. Authenticate and receive authorization info once, at the RAS end of the tunnel and somehow forward the RADIUS reply to the remote end of the tunnel.
3. Authenticate on both ends of the tunnel.

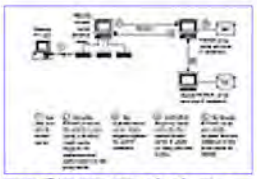

FIGURE 6.6 Interactions among a RADIUS server, proxy server, and clients.

The first model is a poor trust model because it requires the ISP alone to control access to the network, and the second is an adequate trust model but doesn't scale well, due to the way RADIUS authenticates replies. The third option is robust and works well if a RADIUS proxy server is used, which also supports the use of a single user name and password at both ends.

Let's look at the chain of events for creating a tunnel when using RADIUS this way (see Figure 6.7). First, the remote user dials into the remote access server and enters his password as part of the PPP authentication sequence (step 1 in the figure). The remote access server, acting as a RADIUS client, then uses RADIUS to check the password and receives tunnel information from the local RADIUS proxy server; this information would include attributes specifying which PPTP server is to be the endpoint of the tunnel that will be used for this particular user (steps 2 to 5). The remote access server will open the tunneled connection, creating a tunnel if necessary. *Recall that traffic from more than one user can be transmitted in the same compulsory tunnel at the same time.* The PPTP server would reauthenticate the user (step 6), checking the password against the same RADIUS server that was used in the initial exchange (steps 7 and 8). Upon authentication, the PPTP server will accept tunneled packets from the remote user and forward the packets to the appropriate destination on the corporate network.

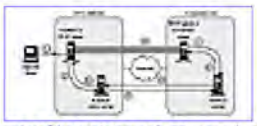

FIGURE 6.7 RADIUS authentication for dynamic tunnels.

Authentication and Encryption

Remote PPTP clients are authenticated by the same PPP authentication methods used for any RAS client dialing directly to a RAS server. Microsoft's implementation of RRAS supports CHAP, MS-CHAP, and PAP authentication schemes. MS-CHAP uses the MD4 hash for creating the challenge token from the user's password.

PAP and CHAP do have definite disadvantages when secure authentication is desired. Both PAP and CHAP rely on a secret password that must be stored on the remote user's computer and the local computer. If either computer comes under the control of a network attacker, then the secret password is compromised. Also, with CHAP or PAP authentication, you cannot assign different network access privileges to different remote users who use the same remote host. Because one set of privileges is assigned to a specific computer, everybody who uses that computer will have the same set of privileges.

In Microsoft's implementation of PPTP, data is encrypted via *Microsoft Point-to-Point Encryption* (MPPE), which is based on the RSA RC4 standard (see Figure 6.8). The *Compression Control Protocol* (CCP) used by PPP is used to negotiate encryption. MS-CHAP is used to validate the end user in a Windows NT domain, and an encryption key for the session is derived from the hashed user password stored on both the client and server. (A MD4 hash is used.) A 40-bit session key normally is used for encryption, but U.S. users can install a software upgrade to use a 128-bit key. Because MPPE encrypts PPP packets on the client workstation before they enter a PPTP tunnel, the packets are protected throughout the link from the workstation to the PPTP server at the corporate site. Changes in session keys can be negotiated to occur for every packet or after a preset number of packets.

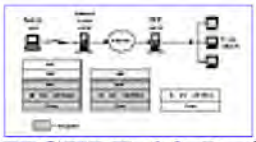

FIGURE 6.8 Packet encryption in PPTP.

LAN-to-LAN Tunneling

The original focus of PPTP was the creation of dial-in VPNs (i.e., to provide secure dial-in access to corporate LANs via the Internet). LAN-to-LAN tunnels were not supported at first. It wasn't until Microsoft introduced their Routing and Remote Access Server for NT Server 4.0 that NT Servers were able to support LAN-to-LAN tunnels. Since then, other vendors also have released compatible PPTP servers that also support LAN-to-LAN tunneling.

As implemented in Microsoft's RRAS, LAN-to-LAN tunneling occurs between two PPTP servers, much like IPSec's use of security gateways to connect two LANs. However, because the PPTP architecture does not make use of a key management system, authentication and encryption are controlled via CHAP, or via MS-CHAP. In effect, one site's RRAS, running PPTP, is defined as a user, with an appropriate password, at the other site's RRAS and vice versa (see Figure 6.9). To create a tunnel between the two sites, the PPTP server at one site is authenticated by the other PPTP server using the stored passwords, much as we described the process earlier for a dial-in user. One site's PPTP server thus looks like a PPTP client to the other server, and vice versa, so a voluntary tunnel is created between the two sites.

Because this tunnel can encapsulate any supported network-layer protocol (i.e., IP, NETBEUI, IPX), users at one site will have access to resources at the other site based on their access rights, defined for that protocol. This means that some form of collaboration between site managers is needed to ensure that

users at a site have the proper access rights to resources at other sites. In Windows NT, for example, each site can have its own security domain and the sites would establish a trust relationship between the domains in order to allow users to access a site's resources.

Using PPTP

Because a major focus of PPTP is to provide secure dial-in access to private corporate resources, the components of a PPTP VPN are organized a bit differently from those of an IPSec VPN (see Chapter 5, "Using IPSec to Build a VPN"). The most important components are those that define the endpoints of a PPTP tunnel. Because one of these endpoints can be your ISP's equipment, this configuration can cut down on the software needed for your mobile clients but requires collaboration between you and your ISP for authentication of users.

FIGURE 6.9 LAN-to-LAN PPTP tunnels.

In general, a PPTP VPN requires three items: a network access server, a PPTP server, and a PPTP client. Although the PPTP server should be installed on your premises and maintained by your staff, the network access server should be the responsibility of your ISP. In fact, if you choose to install PPTP client software on your remote hosts, the ISP doesn't even need to provide any PPTP-specific support.

Figure 6.10 illustrates few differences between the structure of an IPSec VPN and a PPTP VPN. One significant difference is that PPTP enables you to outsource some of the PPTP functions to the ISP. At a corporate site, a PPTP server acts like a security gateway, tying authentication to RADIUS or Windows NT domains. A PPTP client on a user's laptop or desktop computer performs many of the same functions as IPSec client software, although there are no key exchanges.

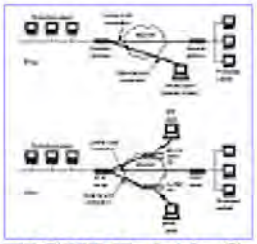

FIGURE 6.10 Comparing IPSec and PPTP architectures.

Previous | Table of Contents | Next

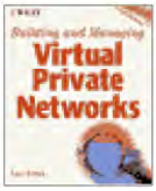

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

PPTP Servers

A PPTP server has two primary roles: it acts as the endpoint for PPTP tunnels, and it forwards packets to and from the tunnel that it terminates onto the private LAN. The PPTP server forwards packets to a destination computer by processing the PPTP packet to obtain the private network computer name or address information in the encapsulated PPP packet.

PPTP servers also can filter packets, using *PPTP filtering.* With PPTP filtering, you can set the server to restrict who can connect to either the local network or to the Internet. In systems like Windows NT 4.0 and RRAS, the combination of PPTP filtering with IP address filtering enables you to create a functional firewall for your network.

Setting up a PPTP server at your corporate site brings with it a few restrictions, especially if the PPTP server is to be placed on the private (i.e., corporate) side of the firewall. PPTP has been designed so that only one TCP/IP port number can be used for passing data through a firewall—port number 1723. This lack of configurability of the port number can make your firewall more susceptible to attacks. Also, if you have firewalls configured to filter traffic by protocol, you will need to set them to allow GRE to pass through.

A related device is the tunnel switch. Tunnel switches are relatively new devices, initially introduced by 3Com in early 1998. A tunnel switch is a combined tunnel terminator and tunnel initiator. The purpose of a tunnel switch is to extend tunnels from one network to another—extending a tunnel incoming from your ISP's network to your corporate network, for example (see Figure 6.11).

Tunnel switches can be used at a firewall to improve the management of remote access to private network resources. Because the tunnel switch terminates the incoming tunnel, it can examine the incoming packets for protocols carried by the PPP frames or for the remote user's name. The switch can use that information to create tunnels into the corporate network based on the information carried in the incoming packets.

PPTP Client Software

As pointed out frequently in this chapter, if the ISP equipment supports PPTP, no additional software or hardware is required on the client end; only a standard PPP connection is necessary. On the other hand, if the ISP does not support PPTP, a Windows NT client (or similar software) can still utilize PPTP and create the secure connection, first by dialing the ISP and establishing a PPP connection, then by dialing once again through a virtual PPTP port set up on the client side.

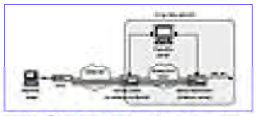

FIGURE 6.11 Example of the use of tunnel switches.

PPTP clients already exist from Microsoft for computers running Windows NT, Windows95, and Windows 98. Network Telesystems also offers PPTP clients for other popular computers, including the Macintosh and computers running Windows 3.1. When selecting a PPTP client, compare its functionality to that of your PPTP server. Not all client software will necessarily support MS-CHAP for instance, which means they won't be able to take advantage of Microsoft's encryption in RRAS.

Network Access Servers

Unlike an IPSec VPN, there are many cases in which a PPTP VPN's design depends on the protocol support offered by the ISP. This support is particularly important if your mobile workers can use a PPP client but do not have PPTP clients installed.

Because ISPs can offer PPTP services without adding PPTP support to their access servers, this approach would require that all clients use a PPTP client on their computers. This approach has its advantages because it enables clients to use more than one ISP if the geographic coverage of a primary ISP isn't adequate. Also recall that remote hosts with a PPTP client can set up voluntary tunnels in the PPTP scheme of things; if you want to control employee access to Internet resources, then you'll have to resort to compulsory tunnels, which require the support of your ISP.

It's unlikely that you'll have any control over the PPTP hardware that your ISP uses, but you should be aware of its capabilities so that you can take the hardware's limitations into account in the design of your VPN.

Network access servers, which are also known as *remote access servers* or *access concentrators*, provide software-based line access management and billing capabilities and run on platforms that offer robustness and fault tolerance at ISP POPs. ISP network access servers generally are designed and built to accommodate a large number of dial-in clients. An ISP that provides PPTP service would have to install a PPTP-enabled network access server that supports PPP clients on a number of platforms, including Windows, Macintosh, and Unix.

In such cases, the ISP server acts as a PPTP client and connects to the PPTP server at the corporate network. The ISP access server thus becomes one of the endpoints for a compulsory PPTP tunnel, with the network server at the corporate site being the other endpoint.

The network access server would choose a tunnel that has not only the appropriate endpoint but also the appropriate level of performance and service. Network access servers can make tunneling choices based on calling number, called number, static port mappings, text-based "terminal server" login, user names (from PAP or CHAP authentication), user-name parsing through DNS, lookups to RADIUS or TACACS+, ISDN call type, or command-line tunnel requests.

Early versions of PPTP devices and software were designed to work with Microsoft's version of PPTP and for remote access only. For instance, it wasn't until the second quarter of 1998 that products other than Windows NT 4.0 could be used as PPTP servers. LAN-to-LAN PPTP tunneling wasn't supported

until Microsoft released their Routing and Remote Access Server (RRAS) in late 1997.

A few vendors already support PPTP (see Table 6.1 for a partial list), with most of the initial equipment designed for ISPs. Since Microsoft's release of RRAS, other vendors also have started providing PPTP servers with similar features. If you're planning to install a PPTP VPN, you'll need to check the interoperability of your equipment with those of the ISP(s) you plan on using, because some features, like MS-CHAP, aren't supported on all devices and client software.

Sample Deployment

To illustrate the use of PPTP in a VPN, we'll create two different scenarios, one strictly for dial-in access (see Figure 6.12) and the second for a LAN-to-LAN VPN (see Figure 6.13). For simplicity's sake, we'll just have two sites—the corporate headquarters and a branch office—for the second example. In both cases, we'll concentrate on the exchange of data between endpoints and not worry about how the information is protected inside the corporate network (using firewalls, for example).

Previous | Table of Contents | Next

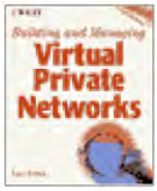

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

TABLE 6.1 Partial List of PPTP Products

Vendor	*Product*
3Com	AccessBuilder 5000, NETBuilder II
Ascend Communications	Max TNT
Bay Networks	Contivity Extranet Switches
Checkpoint Software Technologies	Firewall-1
ECI Telematics	Dial Access Concentrator
Extended Systems	ExtendNet VPN
Freegate Corp.	VPN Remote
Microcom	Access Integrator 1700
Microsoft Corp.	Windows NT Server, RRAS
Network Telesystems	Tunnel Builder
Shiva Corp.	LanRover Access Switch
US Robotics (now 3Com)	Total Control Enterprise Network Hub

Just as with the IPSec example given in Chapter 5, physical security should include ensuring that all hosts reside within the site's physical parameters and all links to outside systems go through the PPTP server and an associated firewall. The connection between the site's internal networks and the external network(s) should be in a locked machine room with restricted access, and only authorized individuals (network managers, for instance) should have access to the encrypting routers.

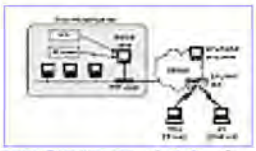

FIGURE 6.12 Sample PPTP dial-in VPN.

In the scenario diagrammed in Figure 6.12, MegaGlobal Corp. has decided to outsource much of the VPN work to its ISP. This means that the ISP providing MegaGlobal Corp.'s Internet connectivity has a RADIUS proxy server and PPTP-enabled network access servers. MegaGlobal Corp. still has to maintain a master RADIUS server and a PPTP server. Because the ISP is presumed to have PPTP-enabled access servers, you don't have to install special PPTP client software on the computers of your mobile workers.

Employing a RADIUS server to control authentication and access rights offers you the ability to centralize control of access, which can be particularly valuable if you're working in a multiprotocol environment. That's because many RADIUS servers have the capability to exchange information with other NOS-based directories, such as Windows NT and *Novell Directory Services* (NDS).

Now let's take a look at a VPN designed just for LAN-to-LAN connectivity, as in Figure 6.13.

In this example, a Windows NT server is installed at each site to serve as a router and PPTP server. In order for the two sites to communicate with each other over a PPTP tunnel, each PPTP server also will have to be configured to be a PPTP client of the other server. If the two sites connect via on-demand dialing, rather than through a permanent network link, the IP address of the ISP's network access server also has to be included in the configuration.

When any branch office traffic destined for the corporate site arrives at the branch office's PPTP server, the server will act as a PPTP client and will create a PPTP tunnel, if one doesn't already exist, to the corporate PPTP server in order to transfer the traffic. If traffic from the corporate site is destined for the branch office, the roles are reversed; the corporate PPTP server takes on the role of a PPTP client and creates a tunnel to the branch office's PPTP server.

FIGURE 6.13 An example PPTP LAN-to-LAN VPN.

As mentioned earlier in this chapter, one of the primary concerns for managing LAN-to-LAN PPTP links is ensuring that users at one site have the appropriate access rights at the other sites. This access can be achieved in Windows NT either by creating a master domain covering all sites or by letting each site be its own domain. In the first case, or any similar situation in which a hierarchy of domains might be used, the tunnels will have to carry added traffic as rights are passed between sites to check a user's traffic. This added traffic might be undesirable; also, using a centralized domain increases the risk of losing authentication between two branch offices if the main domain is unreachable. If independent domains, one for each site, are deployed, then the domain managers will have to establish the appropriate trust relations between sites and exchange user rights accordingly.

Applicability of PPTP

As an interim solution for VPNs, PPTP has a lot going for it, especially if you're running a Windows-only shop. PPTP is an interim solution because most vendors are planning to replace PPTP with L2TP when the protocols are standardized. As you plan to create a PPTP VPN, it would pay to keep an eye on your vendors' plans for L2TP.

PPTP is also better suited for handling dial-up access by a limited number of remote users rather than LAN-to-LAN VPNs. One problem is the need to coordinate user authentication rights across LANs, either via NT domains or RADIUS. Also, the scalability of PPTP servers has often been called into question for large numbers of remote users and for large amounts of traffic, such as might be required for LAN-to-LAN links.

That said, PPTP can still be a good way for you to become familiar with VPNs. A VPN can still be a

good cost-reduction measure, even if it's only focused on remote access costs. (Go back and review Chapter 2, "Virtual Private Networks," if you want to see some numbers.) Plus, if you can find an ISP that supports PPTP on its equipment, you can outsource some of your VPN management to the ISP.

If you're not running a Windows-only shop, then you'll have to bite the bullet and perhaps add management of an NT server to your list of tasks in order to use PPTP. The dependence of PPTP on Windows NT isn't likely to go away, especially with L2TP around the corner. Analyze this option carefully, as the cost savings accompanying an NT server may be more than offset by the support costs, if you're not already familiar with NT.

PPTP's security features aren't nearly as robust as those found in IPSec; see www.counterpane.com for some of the details. On the positive side, that means that security management is less complex for PPTP. But, the placement of the PPTP server with respect to any firewalls, as mentioned earlier, raises security concerns and opens possible holes for attackers.

PPTP's shortcomings make it a reasonable solution for remote access and multiprotocol traffic rather than LAN-to-LAN VPNs. Its popularity on the Windows NT platform and available clients for other popular PC platforms have given it a good headstart for dial-in VPNs. If you need to build a VPN that doesn't suffer from the restrictions of PPTP but aren't ready (or willing) to deploy IPSec, a better solution for VPNs is L2TP (Layer2 Tunneling Protocol), which will be covered in the following chapter.

Summary

We've just covered the details of how PPTP, a popular protocol for dial-up VPNs, works. PPTP systems are rather tightly tied to Windows NT, mainly because so many of the PPTP servers are run on NT servers. But, PPTP can be configured to support either PPP or PPTP clients, making it easier to support a variety of operating systems and clients among your mobile workers. Because it's based on PPP, PPTP is well-suited to handling multiprotocol network traffic, particularly IP, IPX, and NETBEUI protocols.

PPTP's design also makes it easier to outsource some of the support tasks to an ISP. By using RADIUS proxy servers, an ISP can authenticate dial-in users for corporate customers and create secure PPTP tunnels from the ISP's network access servers to your corporate PPTP servers. These PPTP servers then remove the PPTP encapsulation and forward the network packets to their appropriate destination on your private network.

Previous | Table of Contents | Next

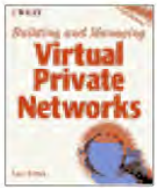

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 7
Using L2TP to Build a VPN

Now that we're turning our attention to the *Layer2 Tunneling Protocol* (L2TP) in this chapter, we're almost finished with the three-letter and four-letter acronyms that make up the alphabet soup of VPNs.

L2TP should be considered the successor to PPTP; it combines many of the features originally defined in PPTP with those created for another protocol, *Layer2 Forwarding* (L2F) originally designed and implemented by Cisco. L2F has seen limited deployment; because L2TP combines the best features of the two protocols, it's been forecast that L2TP will supersede both PPTP and L2F as it becomes a standard sometime this year. Many vendors offering support for PPTP in their products either already include L2TP support as well or have plans to supersede PPTP with L2TP.

This chapter starts out with an overview of the architecture of L2TP and moves on to the details of how the protocol works, including its use of IPSec for encryption. Then we move on to an overview of the types of products you can use to build a VPN using L2TP.

What Is L2TP?

The Layer2 Tunneling Protocol was created as the successor to two tunneling protocols, PPTP and L2F. Rather than develop two competing protocols to do essentially the same thing—PPTP by Microsoft et al. versus L2F by Cisco—the companies agreed to work together on a single protocol, L2TP, and submit it to the IETF for standardization. Because we've already devoted a chapter to PPTP, we'll include a few words about L2F as background for our discussion of L2TP.

Like PPTP, L2F was designed as a tunneling protocol, using its own definition of an encapsulation header for transmitting packets at Layer2. One major difference between PPTP and L2F is that the L2F tunneling isn't dependent on IP and GRE, enabling it to work with other physical media. Because GRE isn't used as the encapsulating protocol, L2F specifications define how L2F packets are handled by different media, with an initial focus on IP's UDP.

Paralleling PPTP's design, L2F utilized PPP for authentication of the dial-up user, but it also included support for TACACS+ and RADIUS for authentication from the beginning. L2F differs from PPTP by defining connections within a tunnel, allowing a tunnel to support more than one connection. There are also two levels of authentication of the user: first, by the ISP prior to setting up the tunnel; second, when the connection is set up at the corporate gateway.

These L2F features have been carried over to L2TP. Like PPTP, the Layer2 Forwarding Protocol utilizes the functionality of PPP to provide dial-up access that can be tunneled through the Internet to a destination site. However, L2TP defines its own tunneling protocol, based on the work of L2F. Work has continued on defining L2TP transport over a variety of packetized media such as X.25, frame relay, and ATM. Although many of the initial implementations of L2TP focus on using UDP on IP networks, it's possible to set up a L2TP system without using IP as a tunnel protocol at all. A network using ATM or frame relay also can be deployed for L2TP tunnels.

Because L2TP is a Layer2 protocol, it offers users the same flexibility as PPTP for handling protocols other than IP, such as IPX and NETBEUI, for example.

Because it uses PPP for dial-up links, L2TP includes the authentication mechanisms within PPP, namely PAP and CHAP; like PPTP, L2TP supports PPP's use of the Extensible Authentication Protocol for other authentication systems, such as RADIUS. Many of the examples of RADIUS use given in Chapter 6, "Using PPTP to Build a VPN," also apply to L2TP.

We'll see later in this chapter that the designers of L2TP were concerned with the end-to-end authentication and data integrity of data passed from the end user to an L2TP server. Because of this concern, they devised ways to invoke IPSec-based authentication and encryption across the PPP link (see Figure 7.1). Using IPSec at the end user's workstation provides stronger security than simply relying on PPP-based authentication and encryption, as PPTP does.

Although Microsoft has made PPTP a popular choice for setting up dial-in VPNs by including support for the protocol within its Windows operating systems, the company also has plans to add support for L2TP within Windows NT 5.0 and Windows 98, which should make it easier for L2TP to become a widely used successor to PPTP. However, unlike PPTP, the feature set of L2TP is defined within the IETF's standards committees and is not necessarily being driven by the features found in Windows NT, as PPTP originally was.

The Building Blocks of L2TP

The components of an L2TP system are essentially the same as those for PPTP: the Point-to-Point Protocol, tunnels, and authentication systems like RADIUS. However, to increase the security of L2TP traffic, IPSec can be used to protect data, which brings key management into play following many of the procedures covered in Chapter 5, "Using IPSec to Build a VPN."

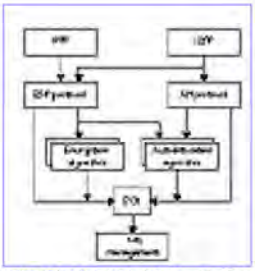

FIGURE 7.1 L2TP's architecture.

PPP and L2TP

PPP is the most common protocol for dial-up access to the Internet and other TCP/IP networks. Working

at Layer2, the Data Link layer, of the OSI protocol stack, PPP includes methods for encapsulating various types of datagrams for transfers over serial links. PPP can encapsulate AppleTalk, IP, IPX, and NETBEUI packets between PPP frames and can send those encapsulated packets by creating a point-to-point link between the sending and receiving computers.

L2TP depends on the PPP protocol to create the dial-up connection between the client and a network access server. L2TP expects PPP to establish the physical connection, perform the first authentication phase of the end user, create PPP datagrams, and close the connection when the session is finished.

When PPP has established the connection, L2TP takes over. First, L2TP determines whether the network server at the corporate site recognizes the end user and is willing to serve as an endpoint for a tunnel for that user. If the tunnel can be created, L2TP takes on the role of encapsulating the PPP packets for transmission over the medium that the ISP has assigned to the tunnel (see Figure 7.2).

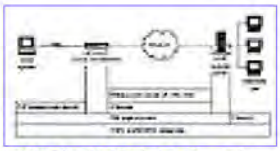

FIGURE 7.2 Example protocols used in an L2TP connection.

As L2TP creates tunnels between the ISP's access concentrator and the client's network server, it can assign more than one session to a tunnel. L2TP creates a *Call ID* for each session and inserts the Call ID into the L2TP header of each packet to indicate to which session it belongs.

It's also possible to create multiple simultaneous tunnels between an ISP's access concentrator and the client's network server. By choosing to assign a single user session to a tunnel rather than multiplex a series of sessions into a tunnel (as in the preceding paragraph), different tunnel media can be assigned to different users according to their *quality-of-service* (QoS) requirements. L2TP includes a tunnel identifier so that the individual tunnels can be identified when arriving from a single source, either an access concentrator or a network server.

Previous | Table of Contents | Next

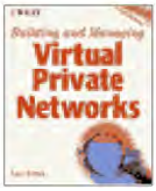

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Much like PPTP, the L2TP protocol defines two different types of messages—control messages and data messages—which it uses for setup and maintenance of the tunnels as well as the transmission of the data. However, unlike PPTP, L2TP transmits both the control messages and data messages as part of the same stream. If tunnels are being transmitted over an IP network, for instance, control and data are sent in the same UDP datagram.

The L2TP control messages are responsible for the establishment, management, and release of sessions carried through the tunnel, as well as the status of the tunnel itself.

In L2TP data messages, the payload packet is essentially the original PPP packet sent by the client, missing only framing elements that are specific to the media. Because L2TP operates as a Layer2 protocol, it must include a media header in the packet description to indicate how the tunnel is being transmitted (see Figure 7.3). Depending on your ISP's infrastructure, this might be Ethernet, frame relay, X.25, ATM, or PPP links.

L2TP also helps reduce network traffic and enables servers to handle congestion by implementing flow control between the network access server, an *L2TP Access Concentrator* (LAC), and the corporate network server, an *L2TP Network Server* (LNS) in L2TP terminology. Control messages are used to determine the transmission rate and buffering parameters that are used to regulate the flow of PPP packets for a particular session over the tunnel. To keep performance high, L2TP tries to keep overhead to a minimum—for example, by compressing packets headers. L2TP uses the same tunnel classes as PPTP (i.e., voluntary and compulsory tunnels) depending on whether the end user uses a PPP client or L2TP client to initiate the connection.

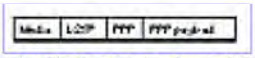

FIGURE 7.3 L2TP packet encapsulation.

Tunnels

Voluntary tunnels are created at the request of the user for a specific use (see Figure 7.4). Compulsory tunnels are created automatically without any action from the user, and more importantly, without allowing the user any choice in the matter.

Voluntary tunnels are just that, set up at the request of the end user. When using a voluntary tunnel, the end user simultaneously can open a secure tunnel through the Internet and access other Internet hosts via basic TCP/IP protocols without tunneling. The client-side endpoint of a voluntary tunnel resides on the user's computer. Voluntary tunnels are often used to provide privacy and integrity protection for intranet traffic being sent over the Internet.

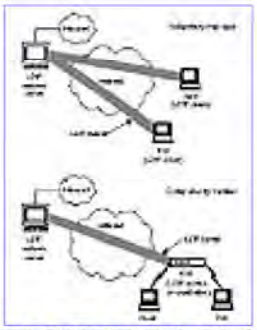

FIGURE 7.4 Voluntary and compulsory tunnels.

Because compulsory, or *mandatory*, tunnels are created without the user's consent, they may be transparent to the end user. The client-side endpoint of a compulsory tunnel resides on the ISP's LAC. All traffic originating from the end user's computer is forwarded over the L2TP tunnel by the LAC. Access to other services outside the intranet would be controlled by the network administrators. Keep in mind that L2TP allows multiple connections to be carried over a single tunnel, which improves L2TP's scalability and reduces the network's overhead for handling tunnels.

Because a compulsory tunnel has predetermined endpoints, and the user cannot access other parts of the Internet, these tunnels offer better access control than voluntary tunnels. If it's corporate policy that employees cannot access the public Internet, for example, a compulsory tunnel would keep them out of the public Internet while still allowing them to use the Internet to access your VPN.

Another advantage to a compulsory tunnel is that multiple connections can be carried over a single tunnel. This feature reduces the network bandwidth required for transmitting multiple sessions, because the control overhead for a single compulsory tunnel carrying multiple sessions is less than that for multiple voluntary tunnels, each carrying traffic for a single session. One disadvantage of compulsory tunnels is that the initial link of the connection (i.e., the PPP link between the end user's computer and the LAC) is outside the tunnel and, therefore, is more vulnerable to attack; this is one of the reasons why L2TP includes provisions for using IPSec to protect traffic, as we'll see in more detail shortly.

Although an ISP could choose to establish statically defined tunnels for its customers, this approach could tie up network resources unnecessarily if the static tunnels are unused or used infrequently. A more flexible approach, that of dynamically setting up the tunnel on a per-user basis when the user connects to the remote access server, or LAC, allows for more efficient use of the ISP's resources. One way to do this is for the ISP to store information about the end users, usually in a RADIUS server.

Using RADIUS to set up and control compulsory tunnels has several advantages. For instance, tunnels can be defined and audited on the basis of authenticated users; authentication and accounting can be based on telephone numbers; and other authentication methods, such as tokens or smart cards, can be accommodated.

In order for RADIUS to be able to control the setup of a tunnel, it has to store certain attributes about the tunnel. These attributes include the tunnel protocol to be used (i.e., PPTP or L2TP), the address of the desired tunnel server, and the tunnel transport medium to be used. In order to take further advantage of RADIUS' capabilities—namely, its capability to track network usage—a few more items are needed—namely the address of the tunnel client, the LAC, and a unique identifier for the tunneled connection. If the user information has to be linked with the customer's RADIUS (or other) database,

then the interaction between the ISP and the customer would be the same as described in Figure 6.7 in the preceding chapter.

Authentication and Encryption

The authentication of a user occurs in three phases in L2TP: the first at the ISP and the second and optional third phases at the corporate site's network server.

In the first phase, the ISP can use the caller's phone number, the number called, or a user name to determine that L2TP service is required and then would initiate a tunnel connection to the appropriate network server. When a tunnel is established, the ISP L2TP Access Concentrator would allocate a new Call ID to identify the connection within the tunnel and would initiate a session by forwarding the authentication information.

PPP's Extensible Authentication

In an effort to accommodate better, more robust methods of authentication within PPP, the IETF has defined the PPP *Extensible Authentication Protocol* (EAP) in RFC 2284. EAP is a general protocol for PPP authentication that supports multiple authentication mechanisms. EAP does not select a specific authentication mechanism at the Link Control Phase but rather postpones this until the Authentication Phase, allowing the authenticator to request more information before determining the specific authentication mechanism. This also permits the use of a back-end server that implements the various mechanisms while the PPP authenticator merely passes through the authentication exchange. By using EAP, you can integrate some of the systems we mentioned in Chapter 4, "Security: Threats and Solutions," like one-time passwords and secure tokens, into the use of PPP; EAP also makes integration of PPP with RADIUS easier.

Previous | Table of Contents | Next

The corporate network server undertakes the second phase of authentication by deciding whether or not to accept the *call*. The call start indication from the ISP might include CHAP, PAP, EAP, or other authentication information; the network server would use this information to decide to accept or reject the call.

After the call is accepted, the network server can initiate a third phase of authentication at the PPP layer. This step would be similar to that used by a company to authenticate remote access users who are dialing in the old way (i.e., using a modem).

Although these three phases of authentication may guarantee that the end user, ISP, and network server are actually who they say they are, nothing up to this point has been done to protect the data against snooping or alteration. The rest of this section points out where encryption can be applied to protect your data.

The tunnel endpoints may authenticate each other during tunnel establishment. This authentication has the same security attributes as CHAP and offers reasonable protection against replay attacks and snooping during the tunnel establishment process. But, it is still fairly simple for an attacker to snoop and inject packets to hijack a tunnel after an authenticated tunnel has been successfully completed.

On their own, L2TP and PPP authentication and encryption do not meet the security requirements for a VPN. Although L2TP authentication can handle mutual authentication of an LAC and an LNS during tunnel setup, it does not protect control and data traffic on a per-packet basis. This lack of protection leaves tunnels open to a variety of attacks, including snooping data packets, modifying both data and control packets, attempts to hijack the tunnel or the PPP connection, or disrupting PPP negotiations to weaken any confidentiality protection or to gain access to user passwords.

PPP authenticates the client to the LNS, but it also does not provide per-packet authentication, data integrity, or replay protection. PPP encryption does meet confidentiality requirements for PPP traffic but does not address authentication, data integrity, and key-management requirements, making it a weak security solution, which does not assist in securing the L2TP control channel.

If L2TP tunnel authentication is desired, it's necessary to distribute keys. Although manual key distribution might be feasible in a limited number of cases, a key-management protocol will be required for most situations.

For L2TP tunnels over IP, IP-level packet security using IPSec provides very strong protection of the tunnel. This security requires no modification to the L2TP protocol. For L2TP tunnels over frame relay or other switched networks, current practice indicates that these media are much less likely to experience attacks on in-transit data.

Note that several of the attacks outlined may be carried out on PPP packets sent over the link between the dial-up client and the NAS/LAC, prior to encapsulation of the packets within an L2TP tunnel. Even though this is not strictly the concern of the L2TP specification (being a part of how PPP handles the link), L2TP can be a better VPN solution if it protects data from end-to-end. This led to the proposal of using IPSec for encrypting packets, at least for IP-based tunnels.

Because ESP functions are defined on the IP payload, excluding the IP header, the presence of an IP header is not a requirement for the use of ESP. Therefore, L2TP implemented on non-IP networks can transport ESP packets. But, key exchange and negotiation of security associations (see Chapter 5) is another matter. IKE, or *ISAKMP/Oakley* if you want to use the old term for the protocols, messages use UDP transport, which would require that non-IP media used for L2TP tunnels would have to support the transport of UDP datagrams. (Is this a problem?)

Let's look at how IPSec would be implemented within L2TP for compulsory and voluntary tunnels. In the case of a compulsory tunnel, the end user sends PPP packets to the LAC and isn't really aware of the tunnel that's created between the LAC and the LNS at the corporate site. A security association may be set up between the LAC and the LNS based on the end user's requirements and identity, and this association would be known to the LAC and LNS but not to the end user.

Because the end user wouldn't be aware of what security services are in place between the LAC and the LNS for his traffic, the best approach is for the end user to rely on IPSec starting on his computer. But, not all of the endpoints might be IPSec-capable, which might force renegotiations for using only PPP encryption (see Figure 7.5). In both cases, the ISP's LAC could apply the IPSec Authentication Header to traffic traveling through the tunnel it creates, but the encryption choice is left up to the end user—either ESP for IPSec-capable destinations or PPP's encryption scheme for non-IPSec destinations.

In the case of a voluntary tunnel, the end user serves as one endpoint of the L2TP tunnel and, therefore, can negotiate a security association with the LNS at his corporate site. But, negotiation of SAs and keys again depends on whether or not both endpoints are IPSec-capable (see Figure 7.6). Because the end user's computer serves as the endpoint for voluntary tunnels, the IPSec Authentication Header is applied at his workstation, not at the ISP's device, which in this case is a network access server, not an L2TP access concentrator. If the destination is not IPSec-capable, then ESP encryption protects only the packets until they reach the LNS at the corporate site.

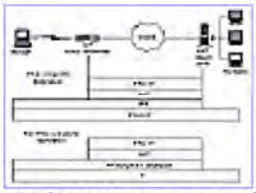

FIGURE 7.5 Packet encryption for compulsory tunnels.

Although IPSec has been specified as the security system of choice for use with L2TP, it may not be suitable for all situations. Although ESP can be used to encrypt non-IP payloads, AH and ESP are designed to be inserted into IP datagrams. The proposed L2TP solution is to always use UDP datagrams for transporting L2TP packets regardless of the medium involved, such as frame relay or ATM. This solution offers the advantage of using only one protocol for securing L2TP traffic, whether it's over IP or non-IP networks.

Other alternatives for non-IP networks are still being investigated. One suggested approach has been to include the *Security Parameters Index* (SPI) (see Chapter 5, "Using IPSec to Build a VPN") for a security association and a cryptographic initialization vector of 128 bits in the L2TP header. In addition, control messages would be defined for negotiating a security association.

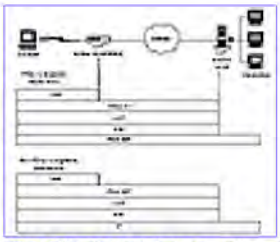

FIGURE 7.6 Packet encryption for voluntary tunnels.

LAN-to-LAN Tunneling

Although the primary focus of L2TP has been dial-up VPNs using PPP clients, it is possible to use L2TP for LAN-to-LAN links within a VPN.

The basic setup of LAN-to-LAN tunneling would occur between two L2TP servers with at least one having a dial-on demand link to their ISP, allowing them to initiate a PPP session whenever traffic is waiting for a destination at another VPN site. This type of arrangement would work best for branch offices that do not generate a great deal of traffic on their VPN links and do not need to stay connected to other VPN sites all the time.

Previous | Table of Contents | Next

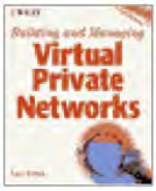

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

In effect, each site serves as both an L2TP access concentrator and a network server, initiating and terminating tunnels as needed (see Figure 7.7). If this were a demand dial-in situation like the one we described for PPTP in Chapter 6, "Using PPTP to Build a VPN," authentication could follow the same three steps as for a remote client (i.e., first by the ISP, then by the LNS at the receiving site, and finally any further authentication of the PPP traffic is set up by the receiving site.)

Running L2TP on Non-IP Networks

Work has been proceeding in the IETF to define PPP framing for interfaces other than asynchronous ones (i.e., modem and serial lines) and ISDN. This work would allow PPP to work over ATM, using the AAL5 and FUNI interfaces of ATM, as well as frame relay, which means that VPN sites could create L2TP tunnels among themselves using dedicated links to their ISPs rather than depending only on dial-up links.

For LANs continually connected via the Internet (i.e., using frame relay, T1, etc., not a dial-up link), most likely a shortcut in the authentication process would exist because an ISP's remote access server would not be involved as a LAC.

Key Management

When two parties want to exchange secure communications, they need to be sure that they're processing the data in the same way. The two parties have to be using the same cryptographic algorithm, the same key length, and the same keys if they're going to successfully exchange secure data; this is handled via a *Security Association* (SA). Although IPSec specifies default algorithms for authentication and encryption, it also allows for other algorithms to be used.

FIGURE 7.7 LAN-to-LAN L2TP tunnels.

A Security Association groups together all the things you need to know about how you communicate securely with someone else. An IPSec SA specifies the following:

- The mode of the authentication algorithm used in the AH and the keys to that authentication algorithm
- The ESP encryption algorithm mode and the keys to that encryption algorithm
- The presence and size of any cryptographic synchronization to be used in that encryption algorithm

- What protocol, algorithm, and key you use to authenticate your communications
- What protocol, encrypting algorithm, and key you use to make your communications private
- How often those keys are to be changed
- The authentication algorithm, mode, and transform for use in ESP plus the keys to be used by that algorithm
- The key lifetimes
- The lifetime of the SA itself
- The SA source address

Although security associations help two communicating parties define the cryptography they'll use to communicate, the procedures for exchanging and negotiating SAs as well as any keys involved in the communications are defined by IKE (or *ISAKMP/Oakley*, its older name). IKE is designed to provide four capabilities:

1. Provide the means for parties to agree on which protocols, algorithms, and keys to use.
2. Ensure from the beginning of the exchange that you're talking to the right person.
3. Manage those keys after they've been agreed upon.
4. Ensure that key exchanges are handled safely.

As you might expect, key exchange is closely related to the management of security associations. When you need to create an SA, you need to exchange keys. So IKE's structure wraps them together and delivers them as an integrated package.

Because IKE is IP-centric, it's easier to graft onto L2TP running over IP networks than over non-IP networks. Some question still exists concerning methods for negotiating security associations and managing keys when using L2TP over non-IP networks.

Using L2TP

Because a major focus of L2TP is to provide secure dial-in access to private corporate resources over the Internet, the components of an L2TP VPN are almost the same as for a PPTP VPN. (See Chapter 6 for more on PPTP.) The most important components are those that define the endpoints of an L2TP tunnel, the L2TP access concentrator, and the L2TP network server (see Figure 7.8). Because one of these endpoints can be your ISP's equipment, the software needed for your mobile clients can be reduced, which requires collaboration between you and your ISP for the first phase of authentication of users.

Although the LNS should be installed on your premises and maintained by your staff, the LAC should be the responsibility of your ISP. In fact, if you choose to install L2TP client software on your remote hosts, the ISP doesn't even need to provide any L2TP-specific support.

At a corporate site, an L2TP network server acts like a security gateway, tying authentication to RADIUS or Windows NT domains. An L2TP client on a user's laptop or desktop computer performs many of the same functions as IPSec client software, although there are no key exchanges.

FIGURE 7.8 Basic L2TP components.

Like PPTP, L2TP offers you the advantage of outsourcing some of the VPN functions to the ISP.

L2TP Network Servers

As with PPTP, an *L2TP network server* (LNS) has two primary roles: it acts as the endpoint for L2TP tunnels, and it forwards packets to and from the tunnel that it terminates onto the private LAN. The L2TP server forwards packets to a destination computer by processing the L2TP packet to obtain the private network computer name or address information in the encapsulated PPP packet. Any computing platform capable of terminating PPP sessions can operate as an LNS.

Unlike PPTP, L2TP is not designed for filtering packets. Instead, the system's architecture leaves that task to your firewall.

When it comes to integrating network servers with your firewalls, L2TP has some advantages over PPTP. First, L2TP does not demand that only one specific port number can be assigned for the firewall to pass L2TP traffic, as PPTP does. (A default port number, 1701, is defined for L2TP, though.) Network managers have the option of selecting a different firewall port number for passing L2TP traffic, making it more difficult for attackers to take over L2TP tunnels or try other attacks based on a known port number. Second, because the L2TP data and control traffic pass over a single UDP channel, firewall setup is simpler. Because some firewalls are not designed to support GRE, compatibility between L2TP and firewall products is less of an issue than for PPTP.

Since IPSec can play a large role in the security of your data, you should keep in mind some of the features and capabilities that we mentioned in Chapter 5 when reviewing L2TP network servers, namely:

- Support separate network connections for plaintext and ciphertext.
- Available key sizes must be consistent with the sensitivity of the information you'll transmit across the data link.
- If you decide that the default crypto algorithms will not meet your needs, the device should support the accepted alternative algorithms.
- Both AH and ESP ought to be supported.
- Manual input of SAs, including wild card SAs, should be supported.
- Mechanisms for protecting secret and private keys should be included.
- A system for changing crypto keys automatically and periodically makes key management easier and more secure.
- A security gateway should include some support for logging failures when processing a header; even better, some kind of alarm for persistent failures should be included.

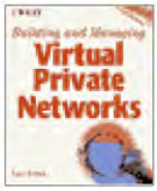

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

L2TP Client Software

If the ISP equipment supports L2TP, no additional software or hardware is required on the client end; only standard PPP software is necessary. Note that this setup would not support data encryption via IPSec, which means that you may have to keep on the lookout for IPSec-enabled PPP clients to make the most of L2TP. On the other hand, if the ISP does not support L2TP, then an IPSec-compliant L2TP client can be used to create tunnels to the corporate LNS.

If you're concerned with proper IPSec support in either PPP or L2TP clients, here are some features to check when evaluating client software:

- Offers compatibility with other IPSec implementations; (i.e., match the site's encrypting server—key exchange protocol, crypto algorithms, etc.)
- Offers a clear indication of when IPSec is working
- Supports downloading SAs (via paper or disks, for instance)
- Has to handle dynamically assigned IP addresses
- Includes mechanisms to protect the keying material from theft (encrypt keys with password, for instance)
- Offers a mechanism to change the crypto key automatically and periodically; includes dynamic assignment of new SPI numbers during rekeying; is compatible with standard IPSec keying protocols; uses a cryptographically strong random-key procedure to generate its keys
- Explicitly blocks non-IPSec traffic

Network Access Concentrators

Unlike an IPSec VPN, in many cases an L2TP VPN's design depends on the protocol support offered by the ISP. This support is particularly important if your mobile workers can use a PPP client but do not have L2TP clients installed. It's also important when you consider what encryption methods to use for protecting your data.

Because ISPs can offer L2TP services without adding L2TP support to their access servers, this approach would require that all clients use an L2TP client on their computers. This approach has its advantages because it enables clients to use more than one ISP if the geographic coverage of a primary ISP isn't adequate. Also recall that remote hosts with an L2TP client can set up voluntary tunnels; if you want to control employee access to Internet resources, then you'll have to resort to compulsory tunnels, which require the support of your ISP.

An ISP that provides L2TP service would have to install an L2TP-enabled network access server that supports PPP clients on a number of platforms, including Windows, Macintosh, and Unix. The ISP access concentrator thus becomes one of the endpoints for a compulsory L2TP tunnel, with the network server at the corporate site being the other endpoint.

The L2TP access concentrator would choose a tunnel that has not only the appropriate endpoint (i.e., your network server) but also the appropriate level of performance and service. Network access servers can make tunneling choices based on calling number, called number, static port mappings, text-based terminal server login, user names (PAP or CHAP authentication), user-name parsing through DNS, lookups to RADIUS or TACACS+, ISDN call type, or command-line tunnel requests.

Your selection of an ISP partner for an L2TP VPN also may hinge on the degree to which you want to protect your data. If you want end-to-end encryption, for instance, you would install IPSec-compliant clients on your mobile workers' computers and expect the ISP to handle encrypted packets from clients all the way to your network server.

If lesser security can be tolerated and you only want to protect your data as it travels through the tunnel over the Internet, then you should deal with an ISP who's installed an L2TP access concentrator that supports IPSec and will encrypt your traffic between the LAC and your LNS.

A few vendors already support L2TP (see Table 7.1 for a partial list). If you're planning to install an L2TP VPN, you should check the interoperability of your equipment with those of the ISP(s) you plan on using.

Sample Deployment

To illustrate the use of L2TP in a VPN, we'll focus on a scenario designed strictly for dial-in access (see Figure 7.9). As with the scenarios in the previous chapters, we'll concentrate on the exchange of data between endpoints and not worry about how the information is protected inside the corporate network (using firewalls, for example).

Just like the IPSec example given in Chapter 5 and the PPTP examples in Chapter 6, physical security should include ensuring that all hosts reside within the site's physical parameters and that all links to outside systems go through the L2TP network server and an associated firewall.

TABLE 7.1 Partial List of L2TP Products

Vendor	*Product*
3Com	AccessBuilder, HiPer Access Router
Ascend Communications	SecureConnect, Pipeline Routers
Bay Networks	Contivity Extranet Switch 1000, 2000, 4000
Checkpoint Software Technologies	Firewall-1
Cisco	IOS
Extended Systems	ExtendNet VPN
Freegate Corp.	VPN Remote
Microcom	Access Integrator 1700

Microsoft Corp.	Windows NT 5.0
Shiva Corp.	LanRover Access Switch

Previous | Table of Contents | Next

The connection between the site's internal networks and the external network(s) should be in a locked machine room with restricted access, and only authorized individuals (network managers, for instance) should have access to the encrypting routers.

In the scenario we've diagrammed in Figure 7.9, MegaGlobal Corp. has decided to outsource much of the VPN work to its ISP. This means that the ISP providing MegaGlobal Corp.'s Internet connectivity has a RADIUS proxy server and an L2TP access concentrator (i.e., an L2TP-enabled network access server). MegaGlobal Corp. still has to maintain a master RADIUS server and an L2TP network server. Because the ISP is presumed to have L2TP-enabled access servers, you don't have to install special L2TP client software on the computers of your mobile workers (unless you want to provide IPSec encryption of the data).

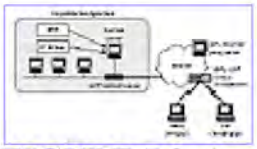

FIGURE 7.9 An example L2TP dial-in VPN.

Applicability of L2TP

L2TP will be the next-generation protocol for dial-in VPNs. It brings together the best features of PPTP and L2F, as well as supporting IPSec for improved data security. As a show of support for L2TP, most vendors of PPTP products are either offering L2TP-compatible products or will be introducing them before long.

Although a great deal of the initial development effort for L2TP has been focused on L2TP over IP, the capability to run L2TP over other networks, such as frame relay or ATM, should add to its long-term popularity. Plus, L2TP still has an advantage over IPSec, because it can transport protocols other than IP.

L2TP's support for non-IP networks also may prove to be a hindrance for some network planners, though. That's because IPSec's key-management scheme, IKE, is designed to work with IP, and translating IKE to other network protocols hasn't become a priority item. The PPPEXT Working Group of the IETF is still working on ways for securing L2TP traffic and managing keys over non-IP networks.

When we discussed PPTP, we mentioned that scalability concerns might arise if a large number of remote users need to be supported or if large amounts of traffic over LAN-to-LAN links might occur. Some of these same scalability concerns may apply to L2TP as well, but L2TP's congestion and flow-control measures should alleviate some of the problems.

Lastly, L2TP tunnels may be better-suited if you want to provide some type of quality-of-service controls to your workers. L2TP enables you to set up multiple tunnels between the same LAC and LNS; each tunnel can be assigned to a specific user, or class of users, and assigned to specific media according to QoS attributes that have been assigned to the user. Recall that IPSec's encryption of the packet header when using tunnel-mode ESP can make QoS assignments based on the user difficult, if not impossible.

Summary

We've now detailed the workings of L2TP, the third (and last) VPN protocol that we'll cover in this book. L2TP should be considered the next-generation VPN protocol, particularly for dial-in VPNs; most vendors already have plans to supplant PPTP-based products with L2TP products.

L2TP offers a number of the advantages of PPTP, particularly for handling multiple sessions over a single tunnel as well as assigning QoS parameters of different tunnels to the same site. In addition, L2TP's capability to run over media like X.25, frame relay, and ATM, while handling multiple network layer protocols, in addition to IP, affords users and ISPs a great deal of flexibility in designing VPNs. L2TP also provides stronger security for your data, because it uses IPSec's ESP for encrypting packets, even over a PPP link between the end-user and the ISP.

Previous | Table of Contents | Next

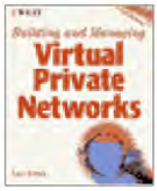

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 8
Designing Your VPN

The planning and design of a VPN should be done with care, because it not only affects the connectivity between different parts of your organization and the security of your data, but it can also affect network traffic at each site. It doesn't make any difference if you're designing a small two- or three-site VPN, a dial-up VPN for hundreds or thousands of remote users, or a huge international VPN; proper planning will help you prepare yourself and your fellow users for deployment and use of the VPN. Proper design, one aligned with your current and future needs, will also help you deal with any problems that might come up along the way.

Although it's difficult to anticipate the special requirements of each type of network, this chapter attempts to cover as many VPN design issues as possible. To achieve that goal, you'll find that we often focus on problems and issues that only larger installations are likely to face. Although these issues may not be especially pertinent to those of you building smaller VPNs, they can prove useful if you're planning to increase the size of your network in the future. What may seem like a small, inconsequential problem on your network now can easily become a monstrous problem as your network grows. Considering that net-work usage and technologies often grow by leaps and bounds and are often unpredictable, it's nice to have some idea of what to expect down the line.

To help you deal with the issues surrounding the design of a VPN, we'll break down the process into three groups of issues and suggestions. First is the needs analysis: what are the requirements for your VPN—for bandwidth, connectivity, applications, users, and so on? Then, we'll move on to many of the issues actually affecting VPN design, such as selecting an ISP, managing addresses, and security options. Finally, we'll cover some steps for deploying your VPN.

Armed with the information and questions presented in this chapter, you should be well-prepared for the following four chapters, which lay out many of the connectivity and product options that have become available for creating VPNs.

Determining the Requirements for Your VPN

In order to design a usable VPN, you need to have some idea of the demands that will be placed on the VPN; in other words, what kind of traffic will be transmitted, what applications will be used, how often will the network be used, and so on. Also since one of the major components of a VPN is security, you have to factor in the type of security you'll require—for data, applications, and computers, as well as

users. Many of the questions and answers are interrelated, but we'll attempt to keep them categorized in some logical fashion.

Let's take a look at site-related network needs first. Some of the questions you'll need to answer for each site include the following:

- How many users are there at each site?
- What kind of connection to the Internet will the site require? Will it be a full-time or on-demand (i.e., dial-up) connection?
- How much network traffic does this site generate? How does the traffic vary hourly and daily?
- If a full-time Internet connection is required, what's the minimum uptime the site can tolerate? Might a second connection be required as backup?
- If an on-demand connection is required, how often will it be required? What kind of reliability is needed? (That is, can busy signals be tolerated? How often?)
- Will the site have to support remote users? How many?

You've probably guessed from some of the questions we raised that network capacity planning is an important issue for VPNs. Actually, it's an important step for just about any major change in your network, whether it be setting up a new WAN link or upgrading servers or routers.

To take capacity planning a step further, you need to know more about the type of traffic generated on your network and the applications that generate the traffic. Let's review some of the pertinent details about different types of applications and the traffic they generate before continuing with our VPN needs analysis.

It's often been the case that corporate LANs have sufficient bandwidth to handle most types of traffic and applications. That was especially true when mainframes and minicomputers held the majority of the data that users required. Accessing that data via terminals or client-server applications led to fairly predictable traffic patterns on networks.

But, that's changed considerably as the World Wide Web and other collaborative applications have become more dominant on many networks. Traffic patterns have become more chaotic and less predictable, with more and more traffic crossing organizational boundaries and their associated subnets, both within and between businesses.

To further complicate the analysis of network capacity, the usage of new types of applications has started to grow. In particular, applications that depend on real-time interactions, such as video conferencing, IP telephony, and other multimedia applications, are becoming more popular. And, they put new demands on both network bandwidth and latency.

Even as you get a handle on these applications and their network demands for your corporate LANs, you have to factor in how this traffic can be accommodated by your WAN links. This probably will have the greatest effect on your VPN plans, because the architectures of your WAN and VPN are likely to be largely the same, at least for LAN-to-LAN VPNs.

WAN links traditionally have less bandwidth than LAN pipes, partly because less traffic is expected to flow between sites on a WAN. Since your WAN links are likely to be a determining factor in the efficiency of your VPN, it's crucial to know what kind of traffic travels on your WAN. With that

information in hand and knowing what kind of uses will be reserved for the VPN, you can determine whether existing WAN bandwidth will be sufficient for your VPN or whether you'll need to upgrade some of the WAN connections. (We'll see later how bandwidth management and QoS enter into the picture.)

Applications and Traffic Types

Because the Internet is a massive conglomeration of different circuits managed by a variety of corporate and academic entities, there's a wide range in the performance of traffic on the Internet. Your traffic may not only be competing with other traffic for the same bandwidth or other network resources at some points in the internetwork, but it also may be subjected to delays that can affect the performance of your applications. Just as nature abhors a vacuum, users will always find ways to use any available bandwidth on a network. Even as new technologies like Gigabit Ethernet make it easier to provide more bandwidth, applications are gobbling up more bandwidth and placing restrictive demands on such data-delivery parameters as network latency and jitter. Thus, real-time data requires some kind of bandwidth reservation based on quality of service as well as priorities related to mission-critical situations.

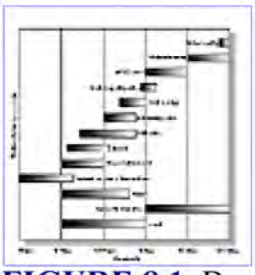

FIGURE 8.1 Bandwidth and latency requirements for different classes of applications.

Some of the simpler multimedia data, such as text combined with graphics, or animation files, do not pose special transmission problems on networks. These files may be larger than the norm, but they don't require synchronization of different parts of the data. But, more complex multimedia data, such as that used in interactive applications—videoconferencing and streaming video, for example—impose special restrictions on networks beyond demands for more bandwidth (see Figure 8.1).

Although bandwidth is the crucial factor when precise amounts of data must be delivered within a certain time period, latency affects the response time between clients and servers. Latency is the minimum time that elapses between requesting and receiving data and can be affected by many different factors, including bandwidth, an internetwork's infrastructure, routing techniques, and transfer protocols. Real-time interactive applications, such as desktop videoconferencing, are sensitive to accumulated delays, usually less than 0.2 seconds end-to-end. Interactive traffic, such as a TELNET terminal session or legacy protocols like SNA, can stand slightly longer latencies, on the order of one second or less. Bulk transfer traffic (an FTP file transfer, for example) can deal with any latency because the services have built-in measures for dealing with the acknowledgment of lost packets, rearranging packet sequences, and so on, but are not time-dependent.

Concerns about the bandwidth of the links to your ISP aren't restricted to LAN-to-LAN VPNs. Bandwidth can become an issue even if you're creating a dial-in VPN, because you'll need adequate bandwidth between your ISP and the VPN server site to handle the anticipated number of simultaneous tunnels from your remote users.

Let's move from site-related data to consider your entire VPN or corporate network. Unless you're planning a dial-in VPN that connects to only one central site, your VPN is going to connect a number of sites together. As you plan your VPN, you should not only have a list of the sites that will be served by the VPN, but you'll need to know their geographic distribution as well. Also, determine whether all the sites will need to interact with each other or if some sites can serve as satellites of other sites. Even though the Internet enables you to create a mesh between all sites, a hierarchy of site functionality and communications capabilities can lead to better traffic control than if you treat each site as the equal of all other sites (see Figure 8.2). Depending on your ISP's capabilities and POP locations, you may find that one architecture is less expensive than the other.

Geography also plays a role in the security of your VPN. If you're creating a multinational VPN, you'll no doubt run into some export restrictions on the cryptographic algorithms and key lengths that you can use for authentication and encryption. The U.S. government may eventually change its stance on restricting the export of long key lengths; some VPN products have been granted export licenses by the United States. In the meantime, be prepared to use systems that support at least two different key lengths and can pick between the two based on the destination of the traffic.

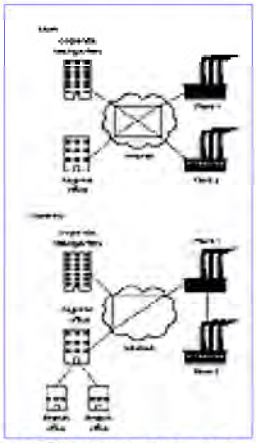

FIGURE 8.2 Mesh versus hierarchy.

While we're on the subject of security, you can improve your understanding of the security needs surrounding your enterprise's data by ranking the relative importance of the different data sources within your company and the effect unauthorized access would have on company operations. It should be

obvious that not all data is of equal importance to your company, but it may be less obvious that not all the data from one source (the CEO, for instance) is always of the same importance. *Warning: this is not a trivial exercise!* But, it could be important to help you decide what data needs to be encrypted for the VPN and what can travel in the clear. Ranking the relative importance of corporate data also can help mold corporate security policies.

Not only should the relative importance of the data be determined, but try to determine the timeliness of the data. Should two-year-old sales data be treated the same as last month's sales data? Probably not. And, it's unlikely that today's purchase orders need to be protected from eavesdropping and alteration for longer than it normally takes to fill those orders and receive payment. When you know the time period for which data needs to remain secret, or at least protected, then you can knowledgeably pick appropriate key lengths and cryptographic algorithms to protect that data. Not all of your data needs to be protected for 20 years, for example.

As you collect all of this data to lay the foundation for the specifics of your VPN, don't forget to try and get a feeling for how the corporation and its data needs will change in the future. Very few of us have crystal balls that work, but certain details in corporate plans can affect how you design your VPN. For example, it's worth knowing that the company plans to increase electronic communications with its business partners or suppliers; a logical course of action then would be to anticipate building an extranet using the VPN as a base. In fact, some writers and business people like to use the words extranet and VPN interchangeably. As we mentioned earlier in this book, we consider an extranet to be a special extension of a company's intranet; it doesn't have to be the same as a VPN.

One last note about extranets. (We'll cover extranets in more detail in Chapter 16, "Extending VPNs to Extranets.") If an extranet is part of the network plans, you'll eventually need to know the networking capabilities of your partners and what applications will be used, which means that someone will have to obtain information from your partners similar to what we've described in this chapter.

Some Design Considerations

You're likely to run into a number of common situations and caveats as you design your VPN. We've broken them down as follows: current network issues, security-related issues, and ISP issues.

Network Issues

We're assuming that you're not building your entire network from scratch; if you are, you're lucky! But, in most cases, you have to take into consideration previous network infrastructure decisions and equipment purchases that cannot be easily changed.

One of the network issues you should take into account in your design is the capabilities of your current routing and security devices. As we'll see in later chapters, it's possible to add hardware and/or software to your routers and firewalls so that they can serve as security gateways for your VPN. But, if your routers and firewalls are already maxed out and have no computational horsepower to spare for VPN functions, it'd be a mistake to plan on adding these functions to your existing equipment. If that's the case, you have three choices:

1. Upgrade your routers or firewalls so that they can support VPN functions.
2. Replace routers or firewalls with newer, more capable equipment.

3. Use a different type of device to provide your VPN services.

The range of hardware and software you can use for this last option is covered in Chapters 11, "VPN Hardware," and 12, "VPN Software." The network locations for devices supporting VPNs are shown in Figure 8.3.

Encryption is a computationally intensive process, but it varies according to algorithm, as Figure 8.4 shows. The strongest encryption available takes the most resources, which may make it unsuitable for many of your existing network devices. But, that doesn't automatically mean that your routers and firewalls cannot do the job; some vendors offer special cryptographic coprocessor cards for routers or firewalls to give them the extra horsepower they would need for a VPN.

FIGURE 8.3 Locations on network for VPN functions.

Other equipment you have to consider are your remote access servers and modem banks. One of the current driving forces behind VPNs is the desire to move the management, support, and equipment requirements of remote access from the corporate premises to an ISP. Some of you may well be planning a dial-in VPN only to achieve this, while others will be looking at hybrid VPNs that support both LAN-to-LAN traffic as well as remote tunnels to LANs. In either case, you've got some remote access equipment hanging around. Unfortunately, most of it will be of little use to your VPN. Some remote access servers can be upgraded to support VPN tunnels, but that has to be handled on a product-by-product basis.

Previous | Table of Contents | Next

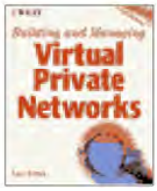

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

You may want to maintain a remote access server even as you deploy a VPN. If your company supports a large number of telecommuters who dial in via local, rather than long-distance lines, a VPN is not a cost-effective option. It will cost you more to set up ISP accounts for your telecommuters than to use your modem bank and remote access server to provide them access to corporate resources from their homes.

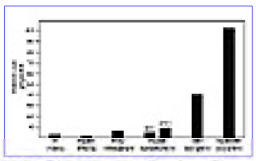

FIGURE 8.4 Computational requirements of cryptographic algorithms.

One component that can still be of use to your VPN is the authentication system you've been using for your remote users. Many VPN devices we cover in Chapters 10 through 12 can use such popular remote access authentication systems as RADIUS, TACACS+, and token-based authentication systems like Axent and SecurID. This compatibility enables you to continue using the same authentication systems as you convert from your current remote access servers to a VPN. In fact, many VPN vendors have recognized that many potential customers are reluctant to change from their existing remote access authentication systems and, therefore, need to support those systems in their products; more and more vendors are adding support for many of the authentication systems we just mentioned.

Keep in mind that your users' computers will be somewhat affected by the VPN. If you're planning a LAN-to-LAN VPN using security gateways, the encryption and decryption of VPN traffic will occur at the gateways, relieving the source and destination computers of that task. But, the encryption/decryption process may introduce some latency that will be noticed by some time-dependent applications. No general agreement exists on the amount of latency that security gateways introduce. If this is likely to be a factor for building your network, it's likely that many of the hardware products covered in Chapter 11, "VPN Hardware," will produce shorter latencies. The cryptographic coprocessors for routers and firewalls that we mentioned earlier will most likely keep latencies low as well.

Remote users are most likely to be affected by the tunneling protocols and encryption algorithms you pick for your VPN. In most cases, those users have nothing to rely on except their laptop to perform encryption and decryption, and their sessions on a VPN are likely to seem somewhat slower than when they don't use a VPN. The weak encryption provided by PPTP and IPSec (only when using short key lengths) might be preferable for remote users, but remember to factor in the importance of the data they're transmitting when making your decision on protocols and key lengths.

Multiplatform issues also can arise when picking the software for your remote clients. Most products

support the more recent Windows operating systems (NT 4.0 and 95), for example, although not all support Windows 3.x or the Macintosh OS.

Two very important issues that need to be resolved in your VPN network design are how the network is to handle routing and name resolution. You can follow one of two directions. First, you can view the network as a single network covering all of your sites. Or, you might choose to consider each site of the VPN as a separate network, joined together by tunnels.

In the first case, you can employ full routing between the different parts of the network connected by tunnels, plus you can set up a single unified name space for the entire corporate DNS. But, your own enterprise network structure may prevent you from implementing such an approach. Much of the problem stems from allocations of IP addresses and the use of *Network Address Translation* (NAT). Companies cannot often acquire a large set of contiguous IP addresses when they want to connect to the Internet, due to the allocation procedures for different classes of IP addresses and the scarcity of IPv4 addresses on the Internet. As an alternative, they could privately assign any addresses they pleased for the internal IP networks and use NAT (see Figure 8.5) to handle translation between the private addresses and a smaller range of addresses that were allocated for public (i.e., Internet) use.

You already may see the problem this approach causes for VPNs. When you attempt to connect these sites with tunnels, it's highly likely that two (or more) networks may have the same addresses, which will break routing services and other network functions.

We'll cover NAT and address management in more detail in Chapter 13, "Security Management," but it's enough to say that there are no simple solutions for combining two or more sites that have privately allocated IP addresses that overlap.

A related issue is that of DNS. It's not unusual to shield DNS entries for internal resources from external uses, but you'll need a way to provide this information to other sites connected via your VPN tunnels. One approach is to limit the number of hosts that can be reached by other parts of the VPN, but this means maintaining dual DNS entries, one set for internal site usage, the other for VPN usage. If NAT is used for translating private addresses, you may have to use DNS spoofing as well. Fortunately, some firewall and VPN products support DNS spoofing, which will make your job a little easier.

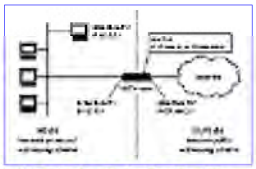

FIGURE 8.5 Example network using NAT.

As part of the solution to potential addressing problems, particularly for larger enterprises and internetworks, you should consider your company's plans for upgrading networks to IPv6, because its large address space can solve many VPN-related problems. Demand for this next generation of IP is slowly growing, and vendors are offering a smattering of products that support IPv6. This migration is likely to be more of a long-term effort rather than something that can be accomplished within the next year. You might not be able to use IPv6 today, but it's something to keep in mind; remember that IPv6 implementations also will include built-in support for IPSec, which may simplify your deployment of client and host software later.

Security Issues

Protecting your data as it travels across a VPN is only one part of ensuring its security. VPN security design, therefore, has to be treated as part of the broader issue of corporate security policy. In general, a corporate security policy should focus on determining who has access to what resources. A good starting point is RFC 2196, the *Site Security Handbook*.

Portions of your existing corporate policy (assuming you have one) may directly impact how you handle your VPN. For instance, policies on user passwords—how often they're changed and so on—already may be enforced for remote access and can be directly translated to VPN access rights and their maintenance.

Access rights is another issue that needs to be extended for VPNs, because you'll have users gaining access to subnets and devices they probably otherwise would not see on the network. In general, a tunnel lets a user onto the network without any restricted access. Depending on which tunnel protocol and which network operating system you're using, you can allocate access rights to the tunnel's user as the tunnel is set up. For instance, many tunnel servers and security gateways use NT user domains for access control, which enables you to control access to Windows network resources. As policy-based management for all types of networks becomes more widespread, this sort of control will extend to all systems and networks.

Previous | Table of Contents | Next

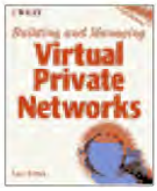

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Another factor affecting your choice of running a CA in house or outsourcing it is your plan for dealing with business partners via an extranet. An in-house CA may prove ideal for issuing digital certificates and keys to your employees, and it also may be suitable if a small number of extranet partners are involved. But, if the extranet involves larger corporate partners, these partners may have their own CAs, and you'll have to find a way of getting the CAs to work together. An outsourced CA may make this easier, but you also can choose to have your internal CA certified by another, external CA, making cross-certification and validation of digital certificates possible outside of your company (see Chapter 4).

The choice between in-house maintenance of a CA versus outsourcing this function underscores an issue that will come up again later in our discussion of ISPs: Who manages your security? Does your company feel that it alone can maintain proper security of its resources? Does it have the appropriate personnel and sufficient resources? Or, is it willing to entrust some of its security to a second party—in this case a certificate authority? Admittedly, the major CAs have spent millions of dollars and large amounts of time to protect their systems, which should make their services quite trustworthy. But, selecting either option comes down to a question of control.

One last note about key management: Be sure to include some type of key recovery mechanism when selecting a key management system. There are good reasons for having a key recovery system in place. As long as critical data is going to be encrypted using public-key systems, situations will arise in which your company may have to recover old data when the original key is no longer available. This situation might happen when employees leave the company, either voluntarily or otherwise, or when someone dies, for instance. In such cases, having a third key that can be used in place of the "lost" public key enables you to recover older protected data.

ISP Issues

It may seem obvious, but you cannot overlook the capabilities of your ISP when you're designing your VPN. After all, we've been talking solely about private networks that are using the Internet as the "plumbing" for the networks in this book. If you don't have a connection to the Internet, you can't have a VPN, at least using the definitions we've adopted for this book.

ISPs can be involved in a VPN in a number of ways. Using PPTP and L2TP for tunneling enables ISPs to offer value-added services as tunnel initiators and proxies for user authentication for your VPN. Other ISPs also are offering full-fledged outsourced VPNs, including security management and installation of the appropriate VPN equipment on your premises (see Chapter 9).

Bear in mind that your current ISP doesn't have to be your VPN provider; you can use two or more ISPs to handle your Internet traffic, and you could use one for open traffic, with a second for your encrypted

traffic.

One reason for using different ISPs might be their geographical coverage. If you're designing a VPN for a multinational corporation, you may need POPs in some countries that very few ISPs service. This issue becomes more important if you're more concerned with LAN-to-LAN tunneling than remote access to a VPN. Multinational remote access to a VPN can be set more easily now, thanks to the new roaming services that have been instituted.

Basically, a roaming service lets travelers or other remote users access the Internet via a local ISP rather than dialing long distance to log on through the corporate ISP. A broker service manages the settlement charges between ISPs and provides client-based access software, including a phone book with a list of local POPs. Initial services have concentrated on PPTP and, to a lesser degree, L2F; future expansion is aimed at IPSec and L2TP.

If you have a large number of remote users to support on your VPN, then roaming services offer you a reasonable alternative to selecting only one ISP for remote access. These services also give you and the VPN users more flexibility as new sales areas and branch offices open up in previously uncovered regions.

Another step to increased flexibility, as well as reliability, is to use more than one ISP for your Internet connections, even for the main sites of your VPN. Although it doesn't happen everyday, ISPs have been known to loose Internet connectivity, sometimes for a day or so. If you're planning to transmit mission-critical data on your VPN, running all your traffic through a single ISP without any backup connection isn't wise. We'll see in the next chapter that *Service Level Agreements* (SLAs) can be negotiated with ISPs to provide refunds when service is lost, but you would rather not lose the connection in the first place.

If you're planning to use PPTP or L2TP to construct a VPN in conjunction with an ISP, then you'll need to know the capabilities of their equipment and how they handle security. (Many of questions revolving around these issues are covered in the next chapter.) If an ISP is going to maintain a proxy RADIUS server for your users, then you want to be sure that the service is secure against unauthorized access, both from outside the ISP and within the ISP—either from ISP staff or other customers of the ISP. We earlier raised the issue of who controls security, in reference to CAs and digital certificates. That question now comes up again, in the context of other forms of user authentication. The recommended approach is for your corporation to maintain a RADIUS (or similar) server for authentication of its employees and to let the ISP's proxy server obtain its database of access rights from that server.

Depending on the uses planned for your VPN, performance can have a number of meanings. At the very least, performance is providing the required bandwidth as needed, when it's needed. Many applications, such as e-mail, FTP, and Web browsing, can function properly with this minimal definition. But, if you're planning to run transactional data or real-time interactive applications over your VPN, then network latencies can become an issue as well. Many of the details surrounding SLAs and monitoring ISP performance are presented in the next chapter.

Providing quality of service on the Internet is still relatively new (see Chapter 15, "Performance Management"). In fact, it's not something you expect on many parts of the Internet and probably won't get for a few more years. ISPs willing to provide some type of QoS and bandwidth management as part of their performance guarantees will do so only for the traffic that flows on their network. That may not

be a problem for your VPN if it's based on a single ISP that runs its own network and can segregate its customers' traffic from other Internet traffic.

But, if you're planning on expanding your VPN to an extranet and your applications require performance guarantees and QoS, you'll most likely be out of luck for the next few years. You could construct an extranet with guaranteed performance if you and your extranet partners are all using the same ISP. At the moment, no one has proposed a way of guaranteeing performances for traffic transmitted over multiple ISP networks. The first agreements of this kind probably will occur between large ISPs that own and control their own networks, such as AT&ampT and UUNET.

Previous | Table of Contents | Next

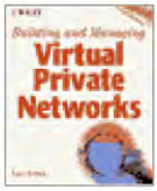

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Planning for Deployment

Although we haven't yet discussed the details of the different devices you'll use to make your VPN an actuality—that's left to Chapters 9 through 12—you can do a few things to prepare for the installation of your VPN.

First, there's the question of the current state of your business's security. It'll do little good if you create a VPN for the secure transmission of corporate data across the Internet if you leave open other ways for attackers to acquire or alter your data. In other words, you need to make sure that your corporation is secure against outside attacks even before you open up any access via a VPN. It's a good idea to initiate a security audit of your corporate security practices, usually by an outside firm, before you install your VPN. If there are weak spots, try to correct them before your VPN is operational. (Another hint: Have these security audits performed periodically. If you don't know where to start with security audits, look at Bernstein et al., *Internet Security for Business*, John Wiley & Sons, Inc., 1996.)

Deploying the various components of your VPN can be a particularly complex task, depending on the number of sites and users comprising the VPN. As we've mentioned previously, some VPN products allow for centralized configuration of all VPN security gateways, which at least ensures that the configuration files will be correct before distribution.

As much as possible, try to treat your branch offices alike. You can use a cookie cutter approach to the VPN products for each site as well as for much of the configuration. This approach not only simplifies the deployment of the VPN, but troubleshooting as well. You may find it difficult to use this approach, however, if each of the VPN sites have very little in common regarding network hardware and capabilities and if you find yourself purchasing different VPN equipment for each site—the real world doesn't always cooperate with your plans. On the other hand, maybe you can use the VPN as an excuse for making each site's network depend on common equipment and configurations.

We've already written much about the distribution of keys and digital certificates, and we'll have more to say in later chapters, but you should consider how users will store and use their digital certificates and how that will affect deploying certificates. You might want to combine your PKI with smart-card technology to provide a personal, flexible, and secure means of identifying people and their capabilities. Some systems are already available for combining smart-card readers with desktop and portable PCs. The combination of smart cards and digital certificates may well prove popular alternatives to other token-based security systems, like SecurID, as more uses for digital certificates make their way into the mainstream.

We're firm believers in pilot projects; before you roll out a corporate-wide VPN, you should plan on

trying a smaller test network. If possible, this network should still function under normal operating conditions but should not include much, if any, mission-critical information. Use this test network to work out the bugs in your configuration and management of the VPN before it becomes an integral part of the rest of the company.

It's also a good idea to change corporate traffic over to the VPN in stages. Bulk traffic, such as e-mail and file transfers, should be shifted first, while transactional and other real-time traffic should be done later once the characteristics of your VPN are known. We started this chapter with a discussion of network capacity planning and performance analysis, and we end the same way. Learn what are the effects of your VPN on your applications, and vice versa, before transferring all intersite traffic to the VPN.

Summary

VPN design needs to take into account not only security issues, but also the bandwidth and latency requirements of your applications as well as national restrictions on cryptographic key lengths.

Deciding between adding software to existing network devices, such as routers and firewalls, and purchasing new devices specifically designed for VPNs depends on the need for performance as well as cost constraints, because encryption is a computationally intensive process. Be sure to factor in the importance of the data being transmitted to determine the period over which it must be protected; not all data has to be secure for years, for instance.

Even if new VPN devices are installed, some existing services, such as authentication servers for remote users, can be adopted for use on the VPN.

One of the most important issues in deploying VPNs and the authentication of users and security gateways is the selection of the infrastructure for distributing digital certificates. Companies can choose to set up their own certificate authority in-house or outsource the operation to a recognized CA. Expect that more uses for certificate-based authentication will arise, making the certificate authority a more important part of your security system as the years pass.

Previous | Table of Contents | Next

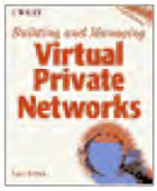

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

PART III
Building Blocks of a VPN

A VPN consist of two main components: The Internet connectivity provided by an ISP and the hardware and software that protects your data by encrypting it for transmission over the Internet. VPN functions can be performed in firewalls, routers, and specially-designed hardware, making deployment of VPN devices relatively easy.

Since a VPN is mission-critical to your business, you should ensure the best possible performance from your ISP, by a guaranteed Service Level Agreement (SLA). While the SLA should define the expected throughput and maximum delays tolerated, you need to plan how your company will monitor network performance to ensure compliance. Also, when selecting the devices that will perform encryption and tunneling for your VPN, you should match your expected WAN throughput with the capabilities (such as speed of encryption) of the devices.

CHAPTER 9
The ISP Connection

Whatever the design of a Virtual Private Network, its success depends greatly on one element—your *Internet Service Provider* (ISP). Because your ISP is responsible for the transmission of your data over the Internet once it leaves your sites, it's important that you have a good working relationship with an ISP that you can trust. Establishing a good working relationship with a service provider depends not only on knowing what your network needs, but also knowing what the service provider can deliver.

But, you can't run your company on gentlemen's agreements alone, so a documented agreement regarding service provider performance, reliability, and liability helps keep relations and expectations on business-like terms. These agreements, *Service Level Agreements* (SLAs), help define what both parties—your ISP and your company—expect of a VPN and, most importantly, the Internet portion of the VPN.

This chapter will cover the different aspects of an ISP's role in VPNs. We'll start out discussing the details of an ISP's capabilities for handling Internet traffic and the requirements that an ISP should be able to fulfill for VPNs. This discussion will be broken down according to the customer's desire to either outsource most, if not all, of the VPN or do most of it in-house. Then we'll go over the details of SLAs, including how they can be monitored and enforced. The last part of the chapter will include an overview

of some of the current VPN services offered by ISPs and NSPs, both here in the United States and internationally.

ISP Capabilities

Before we discuss what services an ISP can provide, let's review the way that ISPs are classified according to their capabilities and hierarchy within the Internet structure.

Types of ISPs

The service providers whose networks make up part of the Internet are classified in tiers according to the capabilities of their networks and the type of Internet connectivity that they provide (see Figure 9.1).

Tier One providers, such as AT&ampT, GTE Internetworking, IBM, MCI, PSInet, and UUNET, own and operate private national networks with extensive national backbones, often architected like the schematic network shown in Figure 9.2. These independent networks meet and interconnect at the Internet *Network Access Points* (NAPs). In other words, the networks interconnect and exchange traffic at the NAPs to form what is essentially the Internet.

FIGURE 9.1 The hierarchy of Internet providers.

The independently created national networks set up by companies like PSInet and UUNET, among others, mostly tie into the NAPs. Some service providers have made their own arrangements for exchanging Internet traffic by sidestepping the NAPs, which can be bottlenecks. These *peering points* help relieve some of the load at the NAPs. The independently created national networks also give these multiservice providers an added advantage when offering services like VPNs to their customers because they can control the traffic that runs on their network and the reliability of the network better than if a series of networks were be used to handle your traffic. (We'll see later that service providers will often offer guarantees of bandwidth and latency as long as traffic is restricted to their network.)

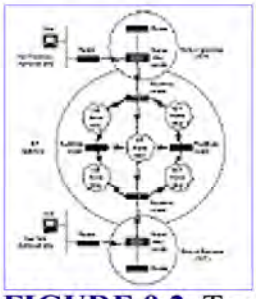

FIGURE 9.2 Typical ISP backbone design.

Note that none of the Internet NAPs provide Internet connectivity to the general public or to business and industry. The NAPs are only points for the orderly exchange of traffic between those organizations that maintain extensive national backbones; a NAP is not a point-to-purchase Internet access. Additionally,

connections to the Internet NAPs are made at a minimum of DS-3 speed (45 Mbps).

A Tier Two provider is a company that buys its Internet connectivity from one of the Tier One providers and then provides residential dial-up access, provides World Wide Web site hosting, or resells the bandwidth. These regional providers typically operate backbones within a state or among several adjoining states. They also may be connected to a NAP, but usually no more than one NAP.

Below Tier Two providers are the individual Internet Service Providers, which can be anything from two or three persons running a dial-up POP to much larger operations supporting as many as 100,000 dial-up customers, for example. These providers generally don't operate a backbone or even a regional network of their own. If they offer national service, they use the POPs and backbone structure of a larger backbone operator with which they're associated.

For a business with a full-time link to the Internet or an individual working at home or on the road and dialing into the Internet, the ISP's POP is an important cog in the use of the Internet. The POP is where the ISP handles the different types of media that its customers use for Internet access and from where the ISP forwards all the customer traffic to its backbone network, which connects to the rest of Internet at some point (see Figure 9.3).

Some POPs contain different equipment for each transmission media they support, such as a modem bank for dial-in sessions and CSU/DSUs for frame relay and *Digital Data Service* (DDS); other ISPs have opted to leave support for the different media to the public network, instead running a leased line to their POPs. In addition to handling different media for customer traffic, the POP includes routers and/or switches to connect the POP's local LAN to the rest of the ISP's network as well as network-management consoles. In some cases, the POP will include servers for hosting mail, news, and Web sites, and RADIUS authentication servers for an ISP's customers.

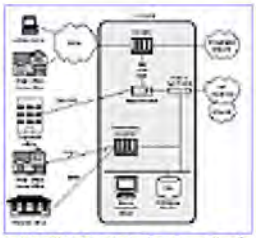

FIGURE 9.3 Schematic of a typical ISP POP.

The fundamental service of an Internet Service Provider is connectivity to the Internet. This connectivity can take the simple form of providing dial-up access for individual users with a modem or ISDN line, or it can be dedicated lines (T1 or T3, for instance) running from your corporate LAN to the ISP's *Point-of-Presence* (POP) and then to the rest of the Internet.

Previous | Table of Contents | Next

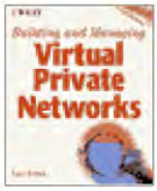

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

The entire range of connectivity options offered by an ISP is important to your VPN design because you're likely to have a range of requirements that differ from site to site and user to user. For example, a large corporation would most likely want high bandwidth connections, such as a T3 line, for its corporate headquarters or main computing center, regional offices could get by with T1 lines, and branch offices might need only ISDN or a dial-up line with a modem. Telecommuting workers might use ISDN or modem lines at home, and workers on the road would likely rely on dial-up access.

Simply having an ISP provide the "pipes" to the Internet may suffice for your VPN design if you're planning on doing everything else in-house. But, many ISPs have recognized the benefit of offering value-added services built atop Internet connectivity. These services can make the design and maintenance of your VPN easier; for example, ISPs offer managed security services using firewalls for protecting your sites as well as full VPN installations, including all necessary CPE equipment such as routers and CSU/DSUs, monitoring software, and authentication servers.

Another thing to keep in mind is that the ISP's help desk and support staff will either become an extension of your own corporate help desk or a replacement for some of its functions as you meld the ISP's services into your VPN. For instance, you should expect the ISP's help desk to handle any problems that your mobile users and telecommuters may have when dealing with modem or ISDN access to the Internet. Plus, whether you use the ISP simply for Internet connectivity or outsource more of your VPN to the provider, the ISP should have a tiered structure for dealing with network problems.

Because the Internet is a massive conglomeration of different circuits managed by a variety of corporate and academic entities, the performance of traffic on the Internet varies greatly. Your traffic may not only be competing with other traffic for the same bandwidth or other network resources at some points in the internetwork, but it also may be subjected to delays that can affect the performance of your applications.

Even as new technologies like Gigabit Ethernet make it easier to provide more bandwidth, applications are gobbling up more bandwidth and placing restrictive demands on such data-delivery parameters as network latency and *jitter*, the variation in latency. Thus real-time data requires some kind of bandwidth reservation based on quality of service as well as priorities related to mission criticality.

With the ever-expanding move to multimedia, more applications now require control of the quality of service they receive from the networks. To support the different latency and bandwidth requirements of multimedia and other real-time applications, networks can use QoS parameters to accept an application's network traffic and prioritize traffic relative to other QoS requests from other applications. QoS provides network services that are differentiated by their bandwidth, latency, jitter, and error rates.

Even if your current application needs do not include a guaranteed latency or prioritization, keep these requirements in mind because they're likely to become more important in the future as applications

change. Even though ISP treatment of QoS is in its infancy, provisioning QoS will become a fact of life among ISP offerings in the near future. We'll cover some of the details in Chapter 15, "Performance Management."

What to Expect from an ISP

If you're shopping for an ISP for your VPN, you should use a few criteria right at the beginning of the selection process before we get into more detailed requirements. Not all of you will have the same requirements of an ISP; because different businesses have different Internet requirements, they can be met by different levels of ISPs.

As for the initial screening criteria, first there's the issue of geographical coverage offered by the ISP. For instance, a multinational corporation probably would find its requirements met by a global service provider, but not by a local or regional ISP. On the other hand, a company with offices only in California or Texas might find that all its needs are met by a regional or local ISP.

Second, there's the type of access your company requires. If you're planning only a dial-in VPN, you'd like to be sure that the prospective ISP can provide both sufficient modem ports and POPs for your workers and have POPs in the areas from which your workers are likely to call. In some cases, this latter requirement is not a severe constraint: Some companies like GRIC Communications and iPass offer a roaming service that allows a series of ISPs to cooperate to offer wider dial-in access. In other words, you can have an account with one ISP but use the POPs of other ISPs in the roaming service when local coverage from your original ISP isn't available. Roaming services can be especially valuable for overseas travellers because they not only keep dial-in costs down for calls back to the United States, but also usually offer more reliable connectivity than many long-distance lines in foreign countries (even among those countries in Europe that have modern, reliable PTTs within their own country).

Third, are you planning to design and implement your VPN entirely with in-house support, or do you want to outsource some, if not all, of the VPN to a service provider? If you're planning on an IPSec-based VPN, any ISP can provide you the connectivity to the Internet that you require (assuming, of course, that they can meet your bandwidth, latency, and location requirements). Not all ISPs have the equipment to handle PPTP or L2TP systems, though. And, only a few ISPs can offer you a turnkey VPN customized to your needs, although the number is growing. We'll get into more details about outsourced VPNs later in this chapter.

Fourth, what are the future plans for your business and the VPN? If your business grows and adds more sites, can your prospective ISP accommodate your growth? Does the ISP cover the geographic areas you're expanding into? Or, perhaps you're planning to open your VPN to partners, distributors, and suppliers, forming an extranet. We'll see later in this chapter that ISP performance guarantees currently cover only traffic serviced by a single ISP and not cross-ISP traffic. If guaranteed performance is important to your partnerships and the extranet, then your choice of ISP may be influenced by which one you and your partners can use.

These initial selection criteria should help you pare down your list of possible ISPs to a select few that you can investigate in detail. The following section lists many of the details about an ISP that you should include in your investigations as you narrow down your search for the appropriate business partner for constructing your VPN.

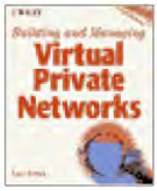

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Learning an ISP's Capabilities

This section contains a fairly extensive (although by no means exhaustive) list of issues that should be raised with any prospective ISP when setting up your VPN. If your company is like most medium- and large-sized businesses in the United States, you already have an ISP that provides you with connectivity to the Internet. If so, you might choose to quickly skim the following set of issues. On the other hand, if you're unhappy with your present ISP or plan to use a different ISP for your VPN, you probably should look over the rest of this section. Although many of the items listed here are applicable to any ISP service and not just VPNs, they do impact the provisioning of value-added services. As you'll see, only a few issues specifically address VPNs or security.

ISP INFRASTRUCTURE

Your first concern should be the ISP's network backbone, because it's going to determine how well your network traffic is handled. The best designs are a full mesh with multiple redundant paths between transfer points; redundant routers and/or switches also should be installed at each of the major transfer points in the network. Each router or switch location in the network should meet data-center quality for environmental controls, including such items as redundant backup power and air conditioning. Maintainability is improved if standard equipment is used, not equipment that is custom-designed and may be hard to replace quickly.

Although the current state of the Internet is such that you'll receive the best VPN services from an ISP with its own national or international network, some VPNs (such as dial-in VPNs and those not requiring low latencies for their applications) can be created to handle traffic that crosses ISP boundaries. Furthermore, you may need to balance VPN traffic with other nonsecure traffic that involves your business on the same ISP, requiring that the reach of your communications involves more of the public Internet. Whatever the reason, moving beyond a single-network situation means that you should pay attention to how your ISP exchanges traffic with other providers. For instance, does the ISP's network extend to all the major NAP peering points and does the ISP have full peering rights with the other major providers?

NETWORK PERFORMANCE AND MANAGEMENT

Aside from the backbone's design, you should understand how the service provider has provisioned bandwidth and how your bandwidth requirements will be treated. For example, how much bandwidth has the service provider already committed to other customers? Also, is the bandwidth you require maintained throughout the system? For example, if you buy a 10-Mbps circuit, does the traffic become

aggregated onto higher bandwidth lines? Is the bandwidth always available? Is it available as a burst speed, which lets you occasionally transmit more traffic than the average load? Ask whether the service provider publishes statistics on the network's traffic loads as well as its reliability statistics.

It's a fact of life that all ISPs occasionally have outages. Find out how long an ISP's outages have lasted and what percentage of the systems and users were affected.

Quality of Service or guarantees of prioritized delivery of different traffic classes can be important when you want to ensure that important traffic always makes it through your network, even if other, lower-priority, traffic doesn't. But, providing QoS on the Internet is a relatively new service that's still largely being handled as an experiment by ISPs. Again, an ISP can more readily guarantee QoS for traffic that is transmitted only on its own network rather than over multiple ISPs. That will no doubt change in the future as policies for defining QoS among ISPs are ironed out. But, if you're interested in differentiating your corporate traffic for different delivery priorities, you should check whether the prospective ISP offers any QoS guarantees.

Proper operation of a service provider depends on an efficient *Network Operations Center* (NOC) that is fully staffed 24 hours a day, 7 days a week (24 × 7). Again, like the transfer points of the network, the NOC should be housed in an environmentally sound facility, including backup power and air-conditioning as well as earthquake protection. It should also be a secure facility and have written plans for dealing with detection of security breaches and procedures for dealing with breaches. A standard system, such as HP Open View or Sun's SunNet Manager, is a good start for monitoring the network but it's also nice to see whether the NOC has developed other tools for monitoring and troubleshooting. Check and see whether the facility has undergone any form of recognized audit.

CONNECTIVITY OPTIONS

Most ISPs specializing in business-to-business services sell a full range of connectivity options, with bandwidth products ranging from 56 Kbps through T3 speeds being common. When planning for a particular bandwidth connection, determine how the connection is actually handed off to you. For example, some ISPs supply a T1 line in an Ethernet format, allowing easy integration into your network. Others present the T1 link in a raw serial format, requiring gateway equipment to transform it into a protocol you can use.

To ensure that your connection has an adequate continuous bandwidth throughout the network system, ask your ISP for a network schematic, with bandwidths listed on each network segment.

Confirm what is included in the standard service price. Some ISPs require that you purchase the routers and CSU/DSU devices, while others will supply and manage them for you. The ISP-supplied equipment normally is configured, monitored, and diagnosed for problems via the ISP's NOC.

Most ISPs have three costs in their access service: installation charges, basic connection bandwidth service you subscribe to, and the local loop charges required to connect your location to the ISP's *Point-of-Presence* (POP).

Find out what assistance the ISP will provide in addressing user connection issues, especially for your dial-in users. Also determine whether they have a tiered support system for their help desk.

SECURITY AND VPNS

On the security side of things, a professional security officer should be a member of the staff, and the ISP should have a written set of security policies. The ISP should have active monitoring tools protecting its own systems and have experts on the staff capable of configuring firewalls and monitoring devices.

Firewall-management services should include firewall selection, installation and setup according to your policy criteria, as well as round-the-clock monitoring of the firewall for attempted security breaches. An ISP should have a written escalation procedure for handling security breaches and guaranteed response times for emergency security breach notification. Reports should cover inbound traffic with attempted break-ins as well as outbound sites visited. Because the service provider will be responsible for the firewall, he should be responsible for updating the software when new vendor releases come out; this should be handled automatically whenever possible. But, because your company is responsible for setting the policies that the firewall is enforcing, you should have the freedom to periodically change the policy rules.

Previous | Table of Contents | Next

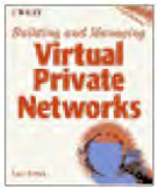

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Turning to VPNs, the details of the ISP's operations that you'll need to know will depend on how much of the VPN you want the ISP to handle. In other words, if the ISP is simply providing the pipe, many issues surrounding encryption keys and certificate authorities aren't pertinent. If the ISP is to handle a dial-in VPN for you, then questions about RADIUS proxy servers and authentication updates become more important.

Some of the important ISP issues you should resolve include: The encryption algorithms supported by the ISP's system, whether the system can switch between algorithms automatically, whether the system conforms to IPSec, and whether the ISP uses a system that is approved for use and export outside the United States. If you're going to use a system that follows IPSec protocols, then you should determine whether key exchanges are only handled manually or if they can be done automatically and how rekeying is handled (automatic rekeying is preferred).

Authenticating legitimate users of your VPN is always important, but the ISP's role in authentication is most important when you're supporting dial-in users, say with PPTP or L2TP as we discussed in Chapters 6 and 7. In such cases, you'll want to know what types of authentication the ISP supports. If you plan on using RADIUS, for example, then will the ISP's server act as a proxy server to your master RADIUS server? The ultimate management of security rights should rest in your company's hands, not those of the ISP. Even if you outsource your entire VPN to a service provider, you should determine and manage the access rights of your employees.

If you choose to use digital certificates to authenticate users and devices, you might want to use the ISP as a *certificate authority* (CA) for managing the certificates. Certificate management would include issuing the certificates and managing revocation lists, as well as maintaining a certificate server for verification of the certificates. Another option is to outsource the certificate management to another firm, other than the ISP. The issues for evaluating certificate management capabilities are the same regardless of the type of firm being reviewed.

If the ISP or another service company acts as a certificate authority, check to see whether your company can act as a backup or concurrent CA to further guard against failures. Be sure that the ISP treats the maintenance of its certificate server with the same care as you'd expect for other crucial network resources. In other words, the certificate server should be located in a secure environment that includes backup power and backup data facilities (usually located at another site).

We've already mentioned in other chapters that encryption methods can have an adverse effect on throughput because they're computationally intensive. Software-based encryption can often be slower than hardware-based encryption, for example, and neither may be able to keep up with the demands of a high-bandwidth pipe. Find out how the ISP implements encryption and what's the maximum throughput

the system can handle; ask for some benchmark data.

Security breaches and breakdowns can happen at any time. Check with the ISP to determine whether the NOC provides around-the-clock monitoring of the entire encryption system. Also find out who's responsible for changing any system firmware—you or the NOC?

One of the advantages of a VPN is supposed be its flexibility; you should be able to easily and quickly add new sites to your VPN, for example. Find out how easy it is to add offices to a VPN. How long does it take for the ISP to add a new site? Are there any written guarantees on maximum times for establishing a new link?

Even though one of your primary concerns in designing a VPN is securing all of its traffic, you also might want to let your users communicate in the clear with other sites that are not a part of the VPN. See whether the ISP's architecture allows that; if so, can users distinguish between plaintext and secured communications easily, and can they switch between the two modes easily?

As always, reports and accounting are important to monitoring the VPN. See whether the ISP provides throughput reports for VPN traffic, and with what frequency (daily, weekly, monthly?). What type of accounting and billing reports does the ISP offer? Does the ISP provide reports on a user-by-user basis or site-by-site basis, for example?

Service Level Agreements

Whenever you're planning to outsource part of your network's operations to another firm, you need some way of ensuring that your expectations regarding network performance, maintenance, and problem resolution are met. An increasingly popular method for documenting your expectations and what a service provider is willing to provide is the *Service Level Agreement* (SLA). Service Level Agreements are a relatively new development in the telecommunications world. SLAs originally were designed for private voice networks and later extended to frame relay data services. Now we're seeing SLAs applied to the world of Internet VPNs, with good reason—everyone needs a way of determining what kind of performance guarantees they're getting for their VPN dollar.

The SLAs that document service-level guarantees have one main purpose: They help keep conflicts between you and your service provider to a minimum by setting reasonable expectations of service. SLAs benefit you, the client, by providing effective grading criteria and protection from poor service. They benefit the service provider by providing a way of ensuring that expectations are set correctly and will be judged fairly. Remember, SLAs include some kind of monetary reimbursement for lost or poor service, but that's a last resort; you'd really rather have good service than compensation for poor service.

Three basic items should be covered in every SLA: availability, effective throughput, and delay. Other items to consider include the mean time to respond to problems and the mean time to repair or restore service.

Network availability is a simple measure of the uptime of the network links available to you, complicated only by the fact that it's measured over all your sites. If you measured network availability over a month's time, the formula would look like this:

$$\frac{(24 \text{ hours} \times \text{days in month} \times \text{number of sites}) - \text{network outage time}}{(24 \text{ hours} \times \text{days in month} \times \text{number of sites})}$$

Even for so simple a measure as network availability, check to see what's included in the service provider's definition. Availability guarantees should include all components of the provider's network, the local loop to the network, and any CPE equipment provided by the service provider (such as a CSU/DSU and router). Excluded items may include a customer-provided CSU/DSU, router, or other access device; the local loop when provided by the customer; network downtime caused by the carrier's scheduled maintenance; customer-induced outages; dial-in links; and *acts of god.*

Previous | Table of Contents | Next

Note that there's an important distinction between network-based availability and site-based availability. For a network consisting of 10 sites, an average network availability of 99.5 percent would allow 36 total hours of downtime in a 30-day month. If the SLA is written around a site-based availability instead of being network-based, then any one site can be down for only 3.6 hours in the month. The distinction can be very important when computing downtime.

When dealing with measurements of throughput, traffic load and delay should be measured when the impact is at its highest (i.e., at times of peak traffic load). Because service providers will often exclude certain data, such as data loss during provider maintenance, dial-up lines, or new circuits added during the month, be sure that you understand which data has been included in any measurement so that your own cross-check measurements will correspond to those performed by the service provider.

Unfortunately, as this book was being written, no standards had been created for Internet SLAs, so you'll have to compare what each ISP offers. Even in the frame-relay world, where SLAs are a more mature feature, no standards exist, although many of the metrics quoted in SLAs are the same from provider to provider.

Any of the SLAs currently offered by service providers covers only single-ISP traffic, because that's really the only traffic that the service provider can hope to control. It'll probably be some time before we see SLAs that cover multiple-ISP traffic, because both policies and technologies have to be developed further before the ISPs can routinely work together on guaranteeing reserved resources that affect such things as latency and other QoS parameters. To start, all service providers will have to agree on how to measure availability, delay, and packet loss.

Many service providers are now offering SLAs for their services, but most of these SLAs are not negotiable unless your company is linking more than one or two sites; frame-relay carriers are usually willing to negotiate when a customer has 10 or more sites, and this rule-of-thumb probably is applicable to ISPs as well. But, whatever you do, if you're in the position to negotiate an SLA, don't attempt to negotiate unless you have some background on your network's performance and needs and what level of WAN service is needed for your users.

Preparing for an SLA

As you prepare your company for a Service Level Agreement, you can follow some steps to see that you get the best possible SLA. (These steps are based on similar ones detailed for frame-relay SLAs in a white paper developed by TeleChoice and Visual Networks.

1. Continuously determine what WAN service levels are needed.
2. When service levels are established, verify them. You need to monitor performance in

real-time, review historical performance, and assess any quality-affecting trends that you find.

3. Baseline your network. Understand your applications, peak times, and areas of concentration.
4. Negotiate SLAs if at all possible. Read the fine print and do the calculations. If the negotiated network availability guarantee is 99.5 percent, how many hours of down-time does that mean for the network per month?
5. Formulate a plan for monitoring your service provider.
6. Analyze your network's performance and reliability on a weekly basis.
7. Compare your own measurements of the network's statistics with your service provider's reports every month.

Monitoring ISP Performance

Whether you accept a standard ISP-provided SLA or spend a great deal of effort negotiating a custom SLA, an SLA will mean little if you don't have some means of monitoring the service levels specified in the SLA.

A network management system has many components, as shown in Figure 9.4, but four are particularly important for verifying your provider's performance:

1. Monitoring devices located at the edge of the provider's network.
2. A database to gather information on performance.
3. Applications designed to analyze data and issue reports specific to each customer's use of the network.
4. Web-accessible HTML versions of the reports.

A few key implementation issues have a direct impact on the usefulness of SLAs to the network manager. The first issue is where the measurements are taken: end-to-end or just within the ISP's network cloud (see Figure 9.5). The local loop can have a profound impact on network performance, but it is ignored in a switch-to-switch implementation. Performance measurements and troubleshooting must be performed end-to-end.

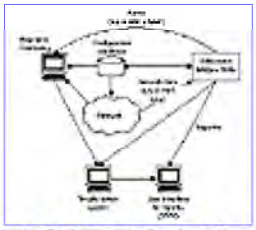

FIGURE 9.4 Network management and monitoring.

The second issue is utilizing a measurement system that is independent of the network you are measuring. Use an objective system that is not biased towards either switch or router architectures. Also keep in mind that how this information is presented is almost as important as the information itself.

Agreeing on definitions of measured parameters and how they're measured is an important task, but one that's not easy to accomplish, particularly because there's no standardization of these metrics among ISPs. Although it'll be some time before standardized metrics for IP network performance and

availability are agreed upon, check out the work of the IETF's working group on *Internet Provider Performance Metrics* (IPPM) to see the latest efforts.

Many of the service providers offering guaranteed service will often locate measurement devices at your CPE. For comparison's sake, you should try to locate your own measuring devices in parallel with those installed by your ISP. You also may find that, before long, ISPs offer direct connections between their management and monitoring environment and customer-management environments (see Figure 9.6), allowing customers direct access to the data that relates to their VPN.

FIGURE 9.5 Measurement areas for SLAs.

Just as there are no standards for Internet SLAs and performance metrics, there's no standard format for the reports that ISPs provide to show that they're complying with an SLA. Whenever possible, see that the reports are delivered regularly, perhaps every week, that they contain information on any gaps the ISP may have had in gathering the data, and that they include an explanation of the process for gathering the data. Their interpretation of the data is also welcome, especially if it includes any warnings about possible future degradation of performance before users are affected.

Previous | Table of Contents | Next

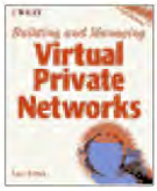

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

There are more than 4,500 Internet Service Providers in the United States. Obviously, attempting to describe their services and fee structures in any detail on an individual basis is beyond the scope of this book. Nor would it necessarily provide timely information for your benefit because mergers and changing technologies keep the market in near-constant flux. If you're shopping for an ISP, a good place to start looking for likely candidates is the *Directory of Internet Service Providers* published by Boardwatch Magazine at different times throughout the year, which provides a fairly complete listing; also check out the Boardwatch Web site at www.boardwatch.com and TheList at thelist.internet.com. Then apply the selection criteria we've listed in this chapter to help select an ISP that will meet your needs.

In-House or Outsourced VPNs?

As we mentioned in Chapter 8, "Designing Your VPN," part of the VPN design process is to decide how much of the implementation effort is to remain in-house and how many tasks are outsourced to the service provider (see Figure 9.7). Obviously, the more of the VPN operations that you outsource to your ISP, the stronger your need for an SLA.

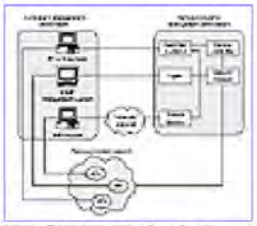

FIGURE 9.6 Integrating corporate and ISP network management.

There are a variety of reasons for outsourcing your VPN. For instance, you may not have enough staff to maintain the CPEs and manage security. Or, you may find that the cost of maintaining your own remote access servers and modem banks is prohibitive and you would be better served with a dial-in VPN.

As Figure 9.7 illustrates, there's no simple dividing line between an in-house VPN and an outsourced one. You'll find that ISPs are willing to offer a variety of services that can be tailored to your needs, ranging from straight connectivity to full installation, configuration, and management of your VPN.

Consider two key areas of operational requirements when dividing the responsibilities between the service provider and your company: administration and configuration management.

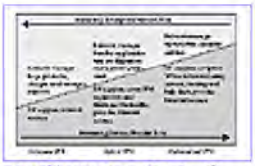

FIGURE 9.7 Outsource efforts versus in-house.

Administrative requirements include such items as account management, help desk, and troubleshooting assistance. Just as important is how you'll coordinate support between the provider and your enterprise. As corporations and service providers move into the shared network architecture that characterizes VPNs, it will be important to create an effective support chain that permits the two-way flow of information.

Another administrative factor to consider is managed access, or the management architecture that will support the public architecture of the Internet and mediate the interconnection between the provider's network and your private enterprise networks. Managed access has a role to play in determining both where monitoring tools and servers are located and how problems are resolved. Finally, there are various billing, usage tracking, and analysis factors to consider. For example, dial-up access servers should allow both enterprise accounts and/or user accounts to be debited for their access time. Network utilization reports should be accessible electronically and dynamically (preferably via the Web); allowing for custom-built reports is also a useful option.

Under configuration management, the key task is that of managing security. Under all circumstances, your corporation should hold-primary responsibility for the policy and configuration of security services.

Commercial VPN Providers

Even though there's a huge number of ISPs, only a few are prepared to handle an outsourced VPN. Since much of the rest of this book is about the details of designing and implementing your own VPN, this seems like the best place to describe some of the typical outsourced VPN services that are available. This list is only a sampling, although the VPN market is young, it is growing rapidly as both potential customers and service providers realize what they may gain from VPNs.

Note that many of the services are based on firewalls, often with product-specific client software for mobile users. Figure 9.8 sketches a typical VPN architecture, based on UUNET's offerings, which we'll describe later in the section.

Prices for these commercial services vary from $750 to $5,000 per month per site, depending on such variables as the speed of the connection, the hardware installed at the site, the desired network availability and latency, and, in some cases, the strength of security desired and the number of remote users.

ANS VPDN Services

Advanced Network Services (ANS) has been offering its *Virtual Private Data Network* (VPDN) service to corporations for a few years, using its own proprietary encryption and tunneling technologies. The VPDN services include SureRemote for remote dial-in access, InterManage for managed security, and InterLock for firewalls.

ANS' proprietary system uses 128-bit RC2 for encryption at domestic sites, but also offers 64-bit RC4 for either domestic or foreign sites. Tunneling is accomplished via a proprietary system based on UDP, and key management and exchange is performed via a proprietary extension to the *Open Shortest Path First* (OSPF) routing protocol. With changes in the market and technology, ANS has announced plans to support open standards as they're finalized, which includes IPSec and digital certificates, perhaps before the end of 1998.

Monitoring and management of the VPN is handled through the ANS Network Operations Center located in Ann Arbor, Michigan. Reports delivered on a weekly basis include bandwidth usage, security logs, and dial-in line availability.

ANS's Service Level Agreement commits to 99.5 percent availability for each site connected to the VPDN, with a network latency of 70 milliseconds or less. This availability covers the CPE at each customer location, the local loop, and the router configuration. In an effort to promote the flexibility of VPNs and accommodate new sites, ANS promises to turn on your VPDN within 32 business days; this process includes hardware procurement, on-site installation, and end-to-end network testing. This time period is measured from the day the contract is signed to the day the sites are reviewed.

FIGURE 9.8 Typical commercial VPN architecture.

AT&ampT WorldNet VPN

AT&ampT first started offering its WorldNet VPN services in late 1997. The network includes more than 300 POPs and 800/888 service for dial-in access, with international access (dial-in or ISDN) in 35 countries.

The first two bundles offered in WorldNet VPN are a firewall-based service, using Check Point Software's Firewall-1 running on Sun's Netra servers, and a dial-in service using Bay Networks' Instant Internet Access Server with a built-in firewall. Like many of the other services profiled in this section, AT&ampT can manage the firewalls for the customer or let the customer manage the firewalls themselves. AT&ampT also will arrange to manage any other optional network equipment, such as routers and CSU/DSUs, or deploy Cisco routers for its customers as needed, upon request.

Previous | Table of Contents | Next

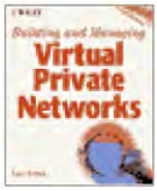

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

As this book went to press, AT&ampT had plans to expand its WorldNet VPN offerings to include IPSec tunneling with digital certificates as well as support for PPTP and L2TP.

WorldNet VPN uses RADIUS based on Novell NDS servers for authentication via CHAP for dial-in users. Proxy RADIUS servers weren't included in the initial rollout of WorldNet VPN, but they are planned for a later release, perhaps by the middle of 1998. Closed user groups also can be defined to restrict dial-in access for groups of users to particular sites. The firewalls are configured to protect internal IP addresses via NAT and include packet filters to prevent address spoofing.

The WorldNet VPN system is monitored and maintained on a full 24 × 7 basis. Remote users can access a hotline for support at any time of any day.

AT&ampT includes a Service Level Agreement as part of the service contract. The main points of the SLA include: if dedicated access connection is down for 10 minutes or more during any single day, AT&ampT will credit the customer for 5 percent of the monthly connection charge (up to a maximum of 25 percent in one month, with the annual maximum credit no more than one full month of service); network uptime is guaranteed to be at least 99.7 percent; and the network latency will be 150 milliseconds or better between any two AT&ampT-managed customer sites.

For customers who purchase optional router service, the end-to-end guarantee covers the total AT&ampT IP backbone, the access router, and the local access service that connects the customer's premises to AT&ampT. Without the router option, the guarantee covers the AT&ampT IP backbone from entry port to exit port.

CompuServe IP Link

CompuServe may be well-known for its international bulletin board services and dial-up Internet access for individuals, but its Network Services Division also offers VPN services for corporate clients, called IP Link and IP Link Plus.

IP Link and IP Link Plus are aimed primarily at dial-in VPNs and only differ from each other in the number of users they support; IP Link is limited to 100 users, while IP Link Plus covers businesses seeking to support more than 100 users.

IP Link uses a Cisco router at the customer's site and leverages CompuServe's extensive network of POPs to provide dial-up access for VPN users. Mobile users are required to have a PPP client on their computer because IP Link is based on L2TP.

CompuServe uses its own authentication system, the CompuServe Authentication Service, which can be

configured for event, challenge-response, or time-based authentication. Event and challenge-response systems are based on Secure Computing's SafeWord system, which also allows use of standard authentication protocols like TACACS+ and RADIUS. The time-based system uses Security Dynamics' ACE/Server as an authentication server and SecurID cards as password generators on the user's computer. VPN traffic is encrypted using the DES algorithm.

GTE Internetworking

The VPN services offered by GTE Internetworking also are based around managed firewall services. This service, called Site Patrol, can be set up with either the Gauntlet firewall from Network Associates Inc. or Firewall-1 from Check Point Software. To accommodate dial-up users, GTE Internetworking uses V-ONE's SmartGate product in conjunction with the site's firewall; SmartGate can be used for authentication, encryption, and authorization of remote users.

Geographic coverage includes 550 local dial-in numbers in the United States, with 240 locations scattered throughout the United States and 79 countries around the world.

Site Patrol's managed firewall services for the VPN include 24 × 7 security monitoring and assistance, and a predefined, three-stage escalation procedure for security breaches. Policies for dealing with security events are worked out as a collaborative effort between GTE security personnel and the customer's staff.

One of the unique features of Site Patrol is that it's not restricted to traffic carried only by GTE Internetworking. Even for cases in which connectivity is provided by an ISP other than GTE Internetworking, the Site Patrol service monitors security and pinpoints security breaches on the other networks.

Summary reports of usage, traffic, and security incidents are delivered to customers on a monthly basis. In addition, GTE Internetworking maintains an archive of historical data that the customer can access to review past performance or incidents to review and formulate new policies when needed.

InternetMCI VPN

MCI (now a part of Worldcom) offers a firewall-based VPN service called InternetMCI VPN, which also supports dial-in access via a firewall-specific client. MCI also can secure corporate data at each site with its managed firewall services, which includes installation, configuration, and monitoring of each site's firewalls.

MCI employs the Firewall-1 product from Check Point Software for both its managed firewall and VPN services. Firewall-1 offers users IPSec with either manual key exchange or Sun's SKIP automatic key exchange, as well as a proprietary encryption scheme called FWZ. More details on Firewall-1 will be presented in Chapter 10, "Firewalls and Routers."

The remote client, Firewall-1 SecuRemote, also is provided by MCI for mobile users wanting dial-in access to corporate sites protected by Firewall-1. User authentication can be based on one-time passwords using S/Key, the SecurID token cards from Security Dynamics, the user's operating system password, or a RADIUS account.

As part of its support for mobile users, MCI maintains a global directory of dial-in numbers that users

can access from any country that MCI services. InternetMCI VPN support is handled on a 24 × 7 basis, including coverage of security and global dial-in access problems.

InternetMCI's help desk support can be configured in one of two ways. If your company wants to maintain its own help desk for VPN support, then MCI will provide support directly to your network managers only. On the other hand, if you don't want to provide in-house support for VPN users, then InternetMCI will offer help desk support to all of your users.

Clients can set up and administer user accounts for their VPN via a Web-based interface.

UUNET ExtraLink

UUNET (also a part of Worldcom) offers its ExtraLink and ExtraLink Remote VPN services based on encrypting routers rather than firewalls. Data within corporate sites is protected by Check Point's Firewall-1, which is managed by UUNET as part of the service. ExtraLink Remote enables mobile users to dial into the VPN via UUNET's worldwide network of more than 845 POPs.

Each site must have a Cisco router installed, because traffic is encrypted by the 56-bit DES algorithm that ships with Cisco's IOS. Access lists also can be created on each router to block sites from communicating with each other, if necessary. The Firewall-1, also installed at each VPN site, is used to handle user authentication; remote clients use Check Point's SecuRemote client software to access the VPN.

Previous | Table of Contents | Next

UUNET also installs a PC-based network management system as part of each site's CPE. The system monitors performance by collecting packet throughput information and relating that information to a central UUNET management site. Reports on performance and availability are provided to customers on a monthly basis.

UUNET's Service Level Agreement promises 99.6 percent availability for VPNs consisting of 3 to 5 sites and 99.9 percent availability for VPNs that consist of 12 sites or more. Latency is guaranteed to be 150 milliseconds or better. If UUNET misses the guarantee in a month, the company promises to refund 25 percent of the total network charge.

Other VPN Providers

The number of VPN service providers is likely to grow as this book goes to press. Some of the current providers we haven't covered in detail include the following:

Concentric Network Corp., whose Enterprise VPN uses VPNet Technologies' VPLink products for hardware encryption at each LAN site

Netcom On-line Communication Services Inc., which uses Livingston's IRX Firewall Router and Milkyway's Black Hole firewall as part of its NETCOMplete for Business service

Pilot Network Services Inc., with its Secure Road Warrior service, which includes 128-bit key encryption

TCG CERFnet, whose Enterprise-Quality VPN offers IPSec tunnel mode and either 4-Mbps or 10-Mbps encryption devices

Future Trends in ISPs

Internet Service Providers in general are looking at value-added services like VPNs as a way to expand their business and get new corporate customers. This shift towards managed services over the Internet means that the ISPs have to deploy more intelligence at their network's edge (i.e., the interface between the customer's LAN and the ISP) and that they have to maintain closer relations with their customers, acting more as partners than simply a business providing a service.

Most of the VPN services offered by ISPs today should be considered first-generation VPNs. Look to the future for improved handling of different classes of traffic, either through classes of service or other QoS approaches such as *ReSource reserVation Protocol* (RSVP) over IP or ATM's built-in QoS classes. As devices become available for automatic mapping of RSVP classes to ATM classes, expect ISPs to deploy them as part of their effort to provide QoS services.

Also look for ISPs to come together over the next few years and reach agreement on how to measure network availability and latency so that SLAs can be extended to traffic that spans multiple ISPs. Both this issue and that of QoS support depends on the ISPs cooperating on various issues, relating to both technologies and policies. It'll happen, but not overnight. Both the work of the IPPM working group to create standard metrics for network performance and efforts by the *Automotive Industry Action Group* (AIAG) to qualify ISPs for what may be the world's largest extranet are leading the charge forward.

Summary

The ISP is an important cog in the design and success of your VPN. The design of your VPN will determine the involvement of your ISP and may limit the ISP to simply providing the pipe to the Internet or might utilize the ISP as a full-fledged designer and maintainer of an outsourced VPN.

When evaluating an ISP for your VPN, you should consider many details, but they generally fall into the following categories: ISP infrastructure, network performance and management, connectivity options, and security.

As you plan the relationship between your company and an ISP for the construction of your VPN, look at using a Service Level Agreement (SLA) to set expectations for the network's performance and how the ISP will handle troubleshooting and network repairs, among other issues of mutual concern. If you do use a SLA, keep in mind that you should track your ISP's performance in parallel with the provider's own measurement systems in order to ensure that the terms of the SLA are being met.

If you choose to outsource your VPN to a service provider, a growing number of companies are capable of doing the job, including ANS, AT&ampT WorldNet, internetMCI, GTE nd UUNET, among others.

Previous | Table of Contents | Next

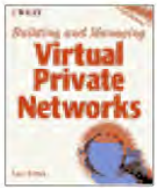

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 10
Firewalls and Routers

After you have a connection to the Internet, the important network devices for your VPN are the ones that control the access to your protected LAN from an external, Internet-based source—the security gateways as we've called them in past chapters. The external source might be another of your corporate LANs tunneling to your site, a mobile worker with a laptop using an ISP-created tunnel, or a business partner tunneling through the Internet to your LAN. Ideally, VPN devices should be able to handle all of these situations equally well; many do, but not all are equally adept at handling the different connectivity situations.

Just as the market definition of VPNs has been fairly confusing, so too has been the classification of the hardware and software required for creating Internet VPNs. Each vendor has his own idea of what a VPN is and how his products fit into the scheme of things (and some of them are right!).

Since each vendor has his own idea of what a VPN device is, classifying VPN hardware and software can be somewhat problematic. As this market begins to mature, we're seeing not only new classes of products, but also a move towards integrating many of what had been individual VPN devices into a single product. This integration inevitably leads to the time-honored argument of whether buying *best-of-breed* individual products or an integrated solution is a better course of action. The information in this and following chapters should help you determine whether modular or integrated solutions will best meet your needs.

VPN hardware and software can be placed at various locations in the network. Consider for the moment how the corporate site of your VPN would be connected to the Internet via an ISP (see Figure 10.1). Starting at the link from the ISP's POP, you would have a CSU/DSU followed by a router, a firewall, and then the corporate LAN. VPN devices can be placed at various locations along this path from ISP to corporate LAN.

Recall that a full-fledged VPN depends on encryption, authentication, and tunneling services. Devices that add these services can be inserted between the CSU/DSU and the router or between the router and the firewall. Other products offer VPN services as part of either the firewall or the router. Some products integrate all of the network services between the ISP and your LAN, bundling WAN links, routing, firewalls, and VPN services into a single device. Lastly, some of the Network Operating Systems (NOS) such as NT Server and NetWare are integrating VPN support into their software.

To start with, we'll concentrate on using firewalls or routers to create VPNs in this chapter, then go on to

dedicated VPN hardware, including stand-alone encryptors and integrated devices, in the next chapter. Chapter 12, "VPN Software," will focus on how NOSs and other software have evolved to support VPNs.

A Brief Primer on Firewalls

Firewalls have long been used to protect LANs from other parts of an IP internetwork by controlling access to resources on the basis of packet type, application type, and IP address. Deployment of firewalls has increased tremendously since the Internet has become more commercialized and as businesses seek to attach their networks to the Internet. If your corporate network is connected to the Internet, you probably already have at least one firewall to control traffic from the Internet.

FIGURE 10.1 Locations for VPN functions.

Firewalls and Security Policies

A firewall is an integral part of your organization's security policy, because it determines what traffic passes between your internal networks and the Internet. (Firewalls also may be used to protect sensitive or restricted subnets from the rest of your corporate network.) In addition to your firewalls, the corporate security policy should include password policies for sensitive systems, data encryption, data backup, and user account management.

Before we discuss how firewalls can be used to support VPNs, let's spend a few pages reviewing some of the salient features of firewalls. If you're already familiar with firewalls, you might choose to skip the rest of this section and go right to the section on firewalls and VPNs.

Types of Firewalls

There are three main classes of firewalls: packet filters, application and circuit gateways (*proxies*), and stateful inspection (or smart filter) firewalls.

PACKET FILTERS

Packet filtering firewalls were the first generation of firewalls. Packet filters track the source and destination address of IP packets permitting packets to pass through the firewall based on rules that the network manager has set (see Figure 10.2).

Two advantages of packet filter firewalls are that they are fairly easy to implement and they're transparent to the end users, unlike some of the other firewall methods we'll discuss shortly. However, even though packet filters can be easy to implement, they can prove difficult to configure properly, particularly if a large number of rules have to be generated to handle a wide variety of application traffic and users.

Packet filtering often doesn't require a separate firewall because it's often included in most TCP/IP routers at no extra charge. Of course, if you're planning to use packet filtering in a router as part of your

security policy, you should ensure that the router itself is secure.

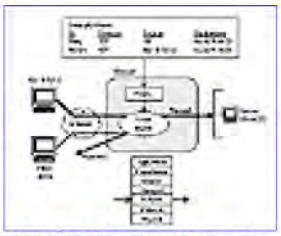

FIGURE 10.2 Packet filtering.

But, packet filtering is not the best firewall security you can get. One of its deficiencies is that filters are based on IP addresses, not authenticated user identification. Packet filtering also provides little defense against man-in-the-middle attacks (see Chapter 4, "Security: Threats and Solutions,") and no defense against forged IP addresses. Also, packet filtering depends on IP port numbers, which isn't always a reliable indicator of the application in use; protocols like *Network File System* (NFS) use varying port numbers, making it difficult to create static filtering rules to handle their traffic.

Packet filters can be used as a part of your VPN, because they can limit the traffic that passes through a tunnel to another network, based on the protocol and direction of traffic. For example, you could configure a packet filter firewall to disallow FTP traffic between two networks while allowing HTTP and SMTP traffic between the two, further refining the granularity of your control on protected traffic between sites.

APPLICATION AND CIRCUIT PROXIES

Since they're based on address information, packet filters look exclusively at some of the lower layers of the OSI model. Better, more secure firewalls can be designed if they examine all layers of the OSI model simultaneously. This principle led to the creation of the second generation of firewalls: application and circuit proxies. These firewalls enable users to utilize a proxy to communicate with secure systems, hiding valuable data and servers from potential attackers.

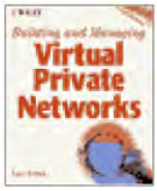

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

The proxy accepts a connection from the other side and, if the connection is permitted, makes a second connection to the destination host on the other side. The client attempting the connection is never directly connected to the destination. Because proxies can act on different types of traffic or packets from different applications, a proxy firewall (or *proxy server*, as it's often called) is usually designed to use proxy agents, in which an agent is programmed to handle one specific type of transfer, say FTP traffic or TCP traffic. The more types of traffic that you want to pass through the proxy, the more proxy agents need to be loaded and running on the machine.

Circuit proxies focus on the TCP/IP layers, using the network IP connection as a proxy (see Figure 10.3). Circuit proxies are more secure than packet filters because computers on the external network never gain information about internal network IP addresses or ports. A circuit proxy is typically installed between your network router and the Internet, communicating with the Internet on behalf of your network. Real network addresses can be hidden because only the address of the proxy is transmitted on the Internet.

Circuit proxies do not examine application data; application proxies, which we'll get to next, do that. When a circuit proxy establishes a circuit between a user and the destination, the proxy doesn't inspect the traffic going through the circuit, which can make the proxy more efficient than an application proxy, but may compromise security.

On the other hand, circuit proxies are slower than packet filters because they must reconstruct the IP header to each packet to its correct destination. Also, circuit proxies are not transparent to the end user, because they require modified client software.

As we mentioned earlier, application proxies examine the actual application data being transmitted in an IP packet (see Figure 10.4). This approach thwarts any attackers who spoof IP packets to gain unauthorized access to the protected network. Because application proxies function at the Application layer of the OSI model, they also can be used to validate other security keys, including user passwords and service requests.

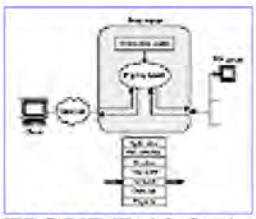

FIGURE 10.3 A circuit proxy.

Proxy firewalls often require two copies of an agent running for each service: one copy to communicate with the internal hosts and one to communicate with the external hosts. Thus, an application proxy may

have two copies each of FTP, HTTP, and telnet agents. A circuit proxy operates in a similar fashion; it may have one copy of TCP for the internal network and one copy for the external network.

Because application proxies operate as one-to-one proxies for a specific application, you have to install a proxy agent for every IP service (HTTP/HTML, FTP, SMTP, and so on) to which you want to control access. This leads to two of the disadvantages of application proxies: a lag usually exists between the introduction of new IP services and the availability of appropriate proxy agents; and the application proxy requires more processing of the packets, leading to lower performance. Furthermore, many of the application proxies require modified client software, although the firewall's operation is becoming transparent to end users in many of the newer application proxy firewalls.

One important differentiating feature of application proxies is their capability to identify users and applications. This identification can enable more secure user authentication, because digital certificates or other secure token-based methods can be used for identifying and authenticating users.

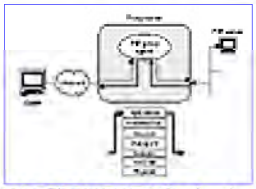

FIGURE 10.4 An application proxy

SOCKS Since circuit-level proxies can offer adequate security for many networks and because some users don't want to pay the price of lower performance found in application-level proxies, a standard for circuit proxies, called *SOCKS*, has been developed. The SOCKS proxy is designed to pass only SOCKS-related traffic, so SOCKS client software has to process all traffic being passed to the proxy for the traffic to be recognized. No other proxies are included in a SOCKS proxy firewall.

SOCKS was designed for TCP-based client/server applications, using a proxy data channel to communicate between the client and the server. In a SOCKS environment, an application client makes a request to SOCKS to communicate with the application server. This request includes the application server address and the user's ID. SOCKS then establishes a proxy circuit to the application server and relays information between client and server. With SOCKS version 5, authentication and support for UDP relay have been added. SOCKS is commonly implemented as a circuit-level proxy that has enhanced features, such as auditing and alarm notifications, so that it offers many of the features expected of a firewall.

The downside to SOCKS is that client applications must be specially coded for SOCKS or application-level proxies. Aventail, one of the main vendors of a SOCKS-capable firewall, tries to address this problem by including a DLL with its Windows client.

STATEFUL INSPECTION

The optimal firewall is one that provides the best security with the fastest performance. A technique called *Stateful Multi-Layer Inspection* (SMLI) was invented to make security tighter while making it easier and less expensive to use, without slowing down performance. SMLI is the foundation of a new generation of firewall products that can be applied across different kinds of protocol boundaries, with an abundance of easy-to-use features and advanced functions.

SMLI is similar to an application proxy in the sense that all levels of the OSI model are examined. Instead of using a proxy, which reads and processes each packet through some data manipulation logic, SMLI uses traffic-screening algorithms optimized for high-throughput data parsing. With SMLI, each packet is examined and compared against known states (i.e., bit patterns) of friendly packets (see Figure 10.5).

One of the advantages to SMLI is that the firewall closes all TCP ports and then dynamically opens ports when connections require them. This feature allows management of services that use port numbers greater than 1,023, such as PPTP, which can require added configuration changes in other types of firewalls. Stateful inspection firewalls also provide features such as TCP sequence-number randomization and UDP filtering.

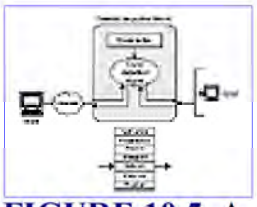

FIGURE 10.5 A stateful inspection firewall.

Firewalls and Port Numbers

Each TCP/IP application is assigned a unique port number used to establish a connection. For a client/server pair, both the client and the server have unique port numbers. Almost all TCP/IP client applications use a randomly assigned port number greater than 1,023 for their end of a connection. If a client/server pair is going to communicate over a firewall, then the firewall has to be configured to open port numbers higher than 1,023, or the client will be unable to establish a connection. But this can cause configuration problems, since some services such as NFS, NIS, and Netware/IP also use ports greater than 1023. If these ports were already opened at the firewall to enable communications between client/server applications, an attacker could disrupt the other services depending on ports greater than 1023.

Previous | Table of Contents | Next

Stateful inspection firewalls are highly secure, which explains why they're being used in more and more VPN bundles. However, these firewalls have to be supplemented with proxies in order to support other important functions, such as authentication.

General Points

It's not possible to say that any one firewall type is always better than another. That's why firewall vendors these days are starting to blend approaches—mixing stateful inspection and proxies, for instance. When deciding on which firewall to select, try to determine what level of protection you need for your traffic based on what part of the packet a firewall processes (see Figure 10.6); also, keep in mind how your firewall is likely to interact with your VPN protocols, which we'll cover in more detail in the next section.

Many of the ISPs offering VPN services (see Chapter 9, "The ISP Connection") either include managed firewalls as part of their VPN offerings or as separate services, giving you the option to outsource your firewall management and monitoring. But, if you're going to manage your own firewalls, one of the best places to gets security updates and advisories is the CERT Coordination Center at www.cert.org, located at Carnegie Mellon University's Software Engineering Institute. Another organization to check is the *International Computer Security Association* (ICSA) at www.icsa.net, which certifies firewall products and can audit your site security.

FIGURE 10.6 What firewalls inspect.

Firewalls and VPNs

Although firewalls should be considered part of your corporate security solution, they are not sufficient on their own for creating a VPN. That's because a firewall cannot monitor or prevent changes to data that may occur as a packet crosses the Internet (data integrity), nor does a generic firewall include encryption.

Furthermore, even if you installed host-based encryption on all your computers (using IPSec, for instance), you'd still need firewalls in your organization. Internet firewalls enforce an enterprise network security policy and are part of a perimeter defense. IPSec on every desktop provides for privacy and authentication but does not ensure that the corporate network security policy is enforced (what services are allowed, when to force virus scanning, etc.). Firewalls are able to enforce a policy that requires private links between networks even if desktop users cannot or do not use an encrypted connection.

Firewalls are often considered to be logical VPN termination points because you can manage your entire network security policy through that single point. However, firewalls are complex devices to install and manage because of the possibility of conflict between rules if you are not careful in establishing or modifying the rule base. Additionally, having firewalls perform VPN services increases your risk in case the firewall fails or is compromised. If you lose your firewall, your VPN goes with it.

Firewalls between busy networks, as on a WAN connection, carry a heavy load just processing the traffic passing through them. Adding encryption and key management may significantly hurt performance, especially when several VPNs are running. Some firewalls, such as Cisco's PIX, move data encryption off the processor and onto a card within the box to improve performance. Check Point Software is making similar arrangements, planning to offer a bundle of the Chrysalis accelerator board (for faster encryption) with its Firewall-1 software later in 1998. Other companies also are bundling firewalls into their hardware. For example, Timestep has ported Check Point's Firewall-1 software to the operating system that's part of their PERMIT security gateway.

IPSec traffic can be handled in two different ways, either as unfiltered packets or as filtered packets (see Figure 10.7). In the unfiltered approach, the IPSec traffic is handled the same way it is in a router—that is, the IPSec-protected data is transferred directly to the internal network without any filtering or controls on its contents. In the filtered approach, the firewall's filter and proxy controls are applied to the IPSec traffic before it is allowed into the internal network. Filtering IPSec traffic can be particularly useful if your security policy is to pass only certain types of traffic between VPN sites, say e-mail and FTP. Filtering also can be useful for controlling the traffic exchanged with business partners if you expand your VPN to an extranet.

It's always a challenge to maintain consistent security policies across different sites. It's particularly important to maintain consistency for the firewalls at all sites, because they control the access to and from each site. Configuration and access rules should be the same for every firewall in the entire enterprise. But, many firewalls require that such consistency be achieved by hand, with administrators at each site carefully updating their own copy of the global rules. Fortunately, some firewalls do maintain their security rules and configuration in a set of files that can be copied from one firewall to another, making your management job easier. But, remember that, if the firewall's configuration files can be transported and installed in the other firewalls, then your next problem is to deliver these files safely. This might have to be done by face-to-face exchanges, a bonded courier, or even secure e-mail, but you should do your best to ensure the security of the exchange and the handling of those files.

Firewalls and Remote Access

Because many firewalls already have strong user authentication mechanisms, they can offer additional functionality by serving as the focal point for dial-in users. Quite a few firewall products can verify the user's identity and establish an encrypted session between the firewall and the dial-in user's computer to protect confidentiality, forming the basis of a dial-in VPN.

In order to make this work, your remote users will have to install appropriate client software on their computers. If you're planning a firewall for terminating PPTP or L2TP, then all your users need is a PPP client for dialing into their ISP. Recall, though, that a PPP client doesn't provide the strongest security for any data transferred between the client and your corporate site; if you want stronger encryption, you'll need to install either a PPTP or, preferably, a L2TP client on the remote user's computer.

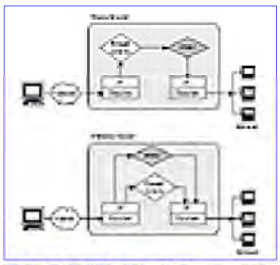

FIGURE 10.7 IPSec in firewalls.

As we've mentioned before, the current market trends emphasize PPTP and L2TP for dial-in VPNs and IPSec for LAN-to-LAN VPNs. But, IPSec is just as well-suited for dial-in VPNs. A number of firewall vendors offer remote client software that uses IPSec. Because IPSec has no standardized way for authenticating users, IPSec remote clients are usually vendor-specific. This means that you'll have to match the IPSec remote client software to your vendor's firewall and that, at least for the near future, you should stick with a single vendor to avoid interoperability problems. This could become particularly important if your mobile workers routinely access more than one VPN site. In such cases, each site should have the same firewall installed, with the same VPN options. Of course, having the same firewall at each site can simplify security policy management as well.

Previous | Table of Contents | Next

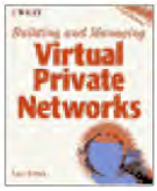

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Product Requirements

If you're going to use a firewall as a security gateway for your VPN, let's review the main issues that you need to consider. We'll review the common requirements first, then discuss IPSec-related requirements, finally touching on issues surrounding PPTP and L2TP.

COMMON REQUIREMENTS

Regardless of which protocol is used for your VPN, you need to consider how the firewall integrates with the rest of your security and network management systems. For example, many of the PPTP and L2TP systems depend on RADIUS or token-based systems for user authentication; if you're already using a particular system for authenticating remote users, then you can simplify the transition by installing a firewall that's compatible with your current system. Alternatively, you may feel the need to improve the security of your authentication systems by installing a two-factor system like SecurID (see Chapter 4, "Security: Threats and Solutions") as you roll out your VPN services. *A firewall's authentication method may become less of a distinguishing factor as more firewall vendors have been bundling stronger authentication systems with their products*, Whatever the reason, system compatibility is important. The same suggestions hold for IPSec implementations as well, even though IPSec has not standardized on a particular authentication method, and most solutions are vendor-specific.

If you're planning to use an authentication system based on digital certificates, then you should give some thought to how the certificates will be distributed and verified. We'll get into this in more detail in the chapter on security management (Chapter 13), but some of the factors to consider include whether a certificate authority should be maintained in-house or outsourced and how will certificates be linked to other services (through a directory service, for example).

Looking ahead, the need to integrate management of the increasing number of user privileges, such as bandwidth, QoS, remote access, and access to servers and other network resources, will drive further use of policy-based management. Policy-based management depends on having a distributed system for identifying and authenticating users that can be managed easily. Directories are thus fast becoming a cornerstone technology for policy-based management, particularly X.500 directories and *Lightweight Directory Access Protocol* (LDAP). It's worth keeping an eye on how the firewall vendors are tying their products, particularly authentication services, to LDAP and X.500 directories, as a first step towards deployment of policy-based management, even if you don't think you will be using such systems for a few more years.

Last, but perhaps just as important, remember that you're most likely going to be installing firewalls at more than one site. You'll be able to maintain a more consistent security policy if the firewall product

you pick supports synchronized administration of multiple sites. This administrator might involve file exchanges, as we discussed earlier, or some other form of remote management. If remote management capabilities are included in a product, be sure that remote access to the firewall is secure.

IPSEC

Since much of IPSec revolves around the use of cryptographic functions, either for encryption or packet authentication, it is important to ensure that a firewall not only supports the proper algorithms, but also the ancillary processes, such as rekeying and security associations. Also, since the IPSec standards are currently undergoing revisions to version 2, you should carefully investigate each product's compatibility with the version 2 specifications, which provide added flexibility and security.

The most secure devices support separate network connections for unencrypted and encrypted traffic, which enables you to provide connectivity for both kinds of traffic (it's highly unlikely that all of your traffic needs to be encrypted) and to maintain a physical separation between your encrypted networks and more open networks. Secure devices should reject all packets without a proper header (an IPSec header, for example) that arrive on the encrypted side, except perhaps for key-exchange protocols.

Check which cryptographic algorithms the VPN firewall supports. For those uses considered to be medium risk, the default DES CBC algorithm for encryption and HMAC-MD5 or HMAC-SHA-1 hash algorithms for authentication will suffice; if your traffic is higher risk, then be sure that automatic rekeying is supported.

On a related matter, even though manual keying is the minimum required by the IPSec specs, you should look for products that allow the cryptographic keys to be changed automatically and periodically or whenever a new connection is established. If interoperability will become a concern, don't settle for a proprietary key-exchange system, but insist on the systems described in IKE (see Chapter 5, "Using IPSec to Build a VPN"). Automatic rekeying further strengthens the security of your traffic by increasing the difficulty of cracking a key when it's intercepted (i.e., the key expires before it can be cracked).

Also be sure which IPSec headers the firewall employs. Although the original IPSec standards did not require the support of both AH and ESP headers, it's preferable to apply both headers to each packet. Some IPSec-compliant devices support only the Authentication Header.

Because security associations are crucial to the operation of IPSec, you should be able to manually input security associations, usually from a file similar to that recommended by S/WAN (see Chapter 5), and, if possible, specify wild card security associations to simplify configuring proxies.

The firewall always has to be treated as a secure device, which includes protecting the secret and private keys used by the device. Personnel with incidental access to the firewall should not be able to obtain the keys, for example.

In the original IPSec specifications, there was no system for countering replay attacks in which an attacker intercepts a series of packets and then retransmits, or replays, them at a later time. However, the revised IPSec standards that were passing through the IETF standards process in late 1998 included an antireplay service in the new Authentication Header; this service can be invoked at the discretion of the receiver to help counter denial-of-service attacks that would be based on these retransmissions. To provide added security to your VPN, see whether the vendor's products supports the new IPSec

antireplay system rather than using a nonstandard variant.

As we've said, every security device should have a way of logging security events (*incidents*) and reporting them. If possible, be sure that the system can generate some kind of alarm if some persistent activity takes place, as this may indicate a systematic attempt at breaching the site's security.

Although you may feel that the IPSec transport mode is sufficient to protect your data, there will inevitably come a time when you'll need the stronger protection afforded by tunnel mode. A proper IPSec-compliant firewall should support both modes.

Previous | Table of Contents | Next

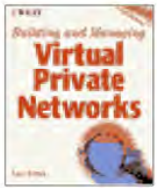

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

PPTP AND L2TP

Because PPTP and L2TP expect tunnels to be terminated at a network server, firewalls aren't usually used as the termination points for these tunnels. Instead, any firewall you install on your network should be configured to pass the traffic from PPTP or L2TP using the properly assigned port number. PPTP traffic uses TCP port 1723, which cannot be changed, and the default port for L2TP is 1701. L2TP does not demand that one specific port number be assigned for the firewall to pass L2TP traffic. Network managers have the option of selecting a different firewall port number for passing L2TP traffic, making it more difficult for attackers to take over L2TP tunnels or try other attacks based on a known port number.

Microsoft's implementation of its PPTP server for Windows NT (the *Routing and Remote Access Server*, or RRAS) does permit packet filtering to be enabled as part of the configuration. If packet filtering is enabled on a server running RRAS, then only PPTP packets will be passed.

AN OVERVIEW OF THE PRODUCTS

Because firewalls often seem to be the logical location for terminating VPN tunnels and enforcing security policies, there are currently more VPN-compatible firewalls than any other class of VPN device. As you'll see in Table 10.1, some of these firewalls are software products designed to run on a variety of operating systems, such as Unix and Windows NT, and a few are hardware devices that can be configured via most of the popular workstations.

Does the operating system for a firewall make a difference? As far as security is concerned, it's becoming less of an issue. The most secure firewalls have traditionally been those written with their own operating systems, and various security holes were found in the more common workstation and server OSs, like Unix and Windows NT. But, vendors offering firewall software that runs on the more common OSs now usually include added code that patches any of the known security holes in the OS. If performance is an issue, Unix might be preferred to Windows NT, but that's an issue we'll leave for others to address in any detail, because benchmarks can change almost every day.

When you're reviewing this list of products to help you pick likely candidates for your own VPN, bear in mind that the market doesn't stand still, and newer versions of some of these products will either include modified features or new features that were introduced since this book was put together. In other words, use the table as a guide, not the "last word."

Note that many vendors of routers and other VPN hardware have started to bundle firewalls with their products; the two most notable bundled firewalls are Check Point Software's Firewall-1 and Network Associates' Gauntlet.

Despite the relatively large number of firewall-based products for VPNs, a major continuing concern is the performance of these products. Even to perform their normal operations, firewalls—especially those that investigate packet contents (application proxies, for instance)—have a lot of computations to perform to get their job done. Add to that the overhead associated with encryption, and many systems will be severely taxed. For example, benchmarks run by Network World indicate that firewall-based VPNs often delay traffic anywhere from 1.3 to 3.3 times that expected without the VPN processing.

Using firewalls to create VPNs is a workable solution for some networks. Firewall-based VPNs are probably best suited to small networks that transfer small amounts of data (on the order of 1-2 Mbytes/sec over a WAN link) and remain relatively static (i.e., don't require frequent reconfiguration). If you're looking for higher performance, there are other, better solutions.

Some companies, like Check Point Software, are now pushing the idea that traffic control, including such tasks as bandwidth management and QoS, should be part of the definition of a VPN. That's a step beyond the definition of VPN that we adopted in this book, but it's a reasonable next step that will most likely gain more support in the future. After all, if you want to treat the tunnels transporting your data over the Internet as your network—whether they're virtual or not—shouldn't you have the same control and policies as you have on your own physical network?

Integrating traffic control with authentication and access control also makes sense over the long run, as policy-based network management becomes more prevalent (and useful). We're only beginning to see the first steps in integrating various management functions and implementations of policy-based management for enterprise networks, and it'll probably be a few more years before we start to see widespread deployment.

Previous | Table of Contents | Next

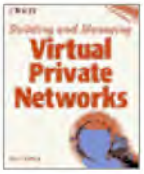

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Routers

If firewalls seem like a logical place for installing VPN functions, then routers are even more so. After all, routers have to examine and process every packet that leaves the LAN, so why not let them handle the encrypting as well?

We've already mentioned one way that routers can be used to protect your LAN from outside attackers, and that's with packet filtering. (Just remember that relying on a router's packet filtering for part or all of the firewall protection means that the router itself must be secure.) But, packet filtering isn't sufficient to secure against many kinds of attacks on your network, which is one of the reasons why other types of firewalls have been developed. Within the context of VPNs, the type of routers in which we're most interested are encrypting routers.

Product Requirements

Many of the requirements for an encrypting router are the same as the ones for firewalls that we presented earlier in this chapter. After all, they're performing the same functions, it's just that the auxiliary network functions differ (routing versus security perimeter defense). Because the requirements are similar, we just briefly list them here.

Encrypting routers are appropriate for VPNs if they do the following:

- Include separate network connections for encrypted and unencrypted traffic
- Support at least the default IPSec cryptographic algorithms (DES CBC, HMAC-MD5, and HMAC-SHA-1)
- Support a cryptographic key length that best matches your security needs
- Allow manual security association configuration
- Restrict access by operations personnel to keys
- Support automatic rekeying at regular periods or for each new connection
- Support the antireplay mechanism of IPSec version 2
- Log failures when processing headers and issue alerts for repeated disallowed activities
- Support both transport mode and tunnel mode IPSec

You've probably noticed that there's a distinct IPSec slant to the preceding list—that's in keeping with the opinion, mentioned previously in this book, that IPSec offers the most secure VPN systems possible. Some routers support either PPTP or L2TP for tunneling, in which case the requirements for interoperability are fewer. Because L2TP defers to IPSec for encryption, L2TP-capable routers should be judged using the same requirements as for IPSec-capable routers.

Because routers are generally designed to investigate packets at Layer3 of the OSI model and not authenticate users, you'll most likely have to add an authentication server in addition to your encrypting router to create a secure VPN. If you're not already using an authentication system, say for remote access, review some of the systems we described in Chapter 4. PAP or CHAP, such as provided in a typical PPTP installation or in Intel's ExpressRouter, is weak authentication. The strongest authentication is a two-factor system, such as that provided

by SecurID or CryptoCard. Many of these systems are designed to work with encrypting routers, but you'll have to check with the vendors to be sure of their compatibility.

TABLE 10.2 VPN-Capable Routers

Product (Company)	IntelExpress Router VPN (Intel)	IOS 11.3 (Cisco)	MicroRouter Series, RISC Router 3500, 3800 (Compatible Systems)	2210 Nways Multiprotocol Routers (IBM)	Pipeline 220 w/ SecureConnect option (Asdend)	NetBuilder II Routers (3 Com)	VPN 500 Series (Bay Nerworks)
Price	$1,299–$5,999	$500–$7,000 acc. to router (not including roter)	Microrouters $1,895–$2,695 RISC Routers: $3,995-$4,495	$2,800–$3,800	$6,495	$10,000+	500n:$3,995 500n:$4,995
Tunneling protocol	proprietary	L2F	IPSec, GRE	L2TP, IPSec	PPTP, L2TP, L2F	PPTP, L2TP	IPSec
Protocols	TCP, UDP, IPX	TCP, UDP, IPX	TCP, UDP, IPX	TCP, UDP, IPX	TCP, UDP, IPX	TCP, UDP, IPX	TCP, UDP
Encryption type	144-bit Blowfish	DES	RSA, DES	IPSec	IPSec,	IPSec, MPPE	DES, 3DES
User authentication type	PAP, CHAP,	RADIUS, TACACS+	CHAP, PAP, RADIUS	—	PAP, CPAP, RADIUS, h/w token cards	PAP, CPAP, RADIUS, ACE, NT domains	ACE, RADIUS, CHAP
Integrated firewall	Yes	Yes	Yes	—	Optional	Yes	No
Compression	Yes	—	Yes	—	No	Yes	Yes
NAT	Yes	—	—	Yes	Yes	Yes	Yes

Previous | Table of Contents | Next

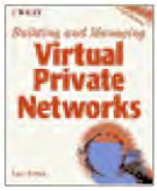

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

AN OVERVIEW OF THE PRODUCTS

If you compare the two product tables, Tables 10.1 and 10.2, in this chapter, you'll see that the number of encrypting routers is much smaller than that for VPN-capable firewalls. One reason for this disparity is that many of the integrated hardware devices for VPNs, which we'll be covering in the next chapter, often incorporate a router into the box.

Although most of the products listed in Table 10.2 support one or more of the standards we've discussed in this book (i.e., PPTP, L2F, L2TP, and IPSec), note that the Intel router uses a proprietary scheme for tunneling. It also uses the Blowfish algorithm for encryption, which, although a good algorithm, isn't included in any of the other standards. Because the Intel router is proprietary, it cannot be used with any other routers to create a VPN; all of your sites would have to have Intel Express routers installed to form a VPN.

As we've said before, the feature sets of these products should not be taken as the last word, because companies are always adding and changing features. As an example, although the IBM routers initially supported L2TP as a tunneling protocol, the company has publicly announced that these routers will also support IPSec in the near future.

Just as with firewalls, routers can take a performance hit when having to perform the added functions of a VPN, particularly encrypting packets. Cisco's *Encryption Service Adapted*, (ESA) is one way of dealing with such performance hits; the ESA is a coprocessor-based encryption engine that relieves the router's regular processor of this task.

Routers are being expected to perform more tasks in current-day networks. Not only might they be expected to handle VPN tunnels and encryption, but they have to handle quality-of-service provisioning as well, for example. These new tasks place new loads on the router, so you should choose carefully when deciding which functions need to be installed on your routers; otherwise, new functionality, such as QoS-based routing and VPNs, may well reduce the performance of your networks rather than enhance them. Routers continue to be logical devices for managing these many different tasks, so you'll need to balance a router's computational power, and that of add-on hardware like Cisco's ESA, against the new tasks you want the router to perform.

Because routers cannot handle all the functions of a VPN, such as user authentication, it's becoming increasingly common for vendors to bundle firewall software with their routers. For example, Bay Networks has embedded the INSPECT engine from Check Point's Firewall-1 into version 11.02 of the Bay Router Services OS so that a router and stateful inspection firewall can be shipped in the same box.

Summary

The three main types of firewalls—packet filters, proxies, and stateful inspection systems—each differ in the security they provide, their performance, and the difficulty of configuration. In general, the more secure a firewall is, the slower it is. And, because any one type of firewall does not cover all the bases for securing your LAN, you'll often find that firewall vendors are integrating the different methods into a single product for maximum security.

Both firewalls and routers can be used as key components for creating a firewall. A wider variety of VPN-capable firewalls than routers is available currently, and this probably will continue, because the idea of adding one security function (i.e., VPN tunneling and encryption) to another (i.e., perimeter security) makes sense to many buyers. But, beware that there may be high performance penalties for combining these two functions, and the available throughput may be inadequate for your higher speed links. Adding encryption coprocessor cards to the firewall/VPN computer can help solve the performance problem.

A small number of stand-alone encrypting routers are available for creating VPNs. They usually have to be supplemented with other products, such as authentication servers, in order to make a complete VPN. Furthermore, as we'll see in the next chapter, many integrated hardware devices, which include routing functionality, are now available for VPNs.

Previous | Table of Contents | Next

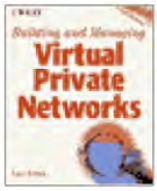

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 11
VPN Hardware

The preceding chapter covered many of the products that focus on using either a firewall or a router as the main building block for creating a VPN. And, as we've already mentioned, each of these classes of products suffers from certain shortcomings. For instance, although many of the products are suitable for small networks with low- to medium-sized bandwidth requirements, very few of the products mentioned in Chapter 10 can handle Ethernet speeds or T3 (44.736 Mbps) WAN links. Furthermore, many products have to be bundled with other companies' products in order to provide all the functions needed for VPNs, particularly authentication as well as encryption and tunneling.

If you'll recall our discussion of classifying VPN products in the beginning of Chapter 10, we mentioned that a number of locations in a network are suitable for performing the basic functions of a VPN, particularly if you have to use different products for different functions. One of the fastest-growing market segments seeking to provide VPN solutions consists of vendors offering integrated VPN hardware, in which a single box includes all of the required VPN functionality, replacing the need for adding software and hardware to an existing firewall or router and, in some cases, any hardware for the WAN link (see Figure 11.1).

One of the purposes of these VPN products is to offload the VPN functions from a firewall or router that doesn't have the computational horsepower to handle functions like encryption. Many of the systems mentioned in this chapter utilize custom-designed ASICs and, in some cases, special cryptographic chips to give them as much of a performance edge as possible.

Not all of the products mentioned in this chapter offer the same features. Some products are aimed at providing a turnkey solution for security, including a firewall. Other VPN hardware ranges from boxes that focus on encryption to turnkey systems that handle all aspects of an Internet connection, including WAN connections, routing, VPNs, DNS, and e-mail services, among others.

Types of VPN Hardware

One of the main differences among the products is their focus on the device that initiates a tunnel. Either a security gateway can create a tunnel to connect the LAN it serves to another gateway, or a single remote host can create a tunnel to connect to a gateway and the LAN it serves; in the past, we've referred to these as LAN-to-LAN VPNs and dial-in VPNs, respectively. We'll continue to use the term VPN gateway to describe products that can handle LAN-to-LAN VPNs and, in most cases, remote access by

individuals. But, to emphasize that some gateways are designed specifically to handle dial-in tunnels, we'll refer to those products as *remote VPN gateways*. We've shied away from using dial-in as a modifier here because the user never dials into the gateway directly, as he would for a remote access server; the only dial-in connection is between the user and his ISP.

FIGURE 11.1 Integrating VPN functions.

The function of a remote VPN gateway is pretty much what you would expect from the name; it is the termination point for tunnels from remote clients. Generally, these products concentrate on PPTP and L2TP, but a few, like the Bay Networks Contivity Extranet Switch, also support IPSec remote hosts using IPSec ESP as the tunneling protocol. Some vendors currently shipping PPTP/L2TP gateways have also committed to adding IPSec support as the standards are approved this year.

One class of products that we won't cover in this book are the remote access concentrators used by ISPs to provide tunneling services such as those used by PPTP and L2TP (see Chapters 5, "Using IPSec to Build a VPN," and 6, "Using PPTP to Build a VPN"). Our focus throughout is on products that you would use to create your corporate VPN.

The Price of Integration

Integrating various functions into a single product can be particularly appealing to businesses that do not have the resources to install and manage a number of different network devices and also don't want to outsource their VPN operations. A turnkey installation can certainly make the setup of a VPN much easier than installing software on a firewall and reconfiguring a router as well as installing a RADIUS server, for example. Of course, this presumes that the configuration software for an integrated VPN box simplifies VPN configuration as well—unfortunately, that's not always the case with the products we've seen.

Aside from differences between LAN-to-LAN and dial-up VPNs, some vendors have two different views on how to extend integrated VPN devices. For some, an integrated device is the ideal location for adding any network service that users may access from other locations. Thus, they bundle Web caching, e-mail servers, and DNS caching with their VPN device. Other vendors, on the other hand, see the VPN device as an appropriate location for controlling the network connection, offering bandwidth management and resource reservation as part of the package.

Integrating many functions into a single box can be too much of a good thing, because that box now becomes a single point of failure. It may be one thing to accept that all security functions controlling communications with the Internet may fail when a single device goes down; at least, a broken communications link means attackers cannot get into your intranet over that link. But, it's another thing entirely to put an e-mail server or a Web server in the same box as your security and external communications link; if it fails, your employees can lose some internal services as well.

One of the biggest problems in dealing with all these devices, at least for most of 1998, is their lack of interoperability. It's rare to find a vendor whose VPN product line can cover the needs for all your

sites—corporate, regional, and branch offices—so interoperability can become very important as you attempt to purchase different-size devices for different sites.

But, expect this situation to change in the near future. We should see IETF approval of the second version of IPSec as a standard sometime in 1998, and many vendors already have promised IPSec-compliant products after the standards are settled. One of the obstacles to prior deployment of IPSec has been the lack of a key-management standard, and that's now being handled in version 2 with IKE. Version 2 will make it easier to find interoperable products before long.

Different Products for Different VPNs

Consider that the important functions of any VPN, as we've mentioned throughout the book, are tunneling, encryption, authentication, and key management. Depending on which protocol you're planning to use for your VPN—PPTP, L2TP, or IPSec—there's a different relative emphasis on each of these functions. PPTP, for example, focuses on tunneling and includes weak encryption, and L2TP supports stronger user authentication and relies on IPSec for strong encryption; IPSec, on the other hand, handles encryption and key management well, but still needs work to be used with strong user authentication.

Depending on a product's features, gateways either can be used in place of some existing network devices, like firewalls, or they may be deployed as additional equipment. In either case, where you put a gateway affects not only ingress to and egress from your network but also the amount of traffic on your network. Another thing to keep in mind is that, although a VPN gateway may well integrate a number of functions and thereby simplify the management of the integrated functions, installing a gateway might force you to reconfigure your existing network devices, such as routers and firewalls.

Previous | Table of Contents | Next

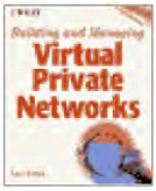

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Recall that if you're planning to use one of the devices covered in this chapter for terminating PPTP or L2TP, then all your users need is a PPP client for dialing into their ISP. Remember that this doesn't provide the strongest security for any data transferred between the client and your corporate site; if you want stronger encryption, you'll need to install either a PPTP or L2TP client on the remote user's computer.

Product Requirements

If you're going to use one of the hardware devices in this chapter as a security gateway for your VPN, you need to consider a few main issues.

To start out, are you going to transfer only IP traffic across your VPN, or will you also have to support IPX and NETBEUI? Many gateways support only IPSec, which is fine for IP-only networks, but that doesn't help if you're running NetWare over IPX, for instance. If you choose not to migrate NetWare to the IP version, don't want to use an IPX-IP gateway, or don't want to replace NetWare entirely to create an IP-only network, then you'll have to use a gateway that supports either PPTP or L2TP, which handle multiple protocols.

Regardless of which protocol is used for your VPN, you need to consider how the product integrates with the rest of your security and network-management systems. For example, many systems depend on RADIUS or token-based systems for user authentication; if you're already using a particular system for authenticating remote users, then picking a gateway that's already compatible with your current authentication system will simplify configuration and management of the gateways.

Alternatively, you may feel the need to improve the security of your authentication systems by installing a two-factor system like SecurID (see Chapter 8) as you roll out your VPN services. Whatever the reason, system compatibility is important.

If you're planning to use an authentication system based on digital certificates, then you should give some thought to how the certificates will be distributed and verified. We'll get into this in more detail in the chapter on security management (Chapter 13), but some of the factors to consider include whether a certificate authority should be maintained in-house or outsourced and how certificates will be linked to other services (through a directory service, for example). Some of the devices covered in this chapter include *Lightweight Directory Access Protocol* (LDAP) links that can be used with certificate servers and, in a few cases, their own LDAP server, for instance.

Looking to the future, network management involves integrating an increasing number of user privileges, such as bandwidth, QoS, remote access, and access to servers and other network resources. This

evolution of network management will drive further use of policy-based management. Policy-based management depends on having an easily managed distributed system for identifying and authenticating users. Directories are thus fast becoming a cornerstone technology for policy-based management, particularly X.500 directories and LDAP. It's worth keeping an eye on how the vendors are tying their products, particularly authentication services, to LDAP and X.500 directories as a first step towards deployment of policy-based management, even if you don't think you will be using such systems for a few more years.

Bear in mind that you're most likely going to be installing these products at more than one site. You'll be able to maintain a more consistent security policy if the product you pick supports synchronized administration of multiple sites. This might involve file exchanges, as we discussed earlier, or some other form of remote management. If remote management capabilities are included in a product, be sure that remote access to the product is secure. One product even bothers to encrypt any requests for log files as well as the reports it creates as a result of those requests.

Check which cryptographic algorithms the product supports. The IPSec default algorithms, the DES CBC algorithm for encryption and HMAC-MD5 or HMAC-SHA-1 hash algorithms for authentication, should suffice for those uses considered to be medium risk; if your traffic is higher risk, then be sure that automatic rekeying is supported.

On a related matter, even though manual keying is the minimum required by the IPSec specifications, you should look for products that allow the cryptographic keys to be changed automatically and periodically or whenever a new connection is established. If interoperability will become a concern, don't settle for a proprietary key-exchange system but insist on the systems described in IKE (see Chapter 5, "Using IPSec to Build a VPN"). Automatic rekeying further strengthens the security of your traffic by increasing the difficulty of cracking a key when it's intercepted (i.e., the key expires before it can be cracked).

Much of key management depends not only on the reliability and security of a certificate authority or certificate server, but also how products react when one part of the key-management process fails or a key is canceled. Some products drop the session immediately when a canceled key is detected, and others will wait until the session is completed; the most secure systems are those that drop the session immediately. Also, to provide additional backup for keys in the case of a CA failure, you can purchase dedicated hardware that stores the appropriate keys for a VPN gateway.

If the gateway supports IPSec, be sure which IPSec headers the product employs. Although the original IPSec standards did not require the support of both AH and ESP headers, it's preferable to apply both headers to each packet. Some IPSec-compliant devices support only the Authentication Header. Because security associations are crucial to the operation of IPSec, you should be able to manually input security associations (usually from a file similar to that recommended by S/WAN; see Chapter 5) and, if possible, specify wild card security associations to simplify configuration. To provide added security to your VPN, see whether the vendor's products support the new IPSec antireplay system rather than using a nonstandard variant.

The gateway always has to be treated as a secure device, which includes protecting the secret and private keys used by the device. Personnel with incidental access to the gateway should not be able to obtain the keys, for example.

As we've said in the past, every security device should have a way of logging security events (*incidents*) and reporting them. If possible, be sure that the system can generate some kind of alarm if some persistent activity takes place, as this may indicate a systematic attempt at breaching the site's security.

Last, if you're considering a product that includes a router, find out what its routing capabilities are. Some of the available products offer limited routing (they may not route IPX protocols, for instance) and need to be deployed in conjunction with routers that already are in place in your network.

An Overview of the Products

At least 17 different VPN hardware products were available as of mid 1998. By the time you're reading this, no doubt even more are available. We've chosen to categorize the products in Tables 11.1 and 11.2 according to their support of LAN-to-LAN VPNs versus remote VPNs.

When you're reviewing this list of products to help you pick likely candidates for your own VPN, bear in mind that the market doesn't stand still, and newer versions of some of these products will either include modified features or new features that were introduced since this book was put together. In other words, use the table as a guide, not the last word.

Previous | Table of Contents | Next

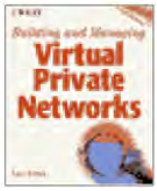

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

It's generally to be expected that these hardware solutions will perform VPN functions, especially encryption, faster than their software counterparts. But, determining the actual performance of these products is difficult. Many vendors will quote throughput values ranging anywhere from 22 Mbps to 60 Mbps in their product literature, for example; use these values only for the roughest comparisons.

Quoted throughput values are usually measured with large packet sizes (1,450 bytes), although normal network traffic often includes many more small packets, which reduces throughput values. Also keep in mind that session initialization places a heavy burden on the processor as the device is calculating session keys and that more sessions lower performance. The products that list support for thousands of sessions cannot set all of them up simultaneously. Just as some vendors have installed special ASICs in their hardware for some of the VPN functions, a few companies also use a dedicated RSA chip for any computations involving public-key cryptography. Cryptographic processing speed will become increasingly important if you're looking for highly secure systems and switch from DES to the slower Triple DES algorithm.

Although many of the hardware devices covered in this chapter are likely to offer you the best performance possible for your VPN, you'll still need to decide how many functions you want to integrate into a single device. Small businesses or small offices without large support staffs, especially those experienced in network security, will benefit from products that integrate all the VPN functions as well as a firewall and perhaps one or two other network services. Some products—usually the more expensive ones—include dual power supplies and failover features to ensure reliability. But, you need to decide which network services are crucial to your company's minute-to-minute operation; after prioritizing those services, you can make the decision whether they should be installed in a single product.

Should you purchase VPN hardware instead of installing additional software and/or hardware on your routers or firewalls? That depends. If you're looking for a low-end solution that doesn't have to process a large amount of traffic, then the products we discussed in Chapter 10, may do the trick. If you don't have a firewall or if you're planning on adding VPN capabilities to branch offices, then some of the integrated boxes described in this chapter will fit the bill; they also can reduce your need for an on-site security specialist or at least reduce some of your network management tasks.

It's hard to beat many of the products in this chapter for throughput and handling large numbers of simultaneous tunnels, which should be crucial to larger enterprises. Some of the products listed in this chapter may seem expensive, but recall when you're making price comparisons that the software prices we mentioned in Chapter 10, "Firewalls and Routers," (and those we'll cover in the next chapter) do not include the prices of the machine on which they run.

Also, don't overlook the importance of integrating the control of other network-related functions, such as

resource reservation and bandwidth control (see Figure 11.6). Some companies already include these features in their products, and it's a step that will most likely gain more support in the future. Integrating traffic control with authentication and access control also makes sense over the long run, as policy-based network management becomes more prevalent and useful. We're only beginning to see the first steps in integrating various management functions and implementations of policy-based management for enterprise net works, and it'll probably be a few more years before we start to see widespread deployment.

FIGURE 11.6 Integrating VPNs and QoS.

Summary

If you're looking for performance, VPN hardware products usually can offer better throughput than their software counterparts. The most basic versions of these products include packet authentication, tunneling, encryption, and key management as well as links to user-authentication systems. More advanced products often package more services into a single box, such as RADIUS and LDAP servers, and support for thousands of simultaneous tunnels.

Previous | Table of Contents | Next

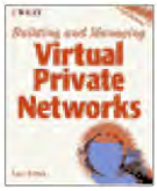

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 12
VPN Software

We now come to the last group of VPN products that we'll cover in this book—software. This group is somewhat of a diverse collection because it covers any software that isn't specifically aimed as an addition to either firewalls or routers, which we covered in Chapter 10, "Firewalls and Routers." Many of the software products covered in this chapter parallel the hardware we covered in Chapter 11, "VPN Hardware," in which a number of different VPN and network services are provided in a bundled product. As you go through this chapter, you'll see that it includes some of the major Network OSs (NOS), such as NetWare and Windows NT, as well as products specifically created for forming and maintaining secure tunnels (AltaVista Tunnel and F-Secure VPN), along with software that can be used for host-to-host tunnels without the need for an intervening security gateway.

It's true that some of the products covered previously, particularly certain firewall products, are also software-based, in which the buyer gets to select the computing platform. But, these products easily fit into the category of firewalls; whereas the products we'll discuss in this chapter cannot be easily categorized. In many ways, this chapter covers a grab-bag of different software products but ones that may be important enough to play a role in the construction of your VPN.

Different Products for Different VPNs

Two classes of software are worth mentioning here. One is composed of the products that provide VPN services for a LAN, much like the hardware that was discussed in Chapter 11. The second class of products are those that can be used for host-to-host tunneling without the need for a security gateway.

The products that provide VPN services for a LAN cover the full gamut of tunneling and VPN approaches, some offering support for the protocols we've covered in this book, and others using proprietary approaches to tunneling and key management.

The evolution of VPN standards, their requisite infrastructures (for digital certificates, for instance), and the current networking marketplace have made LAN-centric solutions a higher priority than host-to-host solutions, which has made the choices for host-to-host software rather small in number so far. Although a few shrink-wrapped products can be used for secure host-to-host connections, some commercially available *software development kits* (SDKs) let developers create their own IPSec-compatible programs.

Tunneling Software

Earlier in this book, when we described tunneling, we pointed out that tunneling was nothing more than encapsulating one packet inside another. In some cases, like with the MBone, the experimental multicasting backbone on the Internet, no effort is made to protect the encapsulated packets. And, with PPTP for example, the amount of protection offered by encryption is rather weak because of the methods employed. IPSec, on the other hand, creates tunnels by applying strong encryption methods to the encapsulated packets.

Now, with VPN software, we see that encrypting encapsulated packets to form tunnels can be done in other ways as well. Of the products covered in this chapter, four use their own proprietary methods for tunneling. And, of course, not one of the methods is compatible with any of the others.

There's much to be said for standards and interoperable products, such as we're seeing with IPSec. Being able to pick and choose among vendors enables you to purchase the best products for your needs without feeling tied to a single vendor; these days, it's highly unlikely that any one vendor has a *lock* on the best networking technology. (Of course, you still have to worry about configuring and managing these different devices if you buy from more than one vendor. Businesses often will go with a single vendor to avoid management and maintenance hassles.)

With the strong move to standardize VPNs using IPSec and L2TP (and PPTP, to a lesser extent), is it wise to use proprietary solutions like the ones mentioned in this chapter? In general, little advantage is gained by using proprietary solutions. A few of these products were some of the first ones created for Internet-based VPNs and thus precede many of the standards efforts. Although we'd much rather use standards-based solutions, we're including the proprietary products for the sake of completeness.

Also keep in mind that vendors change their products over time in response to market pressures. At least two of the products covered here—AltaVista Tunnel and Borderguard—are supposed to include IPSec support before long. Starting out with a proprietary product doesn't keep you from being interoperable with other standards later.

It's also possible to use standard protocols other than IPSec and L2TP to create VPNs. Aventail's use of SOCKS v5 is one such example (see Chapter 10). Another example is DataFellows' use of *Secure SHell* (SSH) in their F-Secure product. SSH is familiar to Unix system administrators for securing communications and has been used on a variety of networks (by NASA and some banks, for instance) for securely transmitting data. Unlike the protocols we've discussed in this book, however, SSH works at the transport layer.

VPNs and NOS-Based Products

Although there will come a time when the authentication and encryption functions of VPNs will be included in each computer as part of the operating system, we're currently forced to rely on using security gateways or remote client software to create VPNs. As a first step to provide VPN support in some of the Network Operating Systems, companies like Microsoft and Novell have started to provide security gateway functions in their NOS software.

We've already mentioned that Microsoft was the first to provide a tunnel server for PPTP in their *Routing and Remote Access Server* (RRAS) product for Windows NT 4.0. Although RRAS is designed to serve as a tunneling server for PPTP (and eventually L2TP) tunnels, either for LAN-to-LAN or host-to-LAN VPNs, it's not a bundling of security services like some other products. For example,

RRAS has a very limited packet filtering system—you either pass PPTP packets or nothing at all. To add the security of a firewall to control access with a finer granularity, you need to add Microsoft's Proxy Server to your server machine. (The Proxy Server was covered in Chapter 10.)

Novell's Borderguard is a series of software modules that can either be used separately or as a unit. Of particular interest to our discussion here are the firewall and VPN modules. The firewall is a fairly generic packet filtering and application proxy type. Borderguard's VPN services use Novell's own tunneling technology to encrypt TCP packets (using the RC2 bulk encryption algorithm). The product also utilizes *Simple Key Management for IP* (SKIP) to exchange keys, although the implementation apparently does not interoperate with key-management servers from other vendors.

Previous | Table of Contents | Next

Although Borderguard may not be suitable for all VPNs, it's an ideal product for companies currently using Netware and IPX, but who want to set up a VPN. Since Borderguard includes an IPX-to-IP gateway, you can set up your VPN without converting your internal IPX networks to IP. And, should you decide to convert to IP, Borderguard could continue to serve as the VPN software as you convert different portions of your network from IPX to IP, because it supports both protocols. However, since only Borderguard installations can take part in the VPN, you wouldn't be able to tie non-Netware sites into the VPN very easily.

Borderguard, RRAS, and Conclave all emphasize the same dilemma: How many different network services should you install on a single computer? There are actually two issues here. First, there's the single-point-of-failure argument, which we discussed in Chapter 11, "VPN Hardware": How many of your network services can you afford to lose if the computer fails? Second, there's the issue of performance: Will installing too many services on a single computer seriously impair the performance of important services? This issue doesn't come up as often for the integrated hardware we discussed in Chapter 11, because many of those products use customized operating systems and hardware for better performance. But, running a variety of network services on a single computer using an OS designed for a variety of computing tasks, which may not be optimized for your network services, could lead to poor performance. (For example, it's often recommended not to install all of Borderguard, especially the Web-caching module, on a single machine.)

Host-to-Host VPNs

Throughout most of this book, we've written about either LAN-to-LAN VPNs or dial-in VPNs, both of which involve some kind of security gateway. But, another kind of VPN involves communications between each individual host; this host-to-host (or end-to-end) VPN doesn't require any security gateways to create tunnels or encrypt packets, because all encryption is done at the host (See Figure 12.1). Rather than use tunnel-mode IPSec, as security gateways would, host-to-host connectivity is set up using transport-mode IPSec (see Chapter 5, "Using IPSec to Build a VPN").

FIGURE 12.1 LAN-to-LAN versus host-to-host connections.

Although the IPSec standards provide for host-to-host connectivity, the majority of products currently available for IPSec VPNs have focused on the use of a security gateway for a number of reasons. First, the market is relatively young, and deployment of security measures at the edges of your network (i.e., routers, firewalls, and VPN gateways) makes it easier to discover the vagaries of VPNs. Second, key

management is considerably easier when a limited number of security gateways are involved, rather than the more numerous individual hosts on all your LANs. (Although managing keys for remote hosts using dial-in VPNs might be comparable to taking care of all the hosts on your network.) Third, the performance penalty for encrypting or decrypting packets can be considerable for many existing computers and is likely to slow down an individual workstation's operations too much to be usable on a routine basis.

All of the preceding factors may lead to slower deployment of encryption at the level of individual hosts, but they don't present insurmountable obstacles to setting up host-level VPNs. Another thing to keep in mind is that IPSec originally was developed in parallel with efforts to define the next version of IP, IPv6, and is a mandatory feature of IPv6, meaning that all IPv6 protocol stacks and drivers will contain IPSec. On the other hand, IPSec is an optional feature in IPv4, and TCP/IP stacks have to be modified in order to use IPSec.

Product Requirements

When selecting VPN software for a LAN, the product requirements are much the same as the ones we've mentioned in the past two chapters. We'll briefly review them here.

Protocol support. First, consider which protocols will be transmitted across your VPN—IP only or IPX and NETBEUI as well? Many gateways support only IPSec, which is fine for IP-only networks, but that doesn't help if you're running NetWare over IPX, for instance.

Integration with existing systems. You also need to consider how the product integrates with the rest of your security and network management systems. For example, many systems depend on RADIUS or token-based systems for user authentication; if you're already using a particular system for authenticating remote users, then picking a gateway that's already compatible with your current authentication system will simplify configuration and management of the gateways.

Digital certificate issues. If you're planning to use an authentication system based on digital certificates, then you should give some thought to how the certificates will be distributed and verified. We'll get into this in more detail in the chapter on security management (Chapter 13), but some of the factors to consider include whether a certificate authority should be maintained in-house or outsourced and how will certificates be linked to other services (through a directory service, for example). Some of the devices covered in this chapter include LDAP links that can be used with certificate servers and, in a few cases, their own LDAP server, for instance.

Multisite maintenance. Keep in mind that you're most likely going to be installing these products at more than one site. You'll be able to maintain a more consistent security policy if the product you pick supports synchronized administration of multiple sites. This might involve file exchanges, as we discussed earlier, or some other form of remote management. (VTPC/Secure, for example, creates a floppy disk with the required configuration for each VPN gateway; these disks are then taken to the gateway computers to complete the installation of the software.) If remote management capabilities are included in a product, be sure that remote access to the product is secure.

Cryptographic algorithm support. Check which cryptographic algorithms the product supports. The IPSec default algorithms, DES CBC algorithm for encryption and HMAC-MD5 or HMAC-SHA-1 hash algorithms for authentication, should suffice for those uses considered to be medium risk; if your traffic is higher risk, then be sure that automatic rekeying is supported.

Automatic rekeying further strengthens the security of your traffic by increasing the difficulty of cracking a key when it's intercepted (i.e., the key expires before it can be cracked). Even some of the systems mentioned in this chapter, which do not use IPSec, support some kind of automatic rekeying procedure.

If the software supports IPSec, be sure you know which IPSec headers the product employs. Although the original IPSec standards did not require the support of both AH and ESP headers, it's preferable to apply both headers to each packet. Software for host-to-host communications should support transport-mode IPSec, although security gateways are more likely to require only tunnel-mode IPSec.

Because security associations are crucial to the operation of IPSec, you should be able to manually input security associations (usually from a file similar to that recommended by S/WAN; see Chapter 5) and, if possible, specify wild card security associations to simplify configuration. To provide added security to your VPN, see whether the vendor's products support the new IPSec antireplay system rather than using a nonstandard variant.

Incident logging. Every security gateway should have a way of logging security events (*incidents*) and reporting them. If possible, be sure that the system can generate some kind of alarm if some persistent activity takes place, as this may indicate a systematic attempt at breaching the site's security.

Previous | Table of Contents | Next


Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Turning to host-based software, as transport-mode IPSec is deployed to individual computers in the future, it's likely that certificate servers, perhaps using LDAP, also will become more widely deployed (see Figure 12.2). Couple that with the increased usage of digital certificates in general—for electronic commerce and secure e-mail, for instance—and a centralized policy for creating and managing certificate servers will be needed. Try to factor that into your planning as you review these products. Much of the work on implementing *Public Key Infrastructures* (PKIs) is still under development.

An Overview of the Products

The bulk of the products in Table 12.1 is software for creating and managing tunnels, either between a pair of security gateways or between a remote client and a security gateway. E-Lock can be used for host-to-host VPN connections; other host-to-host products are part of systems mentioned in Table 11.1.

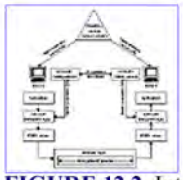

FIGURE 12.2 Interactions between two hosts and a directory server.

It's not possible for a book to include completely up-to-date information on products, especially considering the pace of networking products' introductions and modifications. This table should serve as a guide of what's available, but not as the final word on either the entire marketplace, a vendor's product line, or a particular product's feature set.

In particular, don't make price comparisons using only the data included in the table. Only prices for the server and client portions of the software are given, and client prices frequently change with the number of users. Also, look out for different pricing schemes; for example, Conclave is priced on the basis of the number of users, not the number of sites.

TABLE 12.1 VPN Software

Product	AltaVista Tunnel 98 (Compaq)	BorderManager (Novell)	Conclave (Internet Dynamics)	E-Lock (Frontier Technologies)	F-Secure VPN (DataFellows)	Omniguard/PowerVPN (Axent)	PrivateWire (Cylink)	RRAS/NT Server (Microsoft)	SmartGate (V-ONE)	VTPC/Secure (InfoExpress)
Price	$995+	$2,495+	$2,495 (25 users, no CA) $3,995 (25 users, CA)	Desktop: $249, Director: $,1000	$2,945+	$2,995	$19,000+	free*	$4,995 (Windows), $6,495 (Unix)	$1,495
Platforms—server	Windows NT, Digital Unix	Solaris, SunOS, NetWare, Windows95, NT	Windows NT, Netware, Unix	Windows NT, 95	NetBSD kernel	Windows NT, Solaris, BSDI, HP/UX, Linux	Windows NT, Solaris	Windows NT	BSD Unix, Solaris, HP-UX, Inix, Windows NT	Windows NT, Unix
Platforms—remote access	Windows 95, NT	Windows 95, NT	Windows, Unix, Mac	Windows NT, 95	No	Windows 3.x, 95, NT	Windows 3.x, 95, NT	Windows 95, NT	Windows 3.x, 95, NT, OS/2, Mac (MaXVPN Windows 05, NT only)	Windows 3.x, 95, NT, Solaris, Linux, Mac

Tunneling Protocol	proprietary	proprietary	IPSec, PPTP	IPSec	SSH	PPTP	proprietary	PPTP	tunnel-independent	proprietary
Protocols supported	IP	IP, IPX	IP	IP	IP	IP	IP	IP, IPX, NetBEUI	IP	IP, NetBios
Encryption type	128-bit RC4	128-bit RC2	DES, Triple DES, RC2, RC4	DES, Triple DES	RSA, Triple DES, Blowfish	DES, Triple DES	DES, Triple DES, RC4	RC4	DES, RSA	DES, Triple DES
User authentication	SecurID, LAN segment	RADIUS	SecurID	—	TACACS+, SSH	RADIUS certificates, SecurID	CHAP, PrivateCard	token or MSCHAP, PAP	software ACE, RADIUS, smartcard	TACACS+
Access controls	none	source, destination, protocol	document-level, source, destination	none	none	none	—	PPTP filtering	destination, port, URL	source, destination, protocol, port, user, group
User management integration	NT user domain	NDS	Windows ID, X.509	digital certificates, LDAP	—	Defender software tokens	proprietary directory	NT user domain	proprietary	—
Key management	RSA	SKIP	SKIP	manual, Diffie-Hellman, DNS, LDAP, PKIX	proprietary (RSA)	—	Diffie-Hellman	Diffie-Hellman	—	Diffie-Hellman
# tunnels supported	2,000 (on Unix server)	—	—	—	100	—	—	256	—	unlimited
# nodes supported	unlimited	—	—	—	—	—	—	—	—	unlimited
Integrated firewall	No	packet filter, application proxy	packet filter, application proxy	No	No	No	Yes	No (incl. with NT server)	No	Yes
Remote management	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Certificates	Yes		X.509	X.509	—	—	Yes	No	X.509 (optional)	Yes
Remote access client	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Compression?	Yes	No	No	(planned)	Yes	Yes	—	Yes	—	Yes

Some of the products listed here are bundles of more than one product. If you don't need all of the service listed for a particular product, it's a good idea to check with the vendor to see whether he offers unbundled solutions. For example, Novell has decided to offer some of the modules in Borderguard as separate products.

As we've said before in this chapter, be sure to balance your network's requirements with a product's performance, especially if you're planning to install multiple network services on a single computer. If you find that the best performance is achieved by distributing applications across a few computers, consider how you're going to keep them all secure from both physical and electronic tampering—in such cases, a single piece of hardware, picked from the list in Chapter 11, "VPN Hardware," might be a better choice.

Given that we've said that VPN hardware is likely to provide the best throughput of any VPN products, why would you bother with some of the software-based products we're covering in this chapter?

First, there's the cost. Some of these products are relatively cheap, or even free in the case of Microsoft's RRAS as long as you already have Windows NT. If you already have an appropriate computer on which to install the software, the cost of deploying a VPN is reduced even more.

Second, you already may be familiar with the operating system or NOS on which the software runs, which may make management of the VPN more appealing. Conversely, installing a new OS

just to create your VPN may not make sense; if you're a Unix shop, for instance, you may not want to go through the trouble of buying Windows NT and installing RRAS or Conclave.

Lastly, the services and performance that these products provide may be all you need. If you're building a small VPN or handling small amounts of traffic, you may not require the performance (and price) found in many of the hardware products we've covered in Chapter 11.

Summary

Many of the software products for creating VPNs use proprietary tunneling protocols and nonstandard methods for key exchanges, limiting their interoperability. But, some of those same products are likely to become IPSec-compatible this year, improving their interoperability.

If interoperability isn't an issue, some of these products do a good job of filling certain market niches. For example, Microsoft's RRAS is a good PPTP tunnel server for Windows NT servers (and it's free); Novell's Borderguard is especially useful if you're dealing with IPX as well as IP protocols.

Moving from security gateways to networks on which each computer manages its own VPN sessions and keys is likely to happen more frequently in a few years. In the meantime, a limited number of products support host-to-host VPNs using IPSec transport mode.

Previous | Table of Contents | Next

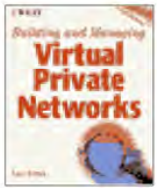

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

PART IV
Managing a VPN

VPN management consists of three main areas: security, IP address allocation, and network performance. Security management includes not only authenticating users from other locations and controlling their access rights but also managing the cryptographic keys associated with VPN devices. VPNs often link together previously isolated networks, which entails extending IP addressing and name management across the entire enterprise and can lead to numbering conflicts.

You should be able to link your VPN's security and address management to existing corporate policies and services. But you will probably have to deploy new technologies to provide performance management on the WAN links usually used for VPNs. You'll not only have to decide how to differentiate network services according to business needs but also how to implement policies for the network to provide different performance for different classes of traffic. Lastly, you'll have to match your performance requirements with what your ISP can provide.

CHAPTER 13
Security Management

Entire books have been written about securing computers and networks and the data that either is stored on devices or flows through them. But, that's not our mission here. Although we'll say a few things about corporate security policies in general, our concern here is covering the issues surrounding the management of VPN-related security. To that end, we'll focus on selecting encryption algorithms and key lengths, distributing keys and associated information in IPSec security associations, as well as user authentication and the control of access rights. Because of the importance of authenticating users and devices with digital certificates, we'll spend some time discussing the details of in-house management of certificates.

As we've discussed in previous chapters, IPSec offers the widest range of options for securing data of any of the protocols and includes perhaps the most complex architecture for negotiating and supporting those options. Because of those options and complexity, much of what we have to say about managing security for VPNs in this chapter will focus on IPSec; coverage of PPTP and L2TP also is included where pertinent.

Corporate Security Policies

There's much more to corporate security than what's covered in this book. A proper security framework for an organization includes seven different elements: authentication, confidentiality, integrity, authorization, nonrepudiation, administration, and audit trails (see Figure 13.1). Networking security is just one part—albeit an important part these days—of corporate security and should fit in with your corporate security policies.

A solid security policy should do the following:

Look at what you're trying to protect

Look at what you need to protect it from

Determine how likely the threats are

Implement measures that will protect your assets in a cost-effective manner

Review the process continually

Improve things every time a weakness is found

A traditional security policy identifies all of the assets in the corporate information infrastructure that are being protected, corporate databases and computer hardware, with overall policies on how to protect these assets. This policy should include everything from physical access to the property, general access to information systems, and specific access to services on those systems.

FIGURE 13.1 The components of a secure system.

But, as information systems have become more distributed, corporate security policies have had to include guidelines governing department LANs as well. This means adding policies on who has access to resources belonging to different departments: Can Sales access the R and D servers, for instance, or who can read the divisional manager's e-mail?

As you define security policies for your network, identify every access point to your information system and define policy guidelines to protect that entrance/exit point. Don't overlook all those modems that employees may have in their offices, which can become inviting access points for hackers as users dial into on-line services.

Some of the questions you should ask when formulating a security policy include the following:

- Which Internet services does the organization plan to use?
- Where will the services be used? Are they to be used on a LAN or via remote access?
- What additional needs (e.g., encryption) may be supported?
- What risks are associated with providing these services and access?
- What is the cost, in terms of control and impact on network usability, of providing protection?
- What assumptions are made about security versus usability?

Integral to your security plan is the capability to monitor compliance and respond to incidents involving violations. As part of your security policy, an emergency response procedure should be defined.

One of the new issues in security policy that arises from VPNs is that of key management. In the past, if a leased-line VPN was used, link-layer encryption may have been used, which didn't require exchanging cryptographic keys. But, the more dynamic nature and added flexibility of Internet-based VPNs requires a wider distribution of keys and more frequent rekeying, which in turn requires more complicated systems for key management. This is especially true when remote users are involved.

Even though we've said very little about the content of traffic passed between hosts, securing content against attacks, such as from computer viruses, also should be an important part of your security policies. Viruses are here to stay, so to prevent costly infections from spreading, antiviral scanning software should be included in your security implementation.

Now let's turn to some of the details of VPN security management.

Selecting Encryption Methods

As you set up your VPN, you'll find that there are two major constraints on securing your data to the desired degree (after you've selected the VPN protocols you'll use). First, even though some protocols like IPSec support a variety of encryption protocols in the specifications, not all products include every encryption algorithm. Second are the country-specific restrictions on exportable key lengths. For instance, here in the United States, you're usually restricted to using 40-bit or 56-bit key lengths with DES for export, although you can use 128-bit key lengths within the United States.

After you've collected and analyzed the corporate data we outlined in Chapter 8, "Designing Your VPN," and selected the products for your VPN from Chapters 10–12, you should be able to select the appropriate algorithms and key lengths.

Previous | Table of Contents | Next

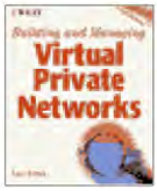

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Protocols and Their Algorithms

Each of the VPN protocols we've discussed in this book—IPSec, PPTP, and L2TP—specify their own list of allowed algorithms for encrypting data.

Although PPTP can use PPP and its negotiable encryption options (including DES and Triple DES) to encrypt data, Microsoft has incorporated an encryption method called *Microsoft Point-to-Point Encryption* (MPPE) for use with PPTP tunnels. MPPE uses the RC4 algorithm with either 40-bit or 128-bit keys, depending on export restrictions. Similarly, L2TP can use PPP to encrypt data, but the preferred method is to use IPSec for this task.

Within IPSec, the default encryption algorithm for use in ESP is DES with an explicit initialization vector. IPSec allows alternative algorithms to be used. These include Triple DES, CAST-128, RC5, IDEA, Blowfish, and ARCFour (a public implementation of RC4).

The choice of supporting algorithms other than DES is left to vendors, so you may find that a vendor's products do not support the alternative algorithm you had planned on using. DES and Triple DES seem to be the most common algorithms supported thus far. There's a definite benefit to having a choice of encryption algorithms: Would-be attackers not only must break the cipher, but they also must determine which cipher they are attempting to break.

Recalling the Oakley modes used in IPSec (Chapter 5), main mode negotiates the encryption method, hash, authentication method, and Diffie-Hellman group between VPN endpoints. The Diffie-Hellman group determines the strength of the keying material; there are 4 Diffie-Hellman groups. Diffie-Hellman Group 1 is strong enough for DES, and Groups 2 and 3 should be used for Triple DES. Because main mode might require six packets, if you're using high-latency satellite connections, for example, it would be better to use the stronger Diffie-Hellman group, even for DES.

Oakley's quick mode also negotiates the algorithms and lifetimes for IPSec. These lifetimes determine how often, based on time or data, another quick-mode negotiation is required. The main-mode lifetime controls the Oakley SA, and the quick-mode lifetime controls the IPSec SA. As an example, the quick-mode lifetime could be set to 15 minutes, or 10 MB, and the main mode lifetime set to 1 hour, or 40 MB, when DES is being used for IPSec. These lifetimes would be increased for Triple DES, because it's more secure than DES, or decreased for ARCFour, because it's less secure than DES. The idea is to balance the strength of the IPSec services and the strength of the underlying cryptographic algorithms against the cost of ISAKMP/Oakley packet overhead; too many changes in keys could affect the efficiency of your network.

Key Lengths

Back in Chapter 8, "Designing Your VPN," we suggested that you determine the sensitivity of your data so that you could calculate how long it will be sensitive and how long it'll have to be protected. When you've figured that out, you can select an encryption algorithm and key length that should take longer to break than the length of time for which your data will be sensitive.

As a starting point, take a look at Table 13.1, which is a condensation of information from Bruce Schneier's book, *Applied Cryptography*, which we also used in Chapter 4, "Security: Threats and Solutions." The table does a good job of illustrating that many of the key lengths currently in use can be broken with a relatively small outlay of funds. This table also helps emphasize this point: know your attacker. If you expect that highly skilled and well-funded industrial spies will be attempting to intercept and decrypt your data, then long key lengths and frequent rekeying are an absolute necessity.

TABLE 13.1 Comparison of Time and Money Needed to Break Different Length Keys

	Length of key in bits				
Cost	***40***	***56***	***64***	***80***	***128***
$100 K	2 secs.	35 hrs.	1 yr.	70,000 yrs.	10^{19} yrs.
$1 M	.2 secs.	3.5 hrs.	37 days	7000 yrs.	10^{18} yrs.
$100 M	2 millisecs	2 mins.	9 hrs.	70 yrs.	10^{16} yrs.
$1 G	.2 millisecs	13 secs.	1 hr.	7 yrs.	10^{15} yrs.
$100 G	2 microsecs	.1 sec.	32 secs.	24 days	10^{13} yrs.

These estimates are for brute-force attacks, that is, guessing every possible key. There are other methods for cracking keys, depending on the ciphers used, which is what keeps cryptanalysts employed, but estimates for brute-force attacks are commonly cited as a measure of the strength of an encryption method.

Remember that this is not a static situation either. Computing power is always going up and costs are falling, so it'll get easier and cheaper to break larger keys in the future. Off-the-shelf processing power (costing around $500 thousand) can crack the 56-bit DES code in 19 days; hackers choosing to invest in custom chips could break the code in a few hours. A student at UC Berkeley used a network of 250 workstations to crack the 40-bit RC5 algorithm in three and a half hours.

Key Management for Gateways

A number of keys are usually required for secure communications between two gateways. First is the key pair that identifies two gateways to each other; these keys might be hard-wired, exchanged manually, or transmitted via digital certificates. Second are the session keys required for authentication and encryption of the packets transmitted between the gateways, using IPSec's AH and ESP headers, for instance. Different keys are required for each IPSec header and are negotiated via security associations (see Chapter 5). If both AH and ESP are used to process packets, for instance, then two SAs are negotiated

between the gateways or hosts.

Identification of Gateways

Before a secure tunnel can be established between two security gateways, or between a remote host and a gateway, these devices have to be authenticated by each other and agree on a key. First, let's look at exchanges between two gateways. This authentication is not the same as the authentication of packets using the AH header; here, we are authenticating the devices themselves.

Previous | Table of Contents | Next

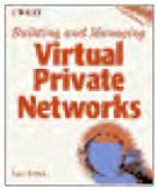

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Gateways using public-key pairs can be authenticated manually. In such cases, the key pairs are usually hard-wired into the device before it's shipped. The network manager then registers the new device with other security gateways on the VPN, giving those gateways the public key so that they can exchange session keys.

If a security gateway isn't shipped with hard-wired keys, the gateway would be set to randomly generate its own key pair. A digital certificate then would be signed with the private key and sent to the appropriate certificate authority, either an in-house certificate server or a third-party CA like VeriSign. When the certificate is approved, that certificate is available from the CA for use by other security gateways and remote clients to authenticate the site before any data is exchanged (see Figure 13.2).

Although these certificates do not need to be standardized (using the X.509 standard, for instance) if only one vendor's products are used for the VPN, interoperability between products is possible when X.509 certificates are employed. More vendors are adopting this approach, which also makes it easier to utilize an outside certificate authority for storing the necessary certificates. This might be a necessity if you're expanding your VPN to include partners in an extranet.

Other gateways and remote hosts usually will obtain the appropriate certificate from a CA to authenticate the destination gateway by using mechanisms such as LDAP or HTTP to retrieve certificate information through a *public key infrastructure* (PKI) Existing PKIs also require checking *certificate revocation lists* (CRL) to ensure the validity of an existing certificate. However, because the CA hierarchy for verifying certificates that we described in Chapter 4, "Security: Threats and Solutions," can become ungainly, and CRLs can become complicated to handle, other mechanisms for verifying certificates are being developed and most likely will become available starting in 1999. In particular, the IETF PKIX working group has been working on definitions for an interoperable PKI and *Online Certificate Status Protocol* (OCSP). OCSP aims to provide an efficient method for handling compromised or revoked certificates (see "Hardware-Based Security" later in this chapter).

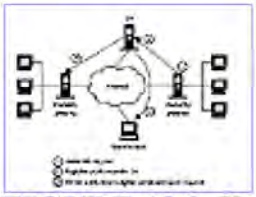

FIGURE 13.2 Key exchanges between gateways.

These systems are based on assigning a public-key pair to each security gateway, and the public key is published in a directory that's accessible to all VPN sites. At the start of each encrypted session, the session key is scrambled by combining the security gateway's private key with the recipient's public key.

Handling Session Keys

If key exchange (such as in IPSec and L2TP) is required between sites, the most basic method is to exchange keys manually. An initial session key is randomly generated by one security gateway and then the network manager has to deliver the key to the administrator of the second device, by telephone, registered mail, or bonded courier, for instance. The second administrator inputs the key to the second security gateway, and a secure session between the two gateways can take place. New keys are generated as required (perhaps once a week) and distributed in the same fashion as before.

This approach is rather cumbersome and not particularly secure; phone lines can be tapped and mail can be intercepted. Dynamic key management, using IKE for instance, is much easier and better suited to frequent key changes and large numbers of sites. Session keys are randomly generated, either by the initiating security gateway or a key-management server, and distributed over the network. The session key itself is scrambled using the recipient's public key before transmission on the net.

Hardware-Based Security

Hardware-based encryption products are less vulnerable to physical attack, which reduces the chances of compromised device keys and, therefore, the need to exchange new keys between gateways. Whether keys are hard-wired or entered from a management workstation, most hardware encryptors are strongly sealed against physical prying and usually erase any stored keys when disturbed.

If session keys are compromised, you'll need a way to revoke a key pair and assign a new one. The procedures for key revocation vary from product to product. How security gateways respond to a revoked key also varies among products. The best, most secure method is to drop the session and log the failed attempt as soon as a revoked key is detected. Some products wait for a session to be completed before denying any further access with that key.

If you're in a situation where your VPN is restricted to shorter length keys than you would like (due to export restrictions, for example), then you should try to improve the security of your VPN sessions by increasing the frequency of rekeying. If keys are used for shorter lengths of time, then any attacker will have less time to acquire the information he needs to break a key; also, the amount of data that could be obtained with a compromised key will be reduced.

Key Management for Users

Generating and distributing keys for LAN-to-LAN VPNs can be a relatively simple process to manage when the number of sites is not very large. Even if the number of sites is in the hundreds, a dynamic system using an outside certificate authority or in-house certificate server should not involve a great deal of management overhead. On the other hand, managing keys for remote users, if they number in the thousands, needs to be as scalable and automated as possible. An automated system also is required if you plan to use the antireplay protection in IPSec. Distribution of the keys and associated information, in particular, may be especially tedious and time-consuming.

Back in Chapter 5, "Using IPSec to Build a VPN," we pointed out that a pair of IPSec devices has to establish a security association with each other in order to communicate. If you're planning to support a large number of remote users with a security gateway, then you'll need to generate the client security

associations centrally and probably in large numbers. The most practical way is to set up a central site to generate all IPSec SA parameters needed and to provide a mechanism to import them into the client. For example, a central site could generate SA data in S/WAN format and send the appropriate information to each client user.

The IPSec architecture is based on the assumption that hosts assign the *Security Parameter Index* (SPI) for inbound IPSec headers, but the essential requirement is simply that inbound SPIs be unique. If you set up a central site for generating keying material and assigning the SPIs to use them, then the client software can use those SPIs for communications without having to generate any on their own.

Previous | Table of Contents | Next

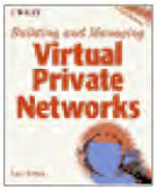

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Protecting Clients Against Theft

Because laptops are particularly susceptible to theft, they pose a special security risk to your VPN because the keys stored on a stolen laptop could be used to access corporate resources via the VPN.

There are essentially three techniques for protecting keys from theft:

1. Store the keys on a removable device like a disk or smartcard and carry it separately from the clients.
2. Encrypt the keys with a secret password or phrase and require the client to verify the password before IPSec can be used (some smartcards can do this as well).
3. Encrypt the keys with a secret password or phrase and let IPSec processing fail if the wrong password is used.

Of these options, the third is the most secure. But, it also can be the most annoying to a legitimate user if he accidentally enters the wrong password because he may not be able to discern the reason for a communications failure.

Remote VPN users have to be authenticated by security gateways in much the same way as security gateways have to identify themselves to each other and be authenticated. The number of options for authenticating users are greater, though, and will be covered in the following section. Since the use of digital certificates for user identification is becoming more popular and is now being supported by more VPN products, we'll discuss the details of managing certificates for users in "Managing an In-House CA" later in this chapter.

Authentication Services

As we discussed in Chapter 4, a variety of ways exist to authenticate users: simple passwords, one-time passwords, challenge/response systems using RADIUS or TACACS+, or two-factor systems using tokens, as well as digital certificates. If you're already supporting remote access via a modem bank and remote access server, for instance, then you already may have an authentication system in place and you'll just need to link it to your security gateway to control the authentication and access rights of your VPN users.

If you're using PPTP, L2F, or L2TP to create tunnels, you might be using your ISP as a tunnel endpoint, as we discussed in Chapters 6, "Using PPTP to Build a VPN," and 7, "Using L2TP to Build a VPN." Should that be the case, the ISP should be running its own authentication server, which, in turn, is a proxy to your authentication server (see Figure 13.3). This enables you to maintain control over setting

authentication parameters and access rights but lets the ISP use that information to provide access to your remote users.

Assuming that you may be setting up an authentication system for the first time as you roll out your VPN, you can choose from any of the different approaches we mentioned earlier. (See Chapter 4, "Security: Threats and Solutions," for more details.) The better systems are RADIUS, token-based authentication, and digital certificates.

RADIUS has three advantages. It's been standardized by the IETF, and many vendors offer products that interoperate. RADIUS also is used by the majority of ISPs for authenticating their customers. Lastly, RADIUS can act as a centralized authentication database, drawing on rights defined for users from various network operating systems such as NT's user domains and NDS trees, making it a good platform for integration.

Both RADIUS and TACACS+ let you define how to authenticate and pass session variables, such as protocol types, addresses, and other parameters. An important feature in both of these systems in the capability to define server-based access control policies. These policies can include time-of-day restrictions, usage quotas, simultaneous login controls (each user can use only one session at a time), and login threshold violations (accounts are locked after X number of consecutive failed logins). RADIUS also can be used for accounting purposes, although quota enforcement requires constant tracking of each user's time online and can become a relatively complex task.

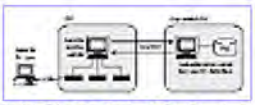

FIGURE 13.3 Main and proxy authentication servers.

A RADIUS server usually consists of three main files: a database of authorized users, a file of client access servers that are authorized to request authentication services, and a set of customized options, called *dictionaries*, for each remote access server or security gateway. If you were configuring RADIUS for use with an ISP and PPTP, for instance, you would add the name or address of the ISP's proxy server to the file of client access servers and define a new dictionary for that server, describing any special authentication and authorization features for the server. (TACACS+ does not include dictionaries in its architecture.)

Token-based authentication usually requires using a special reader attached to the workstation or laptop and a token card that generates special passcodes that are checked by a secure server on the network before access is granted to the user. Before users are permitted to authenticate themselves, token devices request a PIN. Two of the more popular mechanisms for verifying users are a challenge-response system (see Chapter 4) or time synchronization, which depends on synchronized clocks and a frequently changing secret key that the user has to enter when logging in. Although tokens are a very secure method for authentication, because they use a two-factor method (i.e., something the user has—the token card—and something the user knows—the PIN), they can prove to be awkward to use because of the additional hardware that's required.

It's also possible to use digital certificates for authenticating users, although these systems aren't nearly as widespread as RADIUS servers. That's likely to change with time. Some of your employees already may be using personal digital certificates with their Web browsers or e-mail clients. If you're already