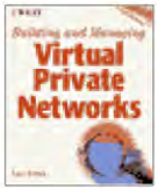

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

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

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).

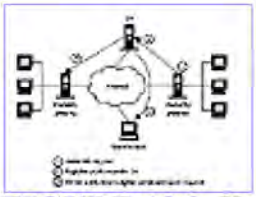

FIGURE 13.2 Key exchanges between gateways.

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

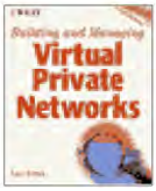

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

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.

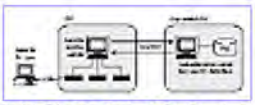

FIGURE 13.3 Main and proxy authentication servers.

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

sending secure e-mail that requires digital certificates and a public-key infrastructure (see Chapter 4), then you could use the same system for issuing and storing the digital certificates required for authentication on the VPN. If not, you'll have to set up an appropriate certificate authority for your users; this could be either a commercial CA like Verisign or an in-house certificate server that you maintain.

Using a third-party certificate authority can make certificate management easier, because it'll be accessible via the Internet and would be more readily accessible to any extranet partners. But, if you're supporting only internal users (i.e., corporate employees), an in-house CA should work just fine. Since your security gateways will be performing authentication of the VPN's users, outside access to the digital certificates isn't necessary. But, you'll still need to secure the computer used for issuing and storing digital certificates, as well as any backup files or tapes (see the next section).

Previous | Table of Contents | Next

Whichever course you choose, you'll need a way to initially distribute to each user the approved digital certificates and the private keys they contain. Outside CAs normally handle certificate distribution via e-mail or HTTP. If you're running your own certificate server, you can do the same, but you also might choose a physical means, such as a floppy disk or smart card. Many companies are employing smart cards for distributing and storing digital certificates because of their portability and the fact that they also can be further secured with a user-specific PIN that makes the card useless if lost or stolen. Plus, these cards shut down after a series of failed login attempts.

Managing an In-House CA

Digital certificates have a limited lifecycle (see Figure 13.4); after they're issued, they should be expired after a reasonable period of time (6 months, for instance) or may be revoked if the owner changes jobs or his private key is compromised. Certificates also can be renewed and need to be backed up in case keys need to be recovered at a later date. If you want to run your own certificate authority in-house, managing the system will involve not only creating key pairs and issuing certificates but also managing those keys and certificates. Certificate management includes maintaining a certificate repository, revoking certificates as needed, and issuing *certificate revocation lists* (CRLs). Key management involves key backup and recovery, automatic updates of key pairs (and their certificates), and management of key histories.

As you plan the deployment of a private certificate server, keep in mind that the infrastructure for digital certificates and certificate management is still in a relatively early stage of development. The use of CRLs for monitoring revoked certificates is also inadequate for dynamic situations, such as you're likely to encounter with remote users accessing a VPN. But, new solutions, like Online Certificate Status Protocol (OCSP), are also under development. Furthermore, commercially available certificate servers still need improvement of their support for administrative tasks. In-house CA systems can be purchased from Certco Inc., Entrust Technologies Inc., GTE CyberTrust Inc., Microsoft Corp., Netscape Communications Corp., Security Dynamics Technologies Inc., and Xcert Software Inc. Some VPN product vendors include a CA as an option, although many of these products are software-based CAs installed on a workstation. Radguard offers a sealed hardware-based CA for use with cIPro.

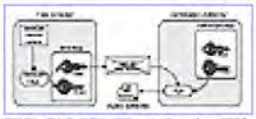

FIGURE 13.4 The lifecycle of digital certificates.

OCSP: A Dynamic Way to Track Certificates

There currently is no practical way to revoke a certificate if the password that unlocks a user's certificate is breached or when the user's private key is compromised. The only solution certificate servers offer right now is the CRL. Ultimately, we'll be looking to the PKIX standard and OCSP for a real solution.

The only way you can use CRLs is by laboriously matching up two lists (the list in your local storage and the one in the CRL) and deleting certificates by hand. OCSP moves away from this static-list model toward a more dynamic one. OCSP defines LDAP and HTTP status queries that are designed to provide fast response time and high availability. In response to a client query, an OCSP server sends a simple status message—valid, invalid, revoked, not revoked, or expired. Using this model, the load is balanced between the client and the server, and it becomes possible to do real-time certificate checking on a per-transaction basis.

Despite these problems, a private certificate server can be installed and managed within your corporation to authenticate both security gateways and users on your network. Let's take a look at some of the features that a useful certificate system should include.

The basic task of a certificate server is to accept requests for new certificates, queue them for their review by the system administrator, and issue the certificates for client retrieval (see Figure 13.5). In general, certificate servers accept certificate requests from a certificate-management workstation, when an administrator is performing batch issues of certificates, or from individuals themselves via HTTP or e-mail. Whenever new requests for certificates are received, they should be matched against certificates held in the directory to prevent accidental obsolescence of valid public keys. The user's certificate can be presented to its owner via HTTP or e-mail as well, or transferred to a disk or smart card for manual distribution.

The private keys of each public-key pair that's issued should be stored within a central repository that's secured against unauthorized access. This repository also should be backed up in a secure fashion (usually as encrypted files), because it becomes part of the key-recovery system should older messages need to be decrypted if a key is lost or compromised and revoked. Backup tapes must be carefully guarded and strictly accounted for.

Since you'll be signing all issued certificates as a certificate authority, a special, dedicated workstation will need to be set up for storing the private keying material (your *root certificate*, as it's called); this same workstation also will include any of the software (and any special hardware, if it's required) for collecting, signing, distributing, and revoking certificates. This workstation should be physically secured against unauthorized access; it should not be treated as a multipurpose computer.

Because one of the purposes of digital certificates is to distribute a public-key pair, the certificate system needs a way of making the public key available to those who need it. The usual method is to store the public keys in a directory. Although large-scale master directories for certificates may be based on X.500, there's been a significant move to use another protocol, *Lightweight Directory Access Protocol* (LDAP), to utilize much of the structure of X.500, but over TCP/IP. Many certificate servers now offered for use at corporate sites are based on LDAP. The increasing popularity of LDAP for directory access also will make it easier for you to link other directory-based services to your digital certificates.

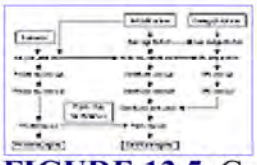

FIGURE 13.5 Generating and certifying a public key.

Certificate servers also have to maintain and make available a Certificate Revocation List which lets users know which certificates are no longer valid. Certificates may be revoked because they were lost, stolen, or because an employee left the company, for example.

For small numbers of digital certificates, a single centralized certificate server will most likely suffice. But, if the number of certificates that the company requires is large, using multiple certificate servers arranged in some sort of hierarchy (perhaps based on departments) will be more manageable and more reliable because the system no longer has a single point of failure. Some certificate servers support multiple levels of administration; one group can perform certificate approval and revocation tasks, for instance, and another group can perform these functions as well as assign certificate authority to subordinate CAs. This makes it possible to set up distributed administration by assigning responsibility for a portion of the directory tree to another CA and set of administrators.

You also should be able to set certain parameters for the clients from your central system; these parameters should include defaults such as approved directory servers and certificate signers.

The task of supporting user access to certificates will become much easier if your system can support more than one method for requesting and receiving a certificate. At a minimum, clients should be able to perform these tasks using HTTP, e-mail, and disk files. As we've said before, smart cards are also becoming an increasingly attractive alternative for distributing and storing digital certificates.

Users should be informed of the need to properly store and protect their certificates. The previous sidebar on protecting clients against theft includes some suggestions for protecting their certificates.

Lastly, expect the client software to provide automatic checking of a certificate's validity using CRL

downloads (in the background, for offline use). When OCSP becomes available, look for software that supports it so that certificates can be immediately verified online.

Controlling Access Rights

Even though a VPN is architected to provide communications between hosts and security gateways, it's likely that you'll still want to maintain some control over the access that each VPN user has to network resources. For instance, if sales department personnel are not allowed access to R and D resources when they're on a hard-wired LAN, they should still be restricted from that access if they dial into the VPN while they're on the road. This means that you'll have to merge the control of the new access routes provided by your VPN with the access controls currently programmed into your routers and firewalls.

VPN traffic can be handled in two different ways by a firewall, either as unfiltered packets or as filtered packets (see Figure 13.6). In the unfiltered approach, the VPN traffic is handled the same way it is in a router; that is, the 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 VPN traffic before it is allowed into the internal network. Filtering VPN 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.

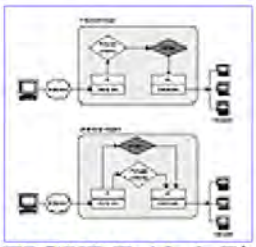

FIGURE 13.6 Firewalls filtering VPN.

If you place the gateway between the Internet and the router, then the router (and subsequent firewalls) can be used to filter both VPN and non-VPN traffic with the same rules; the gateway will provide transparent encryption and decryption services to the entire site. Also, the router doesn't need to be reconfigured to pass special tunnel traffic, which is the case when a gateway is installed behind the router. One caution: If the gateway is on the public, or untrusted, side of the network, you need to ensure that management of the gateway cannot be compromised from someone on the untrusted net. If this link is handling both VPN and non-VPN traffic, then the VPN gateway needs to be configured to pass non-VPN traffic.

When you locate the gateway behind a router or firewall, the control device would have to be configured to pass VPN traffic without filtering. Although this increases the security of the gateway (it's less susceptible to compromising the management port, for instance), it also means that you have less control over the traffic entering the LAN after decryption by the gateway. If you want to filter VPN traffic by destination, time-of-day, or application type, for instance, then you have to duplicate the filters from your router or firewall on the gateway.

Summary

Much of the management of security for VPNs is a straightforward extension of standard corporate security policies, especially for authentication of users and control of their access to network resources. However, VPNs do require added knowledge of the strengths and weaknesses of different encryption algorithms and associated key lengths so that the data transmitted on a VPN is properly protected.

The distribution of keys to authenticate security gateways and remote hosts on a VPN is an important part of VPN management, with many systems employing digital certificates for this task. Either commercial certificate authorities or a private in-house certificate server can be used to issue and control distribution he digital certificates.

Previous | Table of Contents | Next

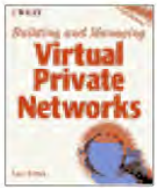

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 14
IP Address Management

The explosive growth in the use of IP for data communications, both within and among enterprises, has led to a number of problems in the allocation and management of IP addresses. Although the original 32–bit address space of IPv4 may have seemed sufficient to handle any network's requirements when it was first described, there is a growing concern that IPv4's address space will soon prove inadequate, at least for the public Internet. (Private internetworks are another matter, as we'll see shortly.) The next generation of IP, version 6 or IPv6, features a 128–bit address space, which should be suitable for some time, but until it's deployed globally, other short–term solutions to address shortages have been put into place. Unfortunately, although these short–term solutions help network managers deal with current day addressing and routing problems, they can cause problems for those of us deploying VPNs.

Although IPSec, as well as many other protocols, may be best suited for use with IPv6, most of us have to deal with the current situation surrounding the continued use of IPv4 and the added complexity that various short–term solutions bring with them. Because IPv4 is likely to stay around for the next few years, VPN design and deployment has to accommodate the complexities surrounding address management even as network engineers look for other solutions that will make addressing easier to use within VPNs.

To help point out some of the addressing problems VPN designers and managers face, this chapter covers the current methods for allocating addresses to network devices, both on public and private networks, as well as the related task of naming network entities via the Domain Name Service (DNS). As we go along, we'll point out some of the special problems VPNs may incur. Wherever possible we'll also discuss some of the current solutions proposed to counter these problems.

Address Allocation and Naming Services

For large enterprises, allocating IP addresses among thousands of workstations and servers and configuring these addresses in TCP/IP software is often a daunting task. In the past, adding, moving, or changing workstations and servers required manual assignment of new IP addresses. Simplistic approaches to tracking addresses, such as a notebook or electronic spreadsheet, may work for small networks, but these approaches quickly prove to be inadequate as networks get larger. Automated servers and related tools have to be employed to ensure that the networks run smoothly. Foremost among these for IP networks are Dynamic Host Control Protocol (DHCP) for address management and DNS for name management; and now, using Dynamic DNS to link the two makes network management easier, although

not foolproof.

A variety of problems can result from inadequate tracking of network addresses. Without proper tracking, addresses can be lost during equipment changes or moves just when network growth is leading to address scarcity. Not knowing which addresses are assigned can lead to mistakenly assigning the same address to two different machines, which leads to loss of connectivity and routing problems.

Another difficult task is allocating addresses for mobile users. Roaming sales reps with laptops may have to be provided with multiple network addresses, one for each router or remote access server that they might access via a dial connection leading to a waste of addresses and further tracking nightmares. Multiple addresses aren't needed when you convert the users of your remote access servers to a VPN, but you may still want to dynamically assign an address rather than use static allocations.

The current state of allocating addresses to companies also causes problems. Companies requiring addresses for more than 1,000 devices cannot obtain Class A or B network addresses and usually are forced to use *Classless Inter–Domain Routing* (CIDR) to combine available Class C addresses. But, using CIDR requires contiguous network numbers, leading to grouping networks by region so that all network numbers within a given region can be represented by a single entry in the routing tables of other regions (see Figure 14.1). If addresses for devices in a given region are not allocated contiguously, then routing table aggregation cannot be performed, and router performance will be reduced.

Static and Dynamic Address Allocation

In the past, an IP address was usually allocated by hand to a network device such as a router, server, or workstation when the device was attached to the network. (Printers and other devices using BOOTP are exceptions.) These addresses corresponded to the subnet in which the device was located—172.52.X.X for Human Resources versus 172.53.X.X for R and D, for instance—and had to be changed if the computer was relocated to another subnet. Furthermore, a device's address was static and didn't change unless someone (usually the network manager) changed the device's configuration file.

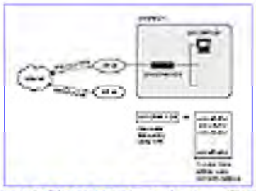

FIGURE 14.1 CIDR and routing table aggregation.

Allocating IPv4 Addresses

IP addresses are divided into three major classes: A, B, and C. (A fourth class, D, is reserved for special uses such as multicasting.) Each ad–dress consists of four octets, or sets of eight binary digits, separated by decimals. The first octet determines the class of the IP address. Class A addresses use the last three octets to specify IP nodes; Class B addresses use the last two octets for this purpose; and Class C addresses use the last octet.

Class A network addresses are the most desirable, because they are large enough to serve the needs of any size enterprise (see Table 14.1). But since fewer than 128 Class A networks can exist in the entire Internet, they are very scarce, and no more Class As are being allocated. Only those organizations that were early users of the Internet (e.g., Xerox Corp., Stanford U., BBN) are in possession of Class A network addresses.

TABLE 14.1 Properties of IPv4 Address Classes

Class	*Network ID*	*# Unique networks*	*Host address ID*	*# Unique hosts*
A	7 bits	128	24 bits	16,777,216
B	14 bits	>16,000	16 bits	65,536
C	21 bits	>2,000,000	8 bits	256

The more than 16,000 possible Class B networks also have become scarce and are now difficult to obtain. There is a large supply (over 2 million) of Class C network addresses, so they are still plentiful. The major problem is that for most organizations, a Class C network is too small (only 256 unique Host IDs). Even a Class B network is not large enough for an enterprise with more than a thousand LANs.

Previous | Table of Contents | Next

But, networks are far from static: Equipment gets changed or upgraded; people and equipment are moved; and networks rearchitected. Manually assigning static IP addresses is time–consuming when any changes are necessary; it also can be an error–prone process. To help deal with this continuing problem, a dynamic method for allocating addresses, the *Dynamic Host Control Protocol* (DHCP) was developed. And, since users normally are more comfortable with names rather than numeric addresses for network devices, the standard naming service, *Domain Name Service* (DNS), also was modified so that it could dynamically link with DHCP and track any changes made by DHCP.

DHCP is designed to provide a centralized approach to the configuration and maintenance of an IP address space, allowing the network administrator to configure various clients on the network from a single location. DHCP permits IP address leases to be dynamically assigned to workstations, eliminating the need for static IP address allocation by network and systems management staff. Pools of available IP addresses are maintained by DHCP servers.

DHCP operation is fairly straightforward. When a DHCP client workstation boots, it broadcasts a DHCP request asking for any DHCP server on the network to provide it with an IP address and configuration parameters. A DHCP server on the network that is authorized to configure this client will offer an IP address by sending a reply to the client. The client can either accept it or wait for additional offers from other servers on the network. Eventually, the client selects a particular offer notifying the proper server. The selected server then sends back an acknowledgment with the offered IP address and any other configuration parameters that the client might have requested.

DHCP servers aren't restricted to assigning only dynamic addresses. A set of addresses can be set aside as static network addresses for assignment to specific clients, such as file and mail servers. The DHCP server treats the lease periods for each of these static addresses as infinite.

The IP address offered to the client by a DHCP server has an associated lease time, which dictates how long the IP address is valid. During the lifetime of the lease, the client usually will ask the server to renew. If the client chooses not to renew or if the client machine is shut down, the lease will eventually expire, and the IP address can be given to another machine.

The *Domain Name Service* (DNS) is the Internet's official naming system and is designed to name various network resources, including IP addresses. DNS is a distributed naming system—the database that translates names to objects is scattered across many thousands of host computers.

Domain name requests (i.e., requests to convert a network name into its corresponding network address) are handled by a hierarchy of DNS servers (see Figure 14.2). Requests are sent first to the local (i.e., lowest level) nameserver in the network hierarchy, with the IP address of this nameserver typically configured in each workstation's TCP/IP software. If this nameserver cannot answer the query, it sends

the request to a higher level nameserver. This higher level nameserver can either resolve the name request itself or obtain information from a lower level nameserver that's unknown to the original requester. For example, the marketing and sales subdomains at Big Company may have nameservers at the same level in the DNS hierarchy, but the only way that users in marketing can obtain name information from the sales nameserver would be to request it via the bigcompany.com higher level nameserver.

In the past, DNS was designed to work with static IP addresses. A relatively new feature, Dynamic DNS (DDNS), has been defined by the IETF (RFC 2136) and now is provided by some vendors for their DHCP servers to automatically pass IP address lease–assignment information to DNS servers. This permits workstations with addresses assigned dynamically by DHCP to be tracked by DNS servers; workstations are then reachable by a name without manually maintaining the DNS database (see Figure 14.3).

Even though DHCP and Dynamic DNS can simplify IP address management, DHCP's dynamic allocation of IP addresses can cause other problems. Some firewalls and other Internet security products track IP addresses on the assumption that an IP address uniquely identifies a computer. If these products cannot map a DHCP–assigned address back to a specific user, an unauthorized user may gain access to the network because the address is authorized to go beyond the firewall, but the user is not.

Similarly, any attempt to debug problems on a live network relies on being able to translate an IP address to a particular user's computer. Other problems that can arise from dynamic IP address assignments might include control of content filtering (i.e., restricting Web browsing to certain sites), as well as billing and chargeback.

Dynamic address assignments can create problems for your security setup unless you're prepared for them. Because firewalls often match access rights to IP addresses, systems supporting DHCP should allow for the reservation of a batch of IP addresses for a specific group of users (a specific named team or department, for instance). As long as those same IP addresses are the ones allowing firewall traversal, use of the firewall can be controlled on a group basis, even when IP addresses are dynamically assigned.

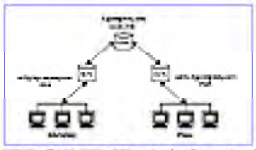

FIGURE 14.2 A hierarchy of DNS servers.

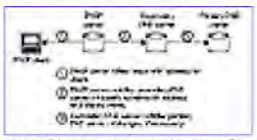

FIGURE 14.3 Coupling DHCP and dynamic DNS.

Although DHCP and DDNS are a nice fit, you can use DHCP without DDNS. In that case, not all devices on your net should have their addresses assigned dynamically. When assigning IP addresses to the file, mail, and other important servers on your net, static address assignment should be used. This makes it possible to use DNS to directly map network names to network addresses. Similarly, workstations that assume server functionality (e.g., personal Web servers) will normally also need static addresses so that they can be tracked by DNS.

Internal versus External DNS

When you're protecting external access to your intranet, say with a firewall or a security gateway, you have to take extra steps to protect your Domain Name Services while still allowing your users to access outside resources when necessary (and allowed). This usually involves setting up what is often called a double DNS scheme.

For a private IP network, your corporate DNS server would have sufficed, because it could take care of all address–name translations for the network, and the lack of a connection to the public Internet would help keep outsiders from discovering the names of corporate computing resources.

The first problem arises when you have a connection to the Internet and some corporate employees need access to resources on the outside. To properly map names to addresses, your corporate DNS server has to communicate with an external DNS server, presumably one hosted by your ISP. But, because you don't want outsiders to access your internal resources, you need to protect your internal DNS server (along with other network resources), so you install a firewall. Since the ISP's DNS server is outside the firewall, and your corporate DNS server is inside the firewall, they cannot readily communicate, which keeps employees from accessing outside resources as well as the reverse.

Previous | Table of Contents | Next

The solution is to install two corporate DNS servers: one on the outside of the firewall and one inside it. This is the double DNS scheme. The next step is to separate the hosts that had been on your sole DNS server into two groups. The first group lists those hosts that you want anyone on the Internet to find, such as your e–mail gateway, public Web site, and anonymous FTP server, for instance; it'll also include the name of the firewall's external interface. The second list contains the set of hosts that only your internal network users will be able to find. Don't forget that the second list also should include the external hosts that your internal users must be able to find.

As you might expect, the external DNS server stores the first list, and the internal DNS server stores the second list. The external DNS server is advertised to the Internet as the authoritative DNS server for your domain, which means that requests from Internet–based hosts will reach the external DNS server, but not the internal DNS server.

The hosts on your internal network use the internal DNS server as their primary DNS server. When they want to access external hosts, the internal DNS server will forward DNS resolution requests to the external DNS server outside the firewall. This is accomplished because the internal DNS server would be configured with a *forwarders entry* telling it where to find the external DNS server. Because the requests have to pass through the firewall, a DNS proxy service is set up on the firewall, allowing it to make a separate connection to the external DNS server on behalf of the internal DNS server (see Figure 14.4).

You may encounter similar situations with your VPN. If you're using an internal DNS server and shielding your DNS entries from the rest of the world, then you'll need a way to provide this information to the other sites and remote users of your VPN so that they can complete connections to appropriate resources. If you intend to allow access to a limited number of hosts on your VPN, then you could try maintaining dual DNS entries: one set for internal usage and the second for VPN use.

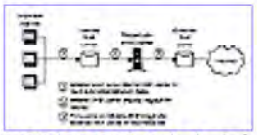

FIGURE 14.4 Linking internal and external DNS.

Private Addresses and NAT

The blocks of IP addresses allocated by the IANA are meant for use on the public Internet. If your company had no intention of using the Internet, but would transmit only IP traffic on its own internetwork, then any range of addresses can be used. Even then, the IETF recommends that only certain ranges be used so that Internet routers would not be confused if the addresses were inadvertently

advertised on the Internet. These blocks, which are defined in RFC 1597, "Address Allocation for Private Subnets," are as follows:

Class A 10.0.0.0–10.255.255.255

Class B 172.16.0.0–172.31.255.255

Class C 192.168.0.0–192.168.255.255

It's possible to use these private addresses for an internal internetwork and still connect to the Internet. To do so, you need to be allocated a block of registered addresses and use a firewall or router that performs *network address translation* (NAT).

NAT converts your inside addressing schemes into the registered addresses prior to forwarding the packets to the public Internet. The translation is fully compatible with standard routing functionality and features; NAT needs to be applied only on the router or firewall that is connected physically to both inside and outside addressing schemes.

NAT is interface independent, meaning that NAT can be applied to any interface on the router that links inside to outside addressing schemes. In Figure 14.5, the host system is using a privatized IP address of 10.2.2.1 as part of the intranet. When the packet reaches the router, NAT translates the 10.2.2.1 address into another address from the NAT IP pool allocated, say 171.69.89.2. It is as if that machine is virtually moved to the outside network segment for outside communication purposes. This network segment resides within the NAT router box itself for this example.

The NAT IP pool is considered part of the outside addressing scheme and not part of the inside addressing scheme.

Remember that NAT requires the capability to translate any part of the headers and packets that reference the addressing scheme. IP and TCP checksums need to be accessible, limiting the encryption of these areas. When the data is encrypted within the IP packets, it is impossible for NAT to perform the internal packet address translations. Thus, hosts using encryption should be assigned legally registered, outside addresses, exempted from NAT.

One significant disadvantage is the loss of end–to–end IP traceability. It becomes much harder to trace packets that undergo numerous packet changes over multiple NAT hops.

If an enterprise uses the private address space, then DNS clients outside of the enterprise should not see addresses in the private address space used by the enterprise, because these addresses would be ambiguous. One way to ensure this is to run two servers for each DNS zone containing both publicly and privately addressed hosts. One server would be visible from the public address space and would contain only the subset of the enterprise's addresses that were reachable using public addresses. The other server would be reachable only from the private network and would contain the full set of data, including the private addresses and whatever public addresses are reachable from the private network.

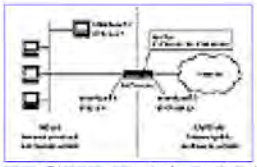

FIGURE 14.5 NAT at boundary router.

NAT configuration can become particularly complex for VPNs, so much so that various working groups within the IETF are still looking for the best solutions to typical uses of NAT and how they affect VPN design.

Multiple Links to the Internet

If you want to increase the reliability of your Internet connections for a VPN, one approach is to use redundant Internet connections (i.e., maintaining two, or more, connections) each served by a different ISP. But, redundant connections do pose problems of their own when configuring routers and firewalls.

Previous | Table of Contents | Next

The simplest method for supporting a second Internet connection is to connect both links to the same router and utilize the *Border Gateway Protocol* (BGP) on the router to decide which of the two ISPs should receive traffic (see Figure 14.6). This solution does not have the highest reliability, because the border router (as the one running BGP is called) can be a single point of failure. It's prefer able to maintain two separate routes to the ISPs, with separate routers and firewalls for each path, as shown in Figure 14.7.

Even this might not be a perfect solution, however. The main problem with this configuration is that most firewalls do not share information about their connections; if one connection point fails, the information about the sessions using it cannot generally be passed on to the second connection point so that it can pick up where the failed connection left off. Most security gateways behave similarly, although at least one product, Bay Networks' Contivity Extranet Switch, has a failover system that provides communications between multiple servers.

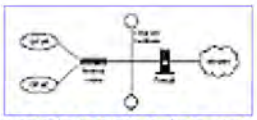

FIGURE 14.6 Multihomed connection to two ISPs.

FIGURE 14.7 Multihomed connection using multiple routers and BGP.

If the security gateways and firewalls are either simple packet filters or can share state, you can use two routers and firewalls to connect to the Internet—provided the internal hosts have registered IP addresses that can be advertised to the Internet. If you've used privatized addresses with NAT to connect to the Internet, this won't work.

IPv6

Although we've focused on the current version of IP, IPv4, throughout most of this book, we would be remiss if we didn't write a few words about the next generation of IP, IPv6.

The current IPv4 address size for a node is only 32 bits, providing for 4,294,967,296 addresses. Although that may have seemed like enough when the protocol was first created in 1978, we're starting to see saturation of the available address space. That's partly due to the class–related method for allocating blocks of addresses—assigning contiguous blocks of addresses for Class A, B, and C networks is easy to implement but does not efficiently distribute addresses, especially for small– and medium–sized

organizations. Some steps, such as *Classless Inter–Domain Routing* (CIDR), have been used to make address allocation more efficient, but these are stop–gap measures that don't address the crucial issue of the size of the address space itself.

But, IPv6 promises to solve that problem. The first notable difference between IPv6 and IPv4 is the length of its address field—it's 128 bits, or four times as long as that found in IPv4. In addition, IPv6 includes built–in support for such options as multicast support, IPSec, and flow control for quality of service, which have all been tested in IPv4 but had to be tacked on in a less–than–optimal fashion. Also, although the IPv6 header is larger than IPv4's, it has fewer fields, which should make routing more efficient as routers will have to do less processing per header (see Figure 14.8).

IPv4's 32–bit addresses are subdivided into four 8–bit groupings called octets, which are then expressed in what's known as the dot notation (i.e., 252.123.345.004). The designers of IPv6 have chosen a similar format composed of eight 16–bit integers separated by colons. Each integer is represented by four hexadecimal digits, as in FEDC:BA98:7654:3210:FEDE:BA98:7655:2130. Some IPv6 addresses can be obtained by prepending 96 zero bits to an IPv4 address. These IPv4–compatible addresses, as they're called, are important if you're planning to tunnel IPv6 packets through an IPv4 network, because the prepended zeros can be easily added to, or removed from, the IPv4 address.

FIGURE 14.8 IPv4 and IPv6 packets.

The larger IPv6 addresses affect just about every part of your network; not only will you have to upgrade IP stacks for client and host computers, but you'll also have to upgrade your DNS servers and routers. DNS servers simply have to be refitted with software that can handle the larger IP addresses, which is a straightforward extension of DNS. But, switching to IPv6 will enable you to create a global addressing scheme for all your VPN sites without resorting to NAT and, therefore, reduce the need for extra reconfiguration of firewalls and DNS servers.

Summary

IP address allocation to networked devices within a company can be a painstaking, time–consuming, and error–prone task if handled manually. One solution to this problem is to utilize dynamic IP address allocation via DHCP. Because address–to–name mapping is an integral of any IP network, it's also necessary to link DNS to DHCP; this is now accomplished via DDNS, or Dynamic DNS. Special DNS configurations using multiple servers are needed if a firewall is used to separate the private corporate network from the rest of the Internet.

Allocation procedures for public IP addresses, along with the limited number of addresses defined in IPv4, have made it necessary to adopt a variety of solutions to simplify routing on the Internet and to use IP addresses on corporate networks. *Network Address Translation* (NAT) has proven to be a popular solution for enterprises wanting to keep a private IP address space for their intranet while still maintaining some connectivity with the public Internet. Unfortunately, NAT also can make it difficult to

easily build VPNs.

Previous | Table of Contents | Next

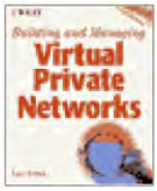

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 15
Performance Management

It's an unwritten law of networking that, like nature and a vacuum, users abhor unused bandwidth and fill it quickly. The resulting network congestion can wreak corporate havoc, preventing high-priority traffic from getting through, frustrating users, and overwhelming network devices.

Furthermore, by combining many different types of traffic, today's multiservice networks, which may handle messaging, transactions, video, telephony, and more, are making it more difficult to allocate bandwidth and control the network.

Network performance and VPNs are inextricably interlinked. If VPN tunnels are to appear transparent to users and applications so that all sites appear as one large enterprise network, then the tunnels cannot act as performance bottlenecks. At the same time, these links may well not have the same bandwidth as that found on each site's LAN, so some care has to be exercised to ensure proper performance between VPN sites.

Since security gateways for a VPN often link two disparate bandwidth domains—that of the LAN and the usually significantly slower WAN—they are chokepoints for the flow of network traffic and can serve as ideal locations for controlling traffic based on application or user priorities. With this in mind, some vendors already have included support for traffic prioritization and bandwidth management within their VPN products; two noteworthy examples are Bay Networks' Contivity Extranet Switches and Check Point Software's Firewall-1.

This chapter covers the basics of network performance and related application requirements as well as methods for offering network services to your customers that can be differentiated on the basis of those application and/or user requirements. Then we'll talk about how policy-based network management can be used to help maintain control over network configurations and bandwidth control. Finally, the chapter discusses the role your ISP plays in supporting your VPN's performance.

Network Performance

Let's investigate the components of network performance before we move on to discuss how it can be managed.

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.

A network can contribute to latency in a number of ways:

1. *Propagation delay*. The length of time it takes information to travel the distance of the line. This type of delay is mostly determined by the speed of light and isn't affected by the networking technology in use.
2. *Transmission delay*. The length of time it takes to send the packet across a given medium. Transmission delay is determined by the speed of light and the size of the packet.
3. *Processing delay*. The time required by a router for looking up routes, changing the header, and other switching tasks.

Another factor, that of jitter, also affects real-time network traffic. Jitter is the variation in the latency. Irregular packet delays due to jitter can introduce distortion, making the multimedia signal unacceptable.

If we take a look at the best-effort delivery offered by IP, we see that IP networks treat every packet independently; a source may transmit a packet to a destination without any prior negotiation or communication. Furthermore, the network has no information that a particular packet belongs to a suite of a packets, such as a file transfer or a video stream. The network will do its best to deliver each of these packets independently. This approach often introduces considerable latency and jitter in end-to-end paths, which aren't compatible with much of the data generated by the newer applications seen on networks that depend on known delays and little, if any, data loss. But, that's unsuitable for real-time applications, such as interactive multimedia, which often cannot tolerate retransmitted packets or indeterminate delays.

Requirements of Real-Time Applications

A wide variety of applications can run on networks. In addition to the bulk transfer applications like ftp, netnews, and e-mail, there are interactive applications ranging from a terminal emulator that requires entering commands to control responses from a remote host or using a Web browser to view pages on another site to interactive simulations between players in a multiplayer network and the even faster interactions required for transaction processing of online orders.

In the past, network managers could predict fairly well what the traffic patterns of their networks would be, because there were only a limited number of servers, legacy databases, and other network resources that most users accessed. But, that's changed considerably over the past few years as the World Wide Web and collaborative applications have changed interactions between users, both within, and among, subnets of an internetwork. At the same time, other new applications, utilizing streaming multimedia, videoconferencing, and so on, have increased the traffic on networks.

Traffic flowing across integrated enterprise networks can be grouped into three basic categories: real-time traffic, interactive traffic, and bulk transfer traffic (see Figure 15.1).

Real-time traffic, such as conversational voice, video conferencing, and real-time multimedia, requires very short latency and controlled jitter. Once minimum bandwidth requirements are met, higher available bandwidth can bring increased quality if the applications are designed to use it.

Interactive traffic, such as transaction processing, remote data entry, and some legacy protocols (e.g.,

SNA), requires latencies of approximately one second or less. Greater latencies cause processing delays as the users must wait for replies to their messages before they can continue their work. Interactive traffic is not sensitive to bandwidth beyond that needed to satisfy their latency requirements.

Bulk transfer traffic accepts virtually any network latency, including latencies on the order of a few seconds; it is more sensitive to the available bandwidth than to the latency. Increased bandwidth can result in sharply decreased transfer times; virtually all bulk transfer applications are designed to use all available bandwidth.

Previous | Table of Contents | Next

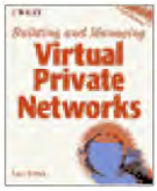

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

With the move to interactive multimedia, applications now require control over the QoS 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 it relative to other QoS requests from other applications. QoS provides network services that are differentiated by their bandwidth, latency, jitter, and error rates.

The increased use of multimedia is not the only reason to differentiate services and control traffic on your network. Some of the traffic flowing on your network is more critical to the running of your business than others; this traffic must go through, even if it means throttling back other, less essential traffic. It thus becomes important to be able to differentiate classes of network traffic and to have a system for dealing with these classes in different ways.

FIGURE 15.1 Networked application categories.

Supporting Differentiated Services

As you might expect, there's more than one approach to providing differentiating services to help deal with network congestion. The five commonly proposed techniques are as follows:

1. Over-provisioning network bandwidth.
2. Bandwidth conservation.
3. Traffic prioritization.
4. Static resource reservation.
5. Dynamic resource reservation.

Over-provisioning isn't exactly a method for differentiating services, but it can help deal with network congestion by allowing the network to handle a larger traffic volume. It's a reasonable solution for local and campus LANs. But, for WANs and VPNs, over-provisioning the bandwidth may not be a viable solution due to the high cost of the added bandwidth. Because there's often a noticeable mismatch in bandwidth (perhaps 100 to 1, or greater) at the LAN-WAN boundary (and hence the LAN-VPN tunnel boundary), some form of traffic management or control is needed.

Bandwidth-conservation techniques improve overall network performance by trying to ensure the most

efficient use of the available network capacity rather than differentiating services. Some of the current conservation techniques are IP multicasting, data compression, and bandwidth-on-demand.

IP multicasting reduces the total amount of traffic on a network by eliminating the forwarding of redundant traffic. (For more details, see *IP Multicasting: The Complete Guide to Interactive Corporate Networks* by John Wiley and Sons, Inc., 1998.) The second technique, *bandwidth compression*, can be accomplished in routers to reduce bandwidth demands for a WAN link by a factor of two to four. Lastly, *bandwidth-on-demand* (BOD) can be used to provide additional bandwidth as needed by using additional analog or digital phone lines when the WAN interface becomes congested.

Each of these approaches may be usable in VPNs, but their applicability depends on the nature of your application demands. IP multicasting only works well when the same data is being transmitted to a number of receivers. If each session travelling across a VPN tunnel is between a different client and a different server, multicasting is of little help. Bandwidth compression may prove more useful, and some vendors (VPNet, for example) already have included it in their security gateways for processing IPSec traffic. But, the bandwidth savings may not be sufficient to meet your needs, and it does not address any latency problems. BOD also can prove useful by providing more bandwidth as needed, but the added links may cause configuration problems on VPNs or may not be able to inexpensively provide sufficient bandwidth (or latency) for your needs.

Traffic prioritization, or *Class of Service* (CoS), is a simple but useful tool for providing differentiated services. Routers can differentiate between service classes according to the precedence field in the header of each packet (IPv4's Type of Service, or TOS, field). This method offers a small fixed number of service classes and only guarantees that packets with higher precedence get better service than packets with lower precedence. Since there is no admission control, there is no mechanism to prevent classes from becoming overloaded.

To improve on CoS support, the major networking product vendors, like Cisco and 3Com, program admission control at edge routers, (i.e., routers that interface between a LAN, such as a branch office's LAN and a core network, such as the Internet or the main corporate network). These edge routers use preset policies, or rules, to assign traffic to classes before the traffic is forwarded to the core network (see Figure 15.2). The routers in the core network use one of a variety of algorithms to process the traffic classes, each of which has its own queue. A common algorithm for processing the queues is called *Weighted Fair Queueing* (WFQ), which prohibits large flows of packets from consuming large amounts of bandwidth, which could keep smaller flows from being transmitted. Because CoS is implemented at edge routers, those same routers can be the security gateways for your VPN, enabling you to combine VPN control and traffic control at the same point.

But, if traffic prioritization using classes is insufficient for your needs, and you choose to allocate network resources between real-time and non-real-time applications, then you have two choices. Either you can statically allocate the resources or you can allow resources to be reserved dynamically.

Static resource allocation enables you to reserve a portion of a network's capacity for a particular type of traffic, usually based on protocol, application, or user. In many enterprise networks, routers are often configured to devote a certain amount of their capacity to SNA traffic, for instance, to accommodate the requirements of legacy data transactions. When the capacity is reserved for a specific protocol or application, the capacity should be large enough to meet the demands of all traffic of that type. If not, the traffic exceeding the allotted capacity will most likely be subject to delays and/or discards. If the allotted

capacity isn't used, it's possible for other traffic to use the remaining bandwidth.

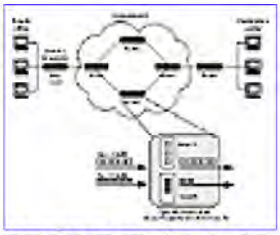

FIGURE 15.2 How class of service works.

Previous | Table of Contents | Next

VPN Performance

When it comes to VPNs, two major factors affect performance:

1. The speed and reliability of transmissions over the Internet.
2. The efficiency of VPN processing at hosts and security gateways.

As we've already discussed, the Internet cannot be used to provide guaranteed response times (i.e., guaranteed latencies and jitter). ISPs offering guaranteed latencies do so by circumventing the public Internet and channelling customer traffic over their own backbone network. This works for VPNs as long as all your sites can be served by the same ISP. As we pointed out back in Chapter 9, "The ISP Connection," no ISP yet offers guaranteed latencies for traffic that travels across more than one ISP's network, although it's likely that the technology and policies to do so will exist in a few years.

For all that we said earlier in this chapter on differentiated services and QoS, most ISPs are not yet prepared to offer support for these technologies. Network product vendors are pushing hard to make the requisite hardware and software available for ISPs, but few have adopted any of the technologies needed to offer customers differentiated services. The notable exceptions include MCI, UUNET, and TCG CerfNet.

Aside from the cash outlay required to purchase and install the devices, a lack of standardization for differentiated services also has contributed to ISP reluctance to adopt the techniques we've outlined. The current thinking is that RSVP will most likely not be implemented across most ISP's backbone networks and the public Internet, partly due to its scalability problems. A more likely solution for offering differentiated services is the Class of Service approach, especially since it appears that the approaches tried by different vendors, particularly Cisco and 3Com, will soon be able to interoperate.

What QoS-related services your ISP offers will have the greatest effect on your company's time-sensitive applications. If all your VPN traffic is going to be file transfers, Web browsing, and e-mail, then you won't need to be concerned with QoS. But, if transactional traffic, interactive multimedia, and IP telephony are going to a part of your VPN's traffic, then you'll need to track developments in QoS technologies and your ISP's deployment of them.

But, there's more to managing the performance of your VPN than utilizing QoS. As we mentioned at the beginning of this section, efficient VPN processing is an important factor. You do not want your security gateways to be chokepoints to network traffic due to their inability to efficiently encrypt and decrypt packets.

Depending on the computational horsepower of your VPN devices and the traffic they must process, you

may have to consider installing multiple gateways at connections experiencing heavy traffic and enabling some form of load balancing between the gateways.

Since security gateways are such strategic points for creating VPN tunnels, you should plan on using them as locations for controlling the traffic that enters the VPN links. For instance, if a gateway can utilize filtering rules based on time of day and application (or user), then you could set up filters to ensure that business-critical traffic is passed with a higher priority during business hours, and Web browsing might have an equal priority later in the day. Of course, setting up and managing all these rules can be a headache as you try to enforce them across numerous VPN sites; as we'll see in the next section, policy-based network management offers a potential solution.

MPLS and ISPs

Another approach, *Multi-Protocol Label Switching* (MPLS), which tags IP traffic so that it can be moved efficiently over switched infrastructures such as ATM is being standardized by the IETF and is initially being offered within Ascend's products for ISPs as part of their MultiVPN product line. Because most of MPLS deployment is aimed at the ISP's backbone network, there's very little an enterprise customer would have to do to use MPLS. Only a few enterprise routers support MPLS, but expect more to become available as the protocol becomes a standard sometime in late 1998.

To make the most of the bandwidth provided by your ISP, you should pay close attention to how your tunnels are created and what traffic they carry. For instance, Layer2 tunnels using L2TP can carry traffic from more than one session. But, if multiple sessions are inserted into a tunnel, it's possible that a higher priority packet can be placed in the channel first, which may disrupt any sequence-sensitive processing of packets, such as header compression. Despite the convenience multisession tunnels may offer, it's best to either restrict tunnels to single sessions or at least insert only sessions of equal priority into the same tunnel.

Aside from affecting how traffic is aggregated, the tunnels your VPN creates also can have an effect on the deployment of any QoS schemes your ISP offers. Simple approaches to tunneling like GRE (Generic Routing Encapsulation), which is used in PPTP, usually map any QoS fields of packet headers to QoS fields in the header of the tunnel packets. But, if tunnel-mode IPSec is used, the original header is encrypted; if the security gateway cannot translate QoS information from the internal host's requests in the inner header to the outer header it generates for the tunnel, then the ISP's network will not be able to provide any quality-of-service support for the tunnel's traffic.

Previous | Table of Contents | Next

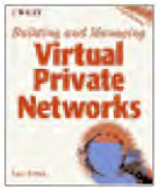

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Policy-Based Management

In past chapters, we've talked about policies in terms of security policies, covering such issues as user authentication and access rights. In the context of this chapter, however, policy-based management has a different meaning—it's using stored rules to manage bandwidth and determine what users get the quality-of-service they require. But, since in the long run all policies are focused on the user or a device, the grand vision of policy-based network management is to use a single management database (distributed or otherwise) to control all aspects of the network, including security, access rights, bandwidth requirements, and so on.

Network management has become more complicated not only as networks scale to larger sizes, but also as they become more complex, with more services coming online and application demands becoming more varied. What's needed is a better way to manage network traffic, setting priorities and bandwidth requirements in a centralized way even as the network itself becomes more distributed and decentralized.

As a form of network management, the major networking vendors, such as Cisco, Bay Networks, and 3Com, have been developing what's called policy-based network management. To help deal with the complexity of their networks, network managers can use policy-based management to implement policies that explicitly address the needs of the ever-expanding range of services.

Policy-based management has come to the fore as switches have become more important in enterprise networks. As switching becomes an integral part of enterprise networks, often displacing routers, users and managers alike are looking for ways to optimize their use of switches, especially when it comes to controlling and distributing network resources. Many networks based on a core or hierarchy of routers aren't capable of prioritizing network traffic based on either user or application priority. Similarly, ATM-based networks can offer *quality-of-service* (QoS) guarantees, but few applications have been developed to take advantage of these QoS requests at the workstation level. And RSVP, which has been developed to provide similar QoS capabilities to IP-based networks, is relatively new and hasn't yet seen wide-spread usage in networks. But, policy-based management offers the promise of working with a variety of network devices to enforce bandwidth management and admission policies for application traffic along the entire path between the source and the destination.

A fundamental tenet of policy-based management is that policies for governing network behavior are set at a high level by network managers, and intelligent network devices use these policies to adapt to network conditions (see Figure 15.4).

It's important that policies for handling priority requests are set in a centralized fashion, usually at a network-wide policy server. Thus, a network manager would set policies to determine which users and

applications get the top priority when congestion slows network performance. When set, these policies can be automatically invoked by the workstations, switches, and other network devices as conditions change in the network.

As an example, see Figure 15.5; here, a network manager established priorities for applications on a user-by-user basis. These priorities, which are stored in a central policy server, are relayed to each appropriate user when the user's workstation starts up and connects to the network. Then, when a particular user launches a particular network application on his workstation, the data packets sent to the network are tagged with the appropriate priority and relayed by the switches according to their priority.

FIGURE 15.4 Basic model of policy-based network management.

This two-dimensional matrix of priorities, sorted by user and by application, enables different users of the network to have different priorities for the same application. This way, it's possible to assign a high priority for the CEO using a SAP application while someone in technical support would have a lower priority for the same SAP application. Similarly, all uses of PointCast or similar push applications can be assigned a lower priority than the use of SAP applications. After the database on the policy server is established, network switches and routers can automatically handle high-priority traffic at the expense of lower priority traffic in the event of network congestion.

Although companies like 3Com, Bay Networks, and Cisco already have developed the first generation of proprietary policy management tools, it's likely that the interoperability and capabilities will be improved over the next few years thanks to an industry effort called the *DEN Initiative* (Directory-Enable Networking). This work, originally started by Cisco and Microsoft, now has more than 20 participating vendors and is defining ways to use directories to store both user profiles and device configuration information. Although much of the initial work is focusing on using Microsoft's Active Directory, which is a part of NT Server 5.0, directories and devices will be able to query each other and exchange information using LDAP. Policy management software then can be used to set rules and store them in a DEN directory; network devices can then automatically make decisions about bandwidth and resource allocation based on the rules propagated from the policy software and user profiles stored in a DEN directory. The first products to support DEN are slated to ship in early 1999.

FIGURE 15.5 Example of policy-based network management.

Previous | Table of Contents | Next

Because the trend in policy-based network management and DEN is to collect all network- and user-related information into one system for centralized management, both the configuration of VPN devices and traffic control of the LAN-WAN links will eventually be integrated into such systems. As we've mentioned earlier in this chapter, some VPN products are already shipping with LDAP capabilities, which will make their integration in policy-based management systems easier. But, since both policy-based management systems and DEN are relatively recent efforts, it'll still be a few years before widespread deployment is possible.

Monitoring ISP Performance and SLAs

We covered many of the details of ISP capabilities and Service Level Agreements in Chapter 9, "The ISP Connection." Remember that SLAs should be used to agree on what are reasonable expectations of service. Three basic items should be covered in every SLA: availability, effective throughput, and delay.

Monitoring your ISP's performance should be done not only to ensure that the conditions of your SLA are being met, but to determine how your VPN is behaving. For instance, if VPN traffic isn't getting through or is being delayed because of congestion at a security gateway—not because of ISP performance—then you might have to consider installing a more powerful gateway or balancing the load between multiple gateways. Alternatively, if some links aren't being heavily used, you may want to renegotiate a slower speed for those links.

Although SLAs may be based on the three items we mentioned earlier—availability, throughput, and delay—your users are going to be most concerned with the performance of their applications over the network. Always keep in mind that your performance measurements should in some way be related to actual user actions, such as the time to download a file or send a message.

Recall from Chapter 9 that where you take your measurements can have an impact on the results you get. Measurements can either be taken end-to-end or just within the ISP's network cloud (see Figure 15.6). 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.

A second issue is utilizing a measurement system that is independent of the network you are measuring. Use an objective system that is not biased toward either switch or router architectures.

Many monitoring tools collect and report on data from SNMP agents. SNMP agents perform the function of accumulating real-time data, and this approach works well for bandwidth-related measurements. Most routers and other network devices are available with SNMP agents that provide most of the information

needed for monitoring availability and utilization.

Other monitoring systems poll devices using specific application protocols such as FTP and HTTP or network-level polling with *Internet Control Message Protocol* (ICMP). But, polling systems can include factors that are beyond the scope (and therefore, control) of your service provider. (An overloaded Web server at a corporate site isn't the responsibility of your ISP, for instance.) A good approach would be to employ ICMP polling and to place the polling device as close as possible to the service being measured (i.e., the LAN-WAN interface in this case).

FIGURE 15.6 Measurement areas for SLAs.

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, enabling customers direct access to the data that relates to their VPN.

Summary

A variety of network applications have different requirements for bandwidth, latency, and jitter, complicating the planning of bandwidth provisioning and traffic control. Many of the newer applications, such as interactive multimedia and videoconferencing, place tighter constraints on network latency and jitter than most legacy applications.

Networks can support this range of applications if they're configured for differentiated services. The five approaches to offering differentiated services are over-provisioning bandwidth, bandwidth conservation, traffic prioritization (or Class of Service), static resource reservation, and dynamic resource reservation.

Because the important components of a VPN are located at the LAN-WAN interface, they not only can become chokepoints for network traffic, but also offer the opportunity for controlling traffic and differentiated services. But, whenever ISPs offer their own support for differentiated services, special attention must be paid to mixing traffic of different priorities in the same tunnel or encrypting packet headers, which defeats most prioritization schemes.

Policy-based network management is a rapidly-developing area that promises to make device configuration and automatic control of traffic easier. Eventually, VPN configuration and user control will be included in policy-based network management systems.

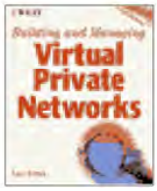

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

PART V
Looking Ahead

VPNs are useful now and have a great deal of potential to be even more useful in the future. Standardization is just now happening, which will improve interoperability and management. Network performance over VPNs will also improve, enabling VPN links to be used for new applications, such as videoconferencing and IP telephony.

To track product interoperability, you should look at the International Computer Security Assn. (ICSA) for their tests of compliance with IPSec standards and at the Automotive Network Exchange (ANX) for valuable information on how the products work together in the real world.

CHAPTER 16
Extending VPNs to Extranets

The Internet and other TCP/IP networks have been around for more than 20 years. But, it's only been in the last few years that the Internet has become a household word and more businesses are paying attention to using TCP/IP and the Web for all types of communications; this includes not only communications within the enterprise, but with customers, suppliers, and business partners. The appeal of using the same protocols and in many cases the same application (a Web browser, for instance) to perform many different tasks and links to many different companies is very real and hard to pass up.

Within the business world, the biggest trend in IP networks (and perhaps the most profitable one thus far) has been to redesign corporate communications around the World Wide Web and intranets. Electronic commerce, particularly using the Internet to buy and sell goods and services, has a lot of potential, and efforts for consumer e-commerce versus business-to-business e-commerce have been following somewhat different paths of development. One of the promising efforts in business-to-business e-commerce has been dubbed an extranet, which involves opening up portions of your corporate network to access by your business partners (see Figure 16.1).

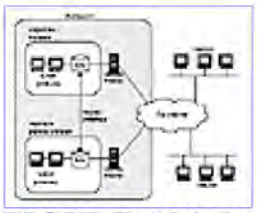

FIGURE 16.1 Intranets, extranets, and VPNs.

As we'll see in this chapter, extranets can require a great deal of coordination between businesses, perhaps more than has ever been attempted previously. And, because you're trying to control outside access to your resources and probably want to secure traffic between you and your partners, you probably already see how important security is to the proper operation of an extranet—that's where the parallel to VPNs come in. If you need the secure transmissions in addition to control of access rights for outsiders, then VPNs can be a good foundation for your extranet.

One difference between extranets and VPNs has been the focus in their evolution. Extranets are motivated more by the need for a particular business application—faster processing of purchase orders or better inventory control, for instance—and VPNs have grown out of the need to provide secure communications over the public Internet, regardless of the application. Because of this capability of VPNs, the applications you plan for your extranet can mesh nicely with the architecture of a VPN; extranet applications can be layered atop the VPN plumbing, as in Figure 16.2.

You don't have to use a VPN to create an extranet; that decision depends on the security requirements of your extranet applications. You could use SSL/TSL to secure communications between a partner's Web browser and a Web server maintained just for your extranet, for instance. Or, trading EDI forms via secure e-mail (using S/MIME, for instance) may be enough for your needs. But, the focus of this chapter will be how you can extend your VPN to become an extranet.

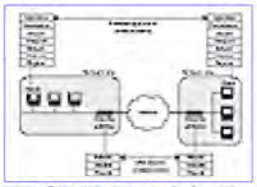

FIGURE 16.2 Extranet applications and VPN networks.

Reasons for an Extranet

Before we discuss some of the details that go into creating an extranet from a VPN, let's spend a few pages delving into extranets.

For many managers, extranets offer many advantages for communications between business partners. First, extranets are usually built using TCP/IP protocols, which means that the difficulty of linking the networks of two (or more) companies is reduced. Furthermore, since the public Internet also uses TCP/IP, partner networks can be linked to each other using the Internet instead of installing expensive leased lines or other links.

Second, using the Internet to link networks gives you more flexibility in forming and dissolving short-term partnerships as needed, which has become increasingly important in today's fast-paced business world. Sometimes you can't wait a month or more for the installation of a leased line, for

instance. Or, a collaborative project between you and another company may involve only a small group of people, which would make the cost of a leased line prohibitive.

Third, many extranets revolve around the use of the World Wide Web, which helps provide a common user interface to many applications across company boundaries. The use of Web browsers has gotten to be pretty pervasive throughout many businesses, and companies are expending a great deal of effort developing applications that use the Web. This not only simplifies the distribution of client software, but also makes access to a wide variety of data easier than in the past with legacy applications.

Some of the arguments for extranets are the same as for VPNs; look back at the discussion in Chapter 1, for instance. But, the business case may be a little different because we're talking about external business communications in extranets rather than the internal communications that VPNs support. We'll shortly see in this chapter what a focus on external partners adds to planning.

Business reasons for an extranet vary, but communications with partners is at the heart of each extranet. It's a question of what kind of data you want to obtain or share: it could be inventory levels, the status of purchase orders and shipments, market data, product information, or just about any kind of business data imaginable.

One of the most popular current uses for extranets is managing the supply chain (see Figure 16.3). The idea is to tie together all the companies involved in your business: suppliers of parts, servicers of equipment, the manufacturer, and distributors, among others. Automating many of the steps in this supply chain across company boundaries can lead to faster order processing, improved inventory tracking and management, more accurate order fulfillment, support for just-in-time manufacturing, and improved customer support.

Other extranets may not be as complex; you might just want to obtain daily point-of-sale data from your distributors, for instance, or provide product information and corporate news updates to them via a Web server.

Large companies like Ace Hardware have used an extranet to provide their network of independently owned retailers access to information that previously had been stored on legacy mainframes and was difficult to access from the outside. In Ace's case, the newly accessible information on the extranet included inventory levels at the warehouse closest to the retailer, inventory-management applications to help plan reorders, and margin management and pricing tools that help store owners maintain profitability.

FIGURE 16.3 Components of a supply chain system.

Another area that's received a lot of attention is the purchasing process. Two different approaches to the process are worth noting: the use of on-line catalogs and the use of *Electronic Data Interchange* (EDI).

Previous | Table of Contents | Next

On-line catalogs have become a standard part of electronic commerce. Within the context of business-to-business e-commerce, suppliers have started to offer customized on-line catalogs for their customers; these catalogs can be based on a customer's past purchasing history or just the business type. Being electronic, they're easier to update and customize. If the catalogs are offered over an extranet, access to the proper catalogs can be easier to control.

EDI has been around since the 1960s, but it's been used mostly by large corporations and their satellite suppliers working together over a private network called a *Value Added Network* (VAN). These VANs offered reliability and security that has been difficult to duplicate on the Internet thus far. EDI data is presented in forms that are defined for a particular type of business, relayed between business parties using e-mail, and then translated into formats that the company's databases can use.

Businesses can use EDI to automate the transfer of information between corporate departments, as well as between companies. For instance, EDI-based data can be transferred between purchasing, finance, and receiving departments (see Figure 16.4) to automate the purchase-and-payment process.

Because the cost of operations on the Internet is lower than on a VAN, interest in using the Internet to transmit EDI data has increased. Some EDI vendors now offer Web-based servers that accept business data in HTML forms and translate the data to EDI formats for transmission either over a VAN or the Internet. Meanwhile, the IETF has been working on a standard for enclosing EDI form data within S/MIME messages. In addition to using EDI over the Internet, another, newer effort by the *World Wide Web Consortium* (W3C) to supplant EDI forms with the *eXtended Markup Language* (XML) will make it easier for companies to relay structured data like purchase orders within IP-based e-mail. Both EDI and XML in IP-based e-mail and on the Web will make it easier for extranet partners to exchange the information they need to tie their businesses together electronically.

FIGURE 16.4 EDI information flow between a buyer and a seller.

Probably the largest combination of VPNs and extranets is the *Automotive Network Exchange* (ANX), which soon will be an extranet linking 8,000 suppliers and 20,000 dealers (see Figure 16.5). Organized and managed by the *Automotive Industry Action Group* (AIAG), this extranet includes certification of IP service providers as well specifications for the capabilities of each ANX member that must be met before it can become part of the extranet. The extranet relies on IPSec for communications over the network and uses EDI for automating commercial transactions between the members.

Turning a VPN into an Extranet

For the remainder of this discussion, we're going to assume that you have deployed a VPN and that you want to use the security features that it offers to create an extranet.

FIGURE 16.5 The Automotive Network Exchange.

The main difference between establishing your VPN and establishing an extranet is obtaining the cooperation of your business partners. Even if the extranet is your company's idea, and the applications developed for the extranet will benefit your partners, you'll still need their buy-in to make the extranet a success. For some companies, this may be the same situation they faced when getting a buy-in for a VPN from branch offices. Ah, politics...

We're not going to get into the details of planning your extranet applications or the development environments you can use for those applications. That could easily be the subject of another book. Because we're mainly concerned about how the extranet can use your VPN's features, we'll assume that the reasons for your extranet and the type of applications you plan to use are already settled.

Most of the concerns you'll face in linking partners to your extranet revolve around compatibility. The five main areas of compatibility are network setup, security policies, authentication servers, VPN protocols, and digital certificates. These compatibility issues may not all be of equal importance, and they will also differ in importance if you're setting up a dial-up extranet versus planning to link corporate LANs together. For instance, providing IP services and client software to your partners is much simpler if they'll be using dial-up connections than if IP LANs need to be installed or reconfigured and security gateways installed.

Regarding network setup, you'll need to know if they have the appropriate IP routers and Internet links to join your extranet, at least if you're establishing LAN-to-LAN links. Because you'll be extending VPN tunnels to their sites, many of the issues we discussed in Chapters 8, "Designing Your VPN," and 13, "Security Management," are pertinent. Does the partner's site have the proper equipment for a security gateway, or does one need to be installed? Or, if the partner already has his own VPN, can that VPN's devices be linked to yours? If you and your partners are all using private IP addresses on your intranets, how will you handle address translation and name services between companies? These are just some of the major issues with which you'll have to deal.

Dial-up extranets are perhaps a bit easier to set up when it comes to network compatibility. Even if your partner doesn't use IP as its main networking protocol, setting up a few desktop or laptop computers with remote access software, modems, and an ISP account is simpler than installing and configuring an IP network and a security gateway. It's also less expensive.

We've said a lot about security policies throughout this book. (See Chapter 13 in particular.) If you're extending your VPN to extranet partners, some agreement on security policies between you and your

partners will be necessary. Whatever you do, try not to compromise your VPN's security or your internal site security to meet the wishes of your partners. If any compromise has to be made, it should favor stronger security. This may take some training on your part, because smaller companies may not have a security policy or they may not understand the value of protecting your resources. In particular, if you're going to issue passwords, security tokens, or digital certificates, be sure to impress on your partners the need to protect them against loss or theft. The question of who's legally liable for loss of data due to a lost password or security token is no doubt an important issue here, but we'll leave the coverage of legal issues to someone else.

The different partners in your extranet also may have different authentication schemes for their employees. Assuming that there's a two-way exchange of data between members of the extranet, the authentication systems will have to be able to handle some outside users (i.e., some of the employees of your partners, for instance). If different partners use different authentication servers, every system that accesses the extranet will require client software, perhaps more than one. Getting your partners to agree on one common method for authentication would simplify the configuration and use of the extranet, as well as reduce support costs in the long-run (especially for your help desk). Using RADIUS, perhaps with proxy servers, is one way to go, because many VPN systems work with RADIUS, and many companies are already adopting it for controlling remote access. Extensions to RADIUS also enable it to work with other authentication systems, such as SecurID tokens, which help make RADIUS an integrator of authentication systems.

Previous | Table of Contents | Next

It's also possible that you'd use digital certificates to authenticate extranets users as well as sites, much the way we described their use for VPNs in Chapter 13, "Security Management." If you do use digital certificates, check for compatibility of the certificates—not all certificates use the X.509v3 standard. The popular encryption program for e-mail, PGP, has its own certificate format that's not compatible with X.509v3, for instance. You'll also need to determine who will issue the certificates. It's possible to use either an in-house certificate server or a commercial certificate authority for your extranet, just as for your VPN. But, if the certificates are being issued by each participating business, then you'll need a way to cross-certify the *certificate authorities* (CAs) (see Figure 16.6). That's relatively easy if everyone's using commercial CAs, because they're usually already cross-certified. If you're maintaining your own certificate server, then you'll have to register your certificate server with an outside CA or your partners' certificate server.

If you're creating an extranet by combining the VPNs of business partners, there will always be the issue of interoperability between the VPNs. IPSec is a big step toward solving these interoperability issues, but don't expect all of today's IPSec-compliant products to work together automatically. The standards for key management are so new that all vendors haven't implemented them, and all the kinks haven't been worked out yet. Each partner's procedures for managing their VPN may not be the same, either, so you'll have to determine what work-arounds, if any, are needed to get each VPN to work with the other VPNs.

Two extranet-related issues don't deal with your VPN infrastructure but deserve mention. Both relate to the fact that your extranet involves sharing resources between partners.

First, some extranets are created for joint project teams, and the data that they work with and generate has to be stored somewhere. When you're designing an inventory control system for an extranet, it may be obvious where the server should reside; with joint project teams, that's less obvious. Some member of the extranet will have to decide to take ownership of the server(s) for the joint project and maintain its security.

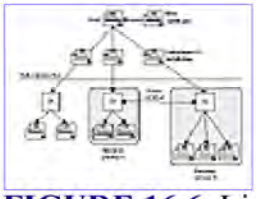

FIGURE 16.6 Linking company CAs together.

Second is the issue of problem resolution. Even though each device on an extranet is managed by someone, the end-to-end connection crosses company boundaries. When something fails, it may be difficult to discover which device is at fault. To counter this, you'll need to institute some kind of

problem-reporting procedures and methods to link together each company's help desk so that they can collaborate on solving any problems that crop up on the extranet. The last thing you want is finger-pointing.

Whether you have an in-house VPN or an outsourced VPN, you can choose to have an ISP maintain your extranet. Service providers offering extranet services usually maintain authentication servers as well as any Web and application servers that you may need for your extranet. Like an outsourced VPN, you should maintain control of user authentication and access rights for the outsourced extranet. Outsourcing the extranet may mean using a very limited portion of your VPN (your authentication servers, for example), but it's a good alternative if you feel better with separate systems for securing your internal communications and for dealing with external partners.

Just as with your VPN, you should plan to roll out your extranet in stages, starting with a pilot project. Pick a few trusted and knowledgeable customers who understand what you're trying to do with your extranet. Be sure that they're willing to deal with glitches in the system in order to make it work. If your extranet applications involve a number of different users (as in supply chain management, for instance), be sure that representatives of each class of user are included in the testing. In any case, make sure that everything that your company is responsible for works on your own network before you bring in any partners for testing.

Summary

Extranets are usually formed between business partners because of a particular business application. Because VPNs form the plumbing for secure networks and can secure any kind of network traffic, regardless of the application, you can build an extranet on top of a VPN. The main step in extending a VPN to an extranet is granting your extranet partners the access rights to specific internal resources and adding them to your authentication systems.

Many different compatibility issues will arise as you attempt to deploy an extranet because you cannot guarantee that each business runs the same types of network. Some of the compatibility issues that you'll need to resolve, revolve around security policies, existing VPNs, types of authentication deployed by the partners, and how keys and digital certificates will be distributed.

Previous | Table of Contents | Next

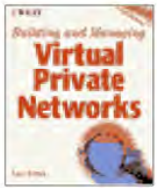

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

CHAPTER 17
Future Directions

Virtual Private Networks using the Internet are an ever-increasing opportunity for businesses, vendors, and ISPs alike. It's been projected by Infonetics Research that the market for VPN products will reach $12 billion in 2001. Many of the business forces motivating the deployment of VPNs, such as cost reductions and changes in telecommunications and networking, will remain in effect for quite a few years. If anything, these forces are likely to get even stronger over time.

Let's take a look at what's likely to happen with various developments affecting VPNs over the next few years. First, we'll look at how businesses will deploy VPNs and the effects ISPs may have on future VPNs. Then we'll say a few words about the state of standards, security, and digital certificates before moving on to managing your VPN. Finally, we'll cover some of the trends in VPN product lines.

VPN Deployment

One of the major uses for VPNs currently is the replacement of existing remote access systems using modems and remote access servers with dial-in VPNs. Not only can a dial-in VPN be less costly, but it provides more flexibility to the mobile user, whether he's a salesperson on the road or a telecommuter.

Dial-in VPNs will continue to be an important use of VPNs for some time. In fact, for some vendors and managers, it's the only VPN they know. The newly formed roaming services, which have added value to remote connections by consolidating access from different ISPs, will continue to make dial-in VPNs useful, especially for multinational businesses. As standardization of VPN protocols continue, expect roaming services to provide client support for the crucial protocols, such as L2TP and IPSec.

Even as companies move out of the pilot project stage with their VPNs, managers are looking for more uses for their VPNs. As voice over IP (or IP telephony) starts to look more attractive, the idea of combining voice and data networking over IP will become a more concrete possibility. Some care in provisioning proper latencies for voice may be required, but many ISPs' offerings are, or will be, suitable for this application. Secure videoconferencing is another application of interest, but this application may require even more constraints on bandwidth and quality-of-service.

Another trend that could drive further VPN usage is the development of the universal mailbox, in which e-mail, faxes, and phone calls can all be received and processed in a single application. The first generation of these systems is already available from Lucent/Octel and Nortel and deployed in some

companies, and the standards that make it easier to transmit faxes and phone messages using e-mail are nearing completion, which will make it easier for these systems to interoperate. Transmitting these various forms of information over data networks makes it easier to secure them with a VPN.

Some companies have also talked about *internal VPNs*, which are local private networks organized around a department or project, designed to restrict access and keep communications secure against internal snooping. Many security experts have pointed out that more security breaches are due to internal attacks than any other single cause. For such applications, installing VPN software on a departmental router or firewall may suffice. Alternatively, as host-based encryption and authentication becomes more widespread, personal tunnels can be set up between team members for secure communications within the enterprise.

ISPs and the Internet

ISPs will continue to play a crucial role in the evolution of VPNs. They have very little to gain if they act only as a transmission service, and they're faced with a bigger business opportunity if they can offer value-added services. In some cases, this might be as simple as offering connections with reduced latency. But, expect ISPs to go beyond this and offer true differentiated services, probably utilizing either Class of Service technology (Cisco and 3Com already offer products for this) or *Multi-Protocol Label Switching* (MPLS), which is being offered by Ascend to ISPs and enterprise networks and also should be provided by other vendors as the protocol becomes an IETF standard.

One continued hitch in these developments is the restriction of differentiated services to a single ISP. Although it is (or soon will be) technically possible to provide such services across ISP domains, most ISPs are reluctant to consider offering such services for fear they'll loose business. Don't expect to see cross-ISP performance guarantees until the providers come up with a method for billing for different classes of traffic as it crosses from one ISP to another, which may take years. ISPs still want to show that their backbone is the best and want your traffic to travel only on their backbone if you want special treatment. (That, of course, is part of the definition of value-added services.)

Not all ISPs will be able to deploy the technologies required for differentiated services; some of them simply won't be able to afford it. Expect to see a continued dichotomy of ISPs, with the multinational and larger national and regional ISPs offering better support for VPNs and differentiated services, and smaller ISPs will continue to offer basic services and some VPN services (such as dial-in VPNs).

Previous | Table of Contents | Next

Although IPSec is more of a site-to-site tunneling protocol that doesn't require any ISP intervention, both PPTP and L2TP provide ISPs with an opportunity to provide value-added services for a VPN. Looking back at Chapters 6 and 7, you may recall that ISPs can use special remote access concentrators to initiate tunnels on behalf of remote callers, which allows for better control of where tunnels are terminated and avoids the need for special client software. With the high demand for dial-in VPNs, expect many ISPs will offer this type of tunneling service.

Hosting and managing outsourced VPNs is another service offered by some ISPs and one that will continue to grow. It's only been within the last year that ISPs have begun offering managed VPNs, and the market is very new. But, considering the complexity that VPNs can entail, competent ISPs, many of whom already offer managed security services for instance, will offer to manage your VPN for you.

Some of these outsourced VPN services also include hosting Web servers or other servers that can either be used internally or form the basis for an extranet. Depending on your ISP, outsourcing a VPN or extranet can include just about anything, including the network equipment, security management, serving as a certificate authority, and hosting any servers that you may need.

VPN Standards

All of the VPN protocols we covered in this book—PPTP, L2TP, and IPSec—will continue to be used over the next few years. PPTP will continue to be used due to the free availability of a Windows client and the relatively low cost of Microsoft's NT Server with RRAS. Small businesses in particular probably will use this solution for some time.

Since L2TP is just now becoming a standard, its deployment probably will take a little longer. Once again, Microsoft may have a hand in increasing its acceptance when L2TP clients get rolled out with Windows 98 and NT 5.0 includes server support. Other companies also are shipping, or will soon ship, products for L2TP, but the added complexity of L2TP and its use of IPSec for encryption may slow down its deployment, at least initially.

For anyone expecting to create a LAN-to-LAN VPN, IPSec will be the protocol of choice. The main components of IPSec, including key management via IKE, are close to standards approval by the IETF as this book goes to press. These standards will be an important step in improving the interoperability of IPSec-based products, which should help further deployment of IPSec VPNs.

IPSec currently may have some shortcomings when dealing with remote users and dial-in VPNs, but extensions to IPSec to simplify user authentication (including links to security tokens and smart cards) and to configure IPSec clients already have been proposed. With the strong interest in remote access to

VPNs, it shouldn't be long before the current IPSec standards are extended to improve remote access.

IPSec client configuration will also take on a larger role when host-to-host tunnels are employed. The main emphasis in the IPSec community thus far has been on LAN-to-LAN tunneling, but some developers are already looking to tunnel creation at individual hosts for added security within the enterprise. For host-based tunneling to be feasible on the scale of large enterprise networks, the infrastructure for key management will need to be improved to handle the large number of keys, so it'll be a while before this type of tunneling becomes widespread. Furthermore, not all corporate networks will need this type of security; security gateways for LAN-to-LAN VPNs will prove to be adequate solutions for many corporations for some time to come.

Although IPSec started out as part of the development of IPv6, it has pretty much taken on a life of its own lately. For instance, this book has discussed the capabilities of IPSec in an IPv4 world, rather than limiting it to IPv6. Deployment of IPv6 will make IPSec a standard feature on any host, but it may be 5–10 years before we see appreciable deployment of IPv6. One of the principal needs for IPv6, insufficient numbers of addresses, has been somewhat alleviated by other, somewhat short-term solutions like CIDR and NAT. Current demand for IPv6 seems low, and vendor response seems correspondingly low. As far as support for IPSec is concerned, only the most basic IPSec functions have been included in commercially available IPv6 protocol stacks, so vendor support will need to be ramped up before deployment of IPv6 increases.

Security and Digital Certificates

Even as companies work to incorporate past cryptographic algorithms into their products, scientists have been developing new algorithms that are both computationally faster and more difficult to break. One of the newest algorithms that's being tested and becoming available in some commercial products (though not for VPNs yet) is *elliptic curve cryptography* (ECC). If ECC becomes more widespread and if cryptanalysts are satisfied with its robustness against attack, it's likely that the IETF will include the algorithm in its list of optional algorithms for use with IPSec.

The U.S. government also has been looking for a replacement for DES as its recommended encryption algorithm, with plans to select a new algorithm early in the next millennium. This choice will no doubt affect which algorithms are routinely selected for encryption on VPNs, as well as current governmental restrictions on exporting encryption products.

Previous | Table of Contents | Next

One of the most active areas of development in the security market these days is that surrounding the use of digital certificates. More products are being developed to use digital certificates for authenticating users, and these products should be readily integrated into VPN systems. One factor that's due to make the integration of use for digital certificates easier is the use of LDAP-compatible directories. X.500 and LDAP directories can be used to store certificates, and LDAP is being increasingly used as the method for accessing those certificates and related user information.

As we mentioned in Chapter 13, "Security Management," some companies are also working on the use of smart cards for carrying digital certificates. Although other card-based security tokens are available in the marketplace, the increasing number of uses for digital certificates and the deployment of LDAP-based directory infrastructures probably will drive the use of card-based certificates past that of other portable security tokens.

But, widespread use of digital certificates is still somewhat hampered by the current *Public Key Infrastructures* (PKIs). The current processes for distributing and checking Certificate Revocation Lists is both awkward and slow and poorly suited to the more dynamic uses like VPNs and secure e-mail, for example. One likely solution to this problem will be the deployment of *On-line Certificate Status Protocol* (OCSP) after it's approved as a standard.

VPN Management

For the next few years, managing VPNs will be one of the biggest concerns as standards and systems evolve. LAN-to-LAN VPNs are easier to manage because most of the VPN processes are transparent to end-users, and key management is largely based on sites rather than individuals. Of course, adding large numbers of remote users requires improved scalability of both your authentication and key-management systems.

Because IKE is only now being approved as the key-management standard for use with IPSec, expect companies to take another year to work out the kinks between their products to improve interoperability across all sizes of networks.

As we've mentioned previously, management of digital certificates is another area that needs improvement. Better handling of revoked certificates and distribution of the certificates themselves will develop over the next few years as more businesses seek to deploy PKIs both for internal uses as well as on extranets.

A relatively new line of products aimed at policy-based network management should eventually ease the management of VPNs, as well as of ordinary networks. But, policy-based network management is a very

young market, and major vendors have only started to ship products in the last half of 1998. It'll still be a few more years before all installed network devices on an enterprise network will be able to take part in policy-based management.

Many policy-based management systems will come to depend on the *Directory Enabled Networks* (DEN) Initiative, which means that LDAP will become the common glue between devices and directories. Because user profiles, device configuration, and bandwidth provisioning can be lumped into DEN's framework, you can anticipate the deployment of DEN on your own networks by looking for VPN devices that include LDAP support. You will have time to plan for this, though, because the first directory and management products for DEN aren't scheduled to ship until early 1999, when Windows NT 5.0 ships.

The DEN Initiative, which started out as a collaborative effort between Cisco and Microsoft and eventually added the support of 20 other companies by early 1998, is now being managed by the *Desktop Management Task Force* (DMTF) to encourage a more open standards environment. Before DEN can be applied to VPN management, the DMTF will have to address links with IPSec as well as the security surrounding the transmission of device configurations from directories. DEN may eventually make network management easier, but it could lead to compromises in your security policies if all the exchanges of information required to keep DEN working aren't secured properly.

Product Trends

With all the hype surrounding VPNs, new developments in VPN products seem to occur on a weekly basis. Although we haven't emphasized the fact in this book, most of the current slate of products are point solutions rather than being highly integrated, which adds to the difficulty of installing, configuring, and maintaining all the pieces of your VPN. The easiest way to get an integrated solution right now is to outsource the VPN to a qualified ISP that also installs and maintain the equipment for your VPN. But, integrated solutions also will become available from vendors as well, now that the efforts on IPSec, key management, LDAP, and network management are converging on standards.

One approach to integration has been the *integrated box* product, in which a security gateway, firewall, and other network services are all bundled into one device (see Chapter 11, "VPN Hardware"). Expect to see more of these products over the next few years, with improved management applications as well. Be aware that many of these first devices included poorly designed management applications that did little to integrate the services they were meant to configure and manage.

Previous | Table of Contents | Next

Many of these integrated devices are combining a wide variety of network services, including Web and e-mail services as well as the security services you'd expect for a VPN. As we mentioned in Chapter 11, you'll need to decide how many of your network services you want installed in a single box. The more services in a single device, the more reliable and secure it has to be. Our own view is that distribution of services among a number of devices is preferable, as long as the services can be managed from a single workstation.

Integrated VPN devices, particularly those that are turnkey systems, require little configuration and can be particularly appealing to small businesses setting up a VPN. It's obvious that some of the products we listed in Chapter 11 are aimed at the small-business market, with more on the way. If you're in the market for this class of products, be sure to review them with future scalability in mind. Buying a product with no expansion possibilities or compatibility with accepted standards can lead to an expensive redesign of your VPN down the road as your business grows.

Although it may not be wise to add e-mail and Web servers to your integrated VPN device, letting the VPN device act as a control point for bandwidth provisioning and quality-of-service deserves consideration, as we pointed out in Chapter 15, "Performance Management." The market for QoS will mature over the next few years, so we expect that more products combining security and QoS will become available (and, of course, will be managed via DEN and/or policy-based management systems).

TABLE 17.1 Important IETF Working Groups

Group Name	*Acronym*	*URL*
Authenticated Firewall Traversal	AFT	www.ietf.org/html.charters/aft-charter.html
Common Authentication Technology	CAT	www.ietf.org/html.charters/cat-charter.html
Differentiated Services		www.ietf.org/html.charters/diff-serv-charter.html
Domain Name System Security	DNSSE	www.ietf.org/html.charters/dnssec-charter.html
Dynamic Host Configuration	DHC	www.ietf.org/html.charters/dhc-charter.html
Electronic Data Interchange-Internet Integration	EDIINT	www.ietf.org/html.charters/ediint-charter.html
Integrated Services	INTSERV	www.ietf.org/html.charters/intserv-charter.html
IPNG (IPv6)	IPNGWG	www.ietf.org/html.charters/ipngwg-charter.html
IP Security Protocol	IPSEC	www.ietf.org/html.charters/ipsec-charter.html

LDAP Service Deployment	LSD	www.ietf.org/html.charters/lsd-charter.html
Multiprotocol Label Switching	MPLS	www.ietf.org/html.charters/mpls-charter.html
Network Address Translators	NAT	www.ietf.org/html.charters/nat-charter.html
One Time Password Authentication	OTP	www.ietf.org/html.charters/otp-charter.html
Point-to-Point Protocol Extensions	PPPEXT	www.ietf.org/html.charters/pppext-charter.html
Procedures for Internet/Enterprise Renumbering	PIER	www.ietf.org/html.charters/pier-charter.html
Public Key Infrastructure (X.509)	PKIX	www.ietf.org/html.charters/pkix-charter.html
Remote Authentication Dial-In User Service	RADIUS	www.ietf.org/html.charters/radius-charter.html
Resource Reservation Setup Protocol	RSVP	www.ietf.org/html.charters/rsvp-charter.html

---

Keeping Up

VPNs are a dynamic market. Standards are still being developed, not only for VPN security but for digital certificates and network management as well, all of which impact your ability to design and deploy a proper VPN.

If you want to track what's going on with some of these standards efforts, check the Web sites listed in Table 17.1. Also check out the vendors themselves and other related resources that we've listed in Appendix B.

Previous | Table of Contents | Next

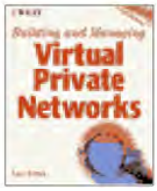

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

APPENDIX A
Resources

Books

Bernstein, Terry, Anish B. Bhimani, Eugene Schultz, and Carol A. Siegel. *Internet Security for Business*. New York: John Wiley & Sons, Inc., 1996.

Comer, Douglas E. *Internetworking with TCP/IP, Vol. I, Principles, Protocols and Architecture*, Third Edition. Englewood Cliffs, NJ: Prentice-Hall, 1995.

Schneier, Bruce. *Applied Cryptography*, Second Edition. New York: John Wiley and Sons, Inc., 1996.

IETF Documents—RFCS

2315 PKCS #7: *Cryptographic Message Syntax Version 1.5*. B. Kaliski. March, 1998. (Status: INFORMATIONAL)

2314 PKCS #10: *Certification Request Syntax Version 1.5*. B. Kaliski. March, 1998. (Status: INFORMATIONAL)

2313 PKCS #1: *RSA Encryption Version 1.5*. B. Kaliski. March, 1998. (Status: INFORMATIONAL)

2307 *An Approach for Using LDAP as a Network Information Service*. L. Howard. March, 1998. (Status: EXPERIMENTAL)

2289 *A One-Time Password System*. N. Haller, C. Metz, P. Nesser, and M. Straw. February, 1998. (Obsoletes RFC1938) (Status: DRAFT STANDARD)

2284 *PPP Extensible Authentication Protocol (EAP)*. L. Blunk and J. Vollbrecht. March, 1998. (Status: PROPOSED STANDARD)

2268 *A Description of the RC2(r) Encryption Algorithm*. R. Rivest. January, 1998. (Status: INFORMATIONAL)

2260 *Scalable Support for Multi-Homed Multi-Provider Connectivity*. T. Bates and Y. Rekhter. January, 1998. (Status: INFORMATIONAL)

2256 *A Summary of the X.500(96) User Schema for Use with LDAPv3*. M. Wahl. December, 1997. (Status: PROPOSED STANDARD)

2251 *Lightweight Directory Access Protocol (v3)*. M. Wahl, T. Howes, and S. Kille. December, 1997. (Status: PROPOSED STANDARD)

2230 *Key Exchange Delegation Record for the DNS*. R. Atkinson. October, 1997. (Status: INFORMATIONAL)

2219 *Use of DNS Aliases for Network Services*. M. Hamilton and R. Wright. October, 1997. (Status: BEST CURRENT PRACTICE)

2215 *General Characterization Parameters for Integrated Service Network Elements*. S. Shenker and J. Wroclawski. September, 1997. (Status: PROPOSED STANDARD)

2212 *Specification of Guaranteed Quality of Service*. S. Shenker, C. Partridge, and R. Guerin. September, 1997. (Status: PROPOSED STANDARD)

2211 *Specification of the Controlled-Load Network Element Service*. J. Wroclawski. September, 1997. (Status: PROPOSED STANDARD)

2210 *The Use of RSVP with IETF Integrated Services*. J. Wroclawski. September, 1997. (Status: PROPOSED STANDARD)

2209 *Resource ReSerVation Protocol (RSVP)—Version 1 Message Processing Rules*. R. Braden and L. Zhang. September, 1997. (Status: INFORMATIONAL)

2208 *Resource ReSerVation Protocol (RSVP)—Version 1 Applicability Statement Some Guidelines on Deployment*. A. Mankin, Ed., F. Baker, B. Braden, S. Bradner, M. O'Dell, A. Romanow, A. Weinrib, and L. Zhang. September, 1997. (Status: INFORMATIONAL)

2207 *RSVP Extensions for IPSEC Data Flows*. L. Berger and T. O'Malley. September, 1997. (Status: PROPOSED STANDARD)

2205 *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*. R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, and S. Jamin. September, 1997. (Status: PROPOSED STANDARD)

2196 *Site Security Handbook*. B. Fraser. September, 1997. (Status: INFORMATIONAL)

2182 *Selection and Operation of Secondary DNS Servers*. R. Elz, R. Bush, S. Bradner, and M. Patton. July, 1997. (Status: BEST CURRENT PRACTICE)

2181 *Clarifications to the DNS Specification*. R. Elz and R. Bush. July, 1997. (Status: PROPOSED STANDARD)

2153 *PPP Vendor Extensions*. W. Simpson. May, 1997. (Status: INFORMATIONAL)

2144 *The CAST-128 Encryption Algorithm*. C. Adams. May, 1997. (Status: INFORMATIONAL)

2139 *RADIUS Accounting*. C. Rigney. April, 1997. (Status: INFORMATIONAL)

2138 *Remote Authentication Dial-In User Service (RADIUS)*. C. Rigney, A. Rubens, W. Simpson, and S. Willens. April, 1997. (Status: PROPOSED STANDARD)

2137 *Secure Domain Name System Dynamic Update*. D. Eastlake. April, 1997. (Status: PROPOSED STANDARD)

2136 *Dynamic Updates in the Domain Name System (DNS UPDATE)*. P. Vixie, Ed., S. Thomson, Y. Rekhter, and J. Bound. April, 1997. (Status: PROPOSED STANDARD)

2132 *DHCP Options and BOOTP Vendor Extensions*. S. Alexander and R. Droms. March, 1997. (Status: DRAFT STANDARD)

2131 *Dynamic Host Configuration Protocol*. R. Droms. March, 1997. (Status: DRAFT STANDARD)

2125 *The PPP Bandwidth Allocation Protocol (BAP) / The PPP Bandwidth Allocation Control Protocol (BACP)*. C. Richards and K. Smith. March, 1997. (Status: PROPOSED STANDARD)

2118 *Microsoft Point-To-Point Compression (MPPC) Protocol*. G. Pall. March, 1997. (Status: INFORMATIONAL)

2104 *HMAC: Keyed-Hashing for Message Authentication*. H. Krawczyk, M. Bellare, and R. Canetti. February, 1997. (Status: INFORMATIONAL)

2085 *HMAC-MD5 IP Authentication with Replay Prevention*. M. Oehler and R. Glenn. February, 1997. (Status: PROPOSED STANDARD)

2078 *Generic Security Service Application Program Interface*, Version 2. J. Linn. January, 1997. (Status: PROPOSED STANDARD)

2065 *Domain Name System Security Extensions*. D. Eastlake, 3rd, and C. Kaufman. January, 1997. (Status: PROPOSED STANDARD)

2040 *The RC5, RC5-CBC, RC5-CBC-Pad, and RC5-CTS Algorithms*. R. Baldwin and R. Rivest. October, 1996. (Status: INFORMATIONAL)

2025 *The Simple Public-Key GSS-API Mechanism (SPKM)*. C. Adams. October, 1996. (Status: PROPOSED STANDARD)

2008 *Implications of Various Address Allocation Policies for Internet Routing*. Y. Rekhter and T. Li. October, 1996. (Status: BEST CURRENT PRACTICE)

1995 *Incremental Zone Transfer in DNS*. M. Ohta. August, 1996. (Status: PROPOSED STANDARD)

1994 *PPP Challenge Handshake Authentication Protocol (CHAP)*. W. Simpson. August, 1996. (Status: DRAFT STANDARD)

1993 *PPP Gandalf FZA Compression Protocol*. A. Barbir, D. Carr, and W. Simpson. August, 1996. (Status: INFORMATIONAL)

1969 *The PPP DES Encryption Protocol (DESE)*. K. Sklower and G. Meyer. June, 1996. (Status: INFORMATIONAL)

1968 *The PPP Encryption Control Protocol (ECP)*. G. Meyer. June, 1996. (Status: PROPOSED STANDARD)

1967 *PPP LZS-DCP Compression Protocol (LZS-DCP)*. K. Schneider and R. Friend. August, 1996. (Status: INFORMATIONAL)

1962 *The PPP Compression Control Protocol (CCP)*. D. Rand. June, 1996. (Status: PROPOSED STANDARD)

1961 *GSS-API Authentication Method for SOCKS Version 5*. P. McMahon. June, 1996. (Status: PROPOSED STANDARD)

1935 *What is the Internet, Anyway?* J. Quarterman and S. Carl-Mitchell. April, 1996. (Status: INFORMATIONAL)

1933 *Transition Mechanisms for IPv6 Hosts and Routers*. R. Gilligan and E. Nordmark. April, 1996. (Status: PROPOSED STANDARD)

1929 *Username/Password Authentication for SOCKS V5*. M. Leech. April, 1996. (Status:

PROPOSED STANDARD)

1928 *SOCKS Protocol Version 5*. M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, and L. Jones. April, 1996. (Status: PROPOSED STANDARD)

1918 *Address Allocation for Private Internets*. Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. February, 1996. (Status: BEST CURRENT PRACTICE)

1912 *Common DNS Operational and Configuration Errors*. D. Barr. February, 1996. (Status: INFORMATIONAL)

1900 *Renumbering Needs Work*. B. Carpenter and Y. Rekhter. February, 1996. (Status: INFORMATIONAL)

1884 *IP Version 6 Addressing Architecture*. R. Hinden and S. Deering, Editors. December, 1995. (Status: PROPOSED STANDARD)

1883 *Internet Protocol, Version 6 (IPv6) Specification*. S. Deering and R. Hinden. December, 1995. (Status: PROPOSED STANDARD)

1881 *IPv6 Address Allocation Management*. IAB and IESG. December, 1995. (Status: INFORMATIONAL)

1865 *EDI Meets the Internet Frequently Asked Questions about Electronic Data Interchange (EDI) on the Internet*. W. Houser, J. Griffin, and C. Hage. January, 1996. (Status: INFORMATIONAL)

1852 *IP Authentication Using Keyed SHA*. P. Metzger and W. Simpson. September, 1995. (Status: EXPERIMENTAL)

1851 *The ESP Triple DES Transform*. P. Karn, P. Metzger, and W. Simpson. September, 1995. (Status: EXPERIMENTAL)

1829 *The ESP DES-CBC Transform*. P. Karn, P. Metzger, and W. Simpson. August, 1995. (Status: PROPOSED STANDARD)

1828 *IP Authentication Using Keyed MD5*. P. Metzger and W. Simpson. August, 1995. (Status: PROPOSED STANDARD)

1827 *IP Encapsulating Security Payload (ESP)*. R. Atkinson. August, 1995. (Status: PROPOSED STANDARD)

1826 *IP Authentication Header*. R. Atkinson. August, 1995. (Status: PROPOSED STANDARD)

1825 *Security Architecture for the Internet Protocol*. R. Atkinson. August, 1995. (Status: PROPOSED STANDARD)

1817 *CIDR and Classful Routing*. Y. Rekhter. August, 1995. (Status: INFORMATIONAL)

1794 *DNS Support for Load Balancing*. T. Brisco. April, 1995. (Status: INFORMATIONAL)

1760 *The S/KEY One-Time Password System*. N. Haller. February, 1995. (Status: INFORMATIONAL)

1702 *Generic Routing Encapsulation over IPv4 Networks*. S. Hanks, T. Li, D. Farinacci, and P. Traina. October, 1994. (Status: INFORMATIONAL)

1701 *Generic Routing Encapsulation (GRE)*. S. Hanks, T. Li, D. Farinacci, and P. Traina. October, 1994. (Status: INFORMATIONAL)

1661 *The Point-to-Point Protocol (PPP)*. W. Simpson, Editor. July, 1994. (Status: STANDARD)

1631 *The IP Network Address Translator (NAT)*. K. Egevang and P. Francis. May, 1994. (Status: INFORMATIONAL)

1591 *Domain Name System Structure and Delegation*. J. Postel. March, 1994. (Status: INFORMATIONAL)

1531 *Dynamic Host Configuration Protocol*. R. Droms. October, 1993. (Status: PROPOSED STANDARD)

1321 *The MD5 Message-Digest Algorithm*. R. Rivest. April, 1992. (Status: INFORMATIONAL)

1320 *The MD4 Message-Digest Algorithm*. R. Rivest. April, 1992. (Status: INFORMATIONAL)

1319 *The MD2 Message-Digest Algorithm*. B. Kaliski. April, 1992. (Status: INFORMATIONAL)

1035 *Domain Names—Implementation and Specification*. P. V. Mockapetris. November 1, 1987. (Status: STANDARD)

1034 *Domain Names—Concepts and Facilities*. P. V. Mockapetris. November 1, 1987. (Status: STANDARD)

IETF Documents—Internet Drafts

IETF Internet-Drafts usually are circulated and stored for six months after publication. At the end of that period, they are either replaced with a new version, replaced with an entirely new Internet-draft, converted to an RFC, or removed. The names and dates of the drafts listed were accurate as of May, 1998.

The documents have been listed according to the working group responsible for their development. The last group, labelled individual submissions, includes any Internet-Draft that may be appropriate to the subject of this book but has not been officially generated by a specific working group.

Access, Searching, and Indexing of Directories

Genovese, Tony, M. Wahl, and Y. Yaacovi. "Lightweight Directory Access Protocol (v3): Extensions for Dynamic Directory Services," 12/23/1997, <draft-ietf-asid-ldapv3-dynamic-07.txt>

Stokes, E., R. Weiser, and Bob Huston. "LDAP Replication Requirements," 11/26/1997, <draft-ietf-asid-replica-req-01.txt>

Authenticated Firewall Traversal

Kayashima, Makoto, Tsukasa Ogino, Masato Terada, and Yoichi Fujiyama. "SOCKS V5 Protocol Extension for Multiple Firewalls Traversal," 11/26/1997, <draft-ietf-aft-socks-multiple-traversal-00.txt>

Michener, J., and Dan Fritch. "Multi-Authentication Framework Method for SOCKS V5," 03/13/1998, <draft-ietf-aft-socks-maf-00.txt>

VanHeyningen, Marc. "Challenge-Handshake Authentication Protocol for SOCKS V5," 01/07/1998, <draft-ietf-aft-socks-chap-01.txt>

VanHeyningen, Marc. "SOCKS Protocol Version 5," 03/05/1998, <draft-ietf-aft-socks-pro-v5-02.txt>

Zorn, Glen, Pat Calhoun, and Jeff Haag. "EAP Authentication for SOCKS Version 5," 03/06/1998, <draft-ietf-aft-socks-eap-00.txt>

Differentiated Services

Brim, Scott, Frank Kastenholz, Fred Baker, John Renwick, Tony Li, and Shantigram Jagannath. "IP Precedence in Differentiated Services Using the Assured Service," 04/10/1998, <draft-ietf-diffserv-precedence-00.txt>

Nichols, Kathleen, and S. Blake. "Definition of the Differentiated Services Field (DS Byte) in the IPv4 and IPv6 Headers," 05/07/1998, <draft-ietf-diffserv-header-00.txt>

Nichols, Kathleen, and S. Blake. "Differentiated Services Operational Model and Definitions," 02/11/1998, <draft-nichols-dsopdef-00.txt>

Domain Name System Security

Eastlake, Don. "DNS Operational Security Considerations," 03/12/1998, <draft-ietf-dnssec-secops-01.txt>

Eastlake, Don. "Domain Name System Security Extensions," 05/05/1998, <draft-ietf-dnssec-secext2-05.txt>

Eastlake, Don. "DSA KEYs and SIGs in the Domain Name System," 01/27/1998, <draft-ietf-dnssec-dss-02.txt>

Eastlake, Don. "RSA/MD5 KEYs and SIGs in the Domain Name System (DNS)," 01/27/1998, <draft-ietf-dnssec-rsa-00.txt>

Eastlake, Don. "Secret Key Establishment for DNS (TKEY RR)," 02/23/1998, <draft-ietf-dnssec-tkey-00.txt>

Eastlake, Don. "Storage of Diffie-Hellman Keys in the Domain Name System (DNS)," 03/12/1998, <draft-ietf-dnssec-dhk-02.txt>

Gudmundsson, O., and D. Eastlake. "Storing Certificates in the Domain Name System (DNS)," 03/06/1998, <draft-ietf-dnssec-certs-02.txt>

Watson, Robert. "DNSsec Authentication Referral Record (AR)," 11/26/1997, <draft-ietf-dnssec-ar-00.txt>

Dynamic Host Configuration

Droms, Ralph, and O. Gudmundsson. "Security Requirements for the DHCP Protocol," 03/13/1998, <draft-ietf-dhc-security-requirements-00.txt>

Rekhter, Y. "Interaction between DHCP and DNS," 03/05/1998, <draft-ietf-dhc-dhcp-dns-08.txt>

Electronic Data Interchange-Internet Integration

Drummond, Rik, M. Jansson, and C. Shih. "MIME-Based Secure EDI," 12/04/1997, <draft-ietf-ediint-as1-05.txt>

Drummond, Rik, M. Jansson, and C. Shih. "Requirements for InterOperable Internet EDI," 04/27/1998, <draft-ietf-ediint-req-05.txt>

IP Performance Metrics

Anon. "Connectivity," 11/24/1997, <draft-ietf-ippm-connectivity-01.txt>

Demichelis, Carlo. "Instantaneous Packet Delay Variation Metric for IPPM," 03/11/1998, <draft-ietf-ippm-ipdv-00.txt>

IP Security Protocol

Adams, R., and R. Pereira. "The ESP CBC-Mode Cipher Algorithms," 03/10/1998, <draft-ietf-ipsec-ciph-cbc-02.txt>

Bhattacharya, P., and R. Pereira. "IPSec Policy Data Model," 02/25/1998, <draft-ietf-ipsec-policy-model-00.txt>

Carrel, D., and D. Harkins. "The Internet Key Exchange (IKE)," 03/13/1998, <draft-ietf-ipsec-isakmp-oakley-07.txt>

Doraswamy, Naganand. "Implementation of Virtual Private Network (VPNs) with IP Security," 03/14/1997, <draft-ietf-ipsec-vpn-00.txt>

Doraswamy, Naganand, and C. Madson. "The ESP DES-CBC Cipher Algorithm with Explicit IV," 02/13/1998, <draft-ietf-ipsec-ciph-des-expiv-02.txt>

Kent, Stephen, and Ran Atkinson. "IP Authentication Header," 05/12/1998, <draft-ietf-ipsec-auth-header-06.txt>

Kent, Stephen, and Ran Atkinson. "IP Encapsulating Security Payload (ESP)," 05/12/1998, <draft-ietf-ipsec-esp-v2-05.txt>

Kent, Stephen, and Ran Atkinson. "Security Architecture for the Internet Protocol," 05/12/1998, <draft-ietf-ipsec-arch-sec-05.txt>

Kent, Stephen, and R. Glenn. "The NULL Encryption Algorithm and Its Use with IPsec," 03/13/1998, <draft-ietf-ipsec-ciph-null-00.txt>

Madson, C., and R. Glenn. "The Use of HMAC-MD5-96 within ESP and AH," 02/18/1998, <draft-ietf-ipsec-auth-hmac-md5-96-03.txt>

Madson, C., and R. Glenn. "The Use of HMAC-SHA-1-96 within ESP and AH," 02/18/1998, <draft-ietf-ipsec-auth-hmac-sha196-03.txt>

Maughan, D., M. Schertler, M. Schneider, and J. Turner. "Internet Security Association and Key Management Protocol (ISAKMP)," 03/11/1998, <draft-ietf-ipsec-isakmp-09.txt.ps>

Orman, H. "The OAKLEY Key Determination Protocol," 07/25/1997, <draft-ietf-ipsec-oakley-02.txt>

Patel, B. "Dynamic Remote Host Configuration over IPSEC Using DHCP," 12/04/1997, <draft-ietf-ipsec-dhcp-00.txt>

Patel, B., and Michael Jeronimo. "Revised SA Negotiation Mode for ISAKMP/Oakley," 12/04/1997, <draft-ietf-ipsec-isakmp-SA-revised-00.txt>

Patel, B., R. Pereira, and S. Anand. "The ISAKMP Configuration Method," 04/23/1998, <draft-ietf-ipsec-isakmp-mode-cfg-03.txt>

Pereira, R. "Extended Authentication within ISAKMP/Oakley," 02/24/1998,

<draft-ietf-ipsec-isakmp-xauth-01.txt>

Piper, D. "A GSS-API Authentication Mode for ISAKMP/Oakley," 12/23/1997, <draft-ietf-ipsec-isakmp-gss-auth-01.txt>

Piper, D. "The Internet IP Security Domain of Interpretation for ISAKMP," 05/13/1998, <draft-ietf-ipsec-ipsec-doi-09.txt>

Piper, D., and D. Harkins. "The Pre-Shared Key for the Internet Protocol," 04/06/1998, <draft-ietf-ipsec-internet-key-00.txt>

Provos, Niels. "The Use of HMAC-RIPEMD-160-96 within ESP and AH," 02/16/1998, <draft-ietf-ipsec-auth-hmac-ripemd-160-96-01.txt>

Simpson, W., Naganand Doraswamy, Perry Metzger. "The ESP Triple DES Transform," 07/03/1997, <draft-ietf-ipsec-ciph-des3-00.txt>

Thayer, Rodney, Naganand Doraswamy, and R. Glenn. "IP Security Document Roadmap," 12/04/1997, <draft-ietf-ipsec-doc-roadmap-02.txt>

Multiprotocol Label Switching

Callon, Ross, George Swallow, N. Feldman, A. Viswanathan, P. Doolan, and A. Fredette. "A Framework for Multiprotocol Label Switching," 11/26/1997, <draft-ietf-mpls-framework-02.txt>

Callon, Ross, A. Viswanathan, and E. Rosen. "Multiprotocol Label Switching Architecture," 04/07/1998, <draft-ietf-mpls-arch-01.txt>

Davie, Bruce, Y Rekhter, A. Viswanathan, S. Blake, Vijay Srinivasan, and E. Rosen. "Use of Label Switching with RSVP," 03/12/1998, <draft-ietf-mpls-rsvp-00.txt>

Next Generation (IPv6) Transition

Anon. "Transition Mechanisms for IPv6 Hosts and Routers," 11/21/1997, <draft-ietf-ngtrans-mech-00.txt>

Durand, A., and Bertrand Buclin. "IPv6 Routing Issues," 04/24/1998, <draft-ietf-ngtrans-6bone-routing-issues-02.txt>

Srisuresh, Pyda, and George Tsirtsis. "Network Address Translation—Protocol Translation (NAT-PT)," 03/06/1998, <draft-ietf-ngtrans-natpt-01.txt>

Point-to-Point Protocol Extensions

Aboba, B., and B. Patel. "Securing L2TP Using IPSEC," 03/11/1998, <draft-ietf-pppext-12tp-security-01.txt>

Calhoun, Pat, W. Mark Townsley, Sumit Vakil, and Donald Grosser. "Layer Two Tunneling Protocol 'L2TP' Security Extensions for Non-IP Networks," 03/18/1998, <draft-ietf-pppext-12tp-sec-03.txt>

Carter, G., "PPP EAP ISAKMP Authentication Protocol," 09/20/1997, <draft-ietf-pppext-eapiskamp-00.txt>

Casner, Stephen, C. Bormann, and M. Engan. "IP Header Compression over PPP," 12/23/1997, <draft-engan-ip-compress-02.txt>

Peirce, Ken, and Pat Calhoun. "Layer Two Tunneling Protocol 'L2TP' IP Differential Services Extension," 03/09/1998, <draft-ietf-pppext-12tp-ds-01.txt>

Peirce, Ken, and Pat Calhoun. "Layer Two Tunneling Protocol 'L2TP' Multi-Protocol Label Switching Extension," 03/09/1998, <draft-ietf-pppext-12tp-mpls-00.txt>

Rubens, Allan, William Palter, T. Kolar, G. Pall, M. Littlewood, A. Valencia, K. Hamzeh, W. Verthein, J. Taarud, and W. Mark Townsley. "Layer Two Tunneling Protocol 'L2TP,' " 04/06/1998, <draft-ietf-pppext-12tp-10.txt>

Valencia, A. "L2TP Header Compression ("L2TPHC")," 12/22/1997, <draft-ietf-pppext-12tphc-01.txt>

Zmuda, James, and W. Nace. "PPP Certificate Exchange Protocol," 12/03/1997, <draft-ietf-pppext-crtxchg-01.txt>

Zmuda, James, and W. Nace. "PPP EAP DSS Public Key Authentication Protocol," 12/03/1997, <draft-ietf-pppext-eapdss-01.txt>

Zorn, Glen, and S. Cobb. "Microsoft PPP C Extensions," 03/11/1998, <draft-ietf-pppext-mschap-00.txt>

Zorn, Glen, and G. Pall. "Microsoft Point-To-Point Encryption (MPPE) Protocol," 04/06/1998, <draft-ietf-pppext-mppe-01.txt>

Public Key Infrastructure (X.509)

Adams, C., and S. Farrell. "Internet X.509 Public Key Infrastructure Certificate Management Protocols," 02/26/1998, <draft-ietf-pkix-ipki3cmp-07.txt>

Adams, C., and M. Myers. "Internet X.509 Public Key Infrastructure Certificate Management Message Formats," 03/12/1998, <draft-ietf-pkix-cmmf-00.txt>

Adams, C., M. Myers, Ambarish Malpani, Rich Ankney, and Slava Galperin. "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol—OCSP," 04/07/1998, <draft-ietf-pkix-ocsp-03.txt>

Branchard, Marc. "Internet Public Key Infrastructure Caching the Online Certificate Status Protocol," 04/17/1998, <draft-ietf-pkix-ocsp-caching-00.txt>

Ford, Warwick, and S. Chokhani. "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework," 04/28/1998, <draft-ietf-pkix-ipki-part4-03.txt>

Ford, Warwick, and P. Hallam-Baker. "Internet X.509 Public Key Infrastructure OPEN CRL DISTRIBUTION PROCESS (OpenCDP)," 04/22/1998, <draft-ietf-pkix-ocdp-00.txt>

Fox, Barbara, Xiaoyi Liu, M. Myers, and Jeff Weinstein. "Certificate Management Messages over CMS," 03/12/1998, <draft-ietf-pkix-cmc-00.txt>

Housley, Russ. "Internet X.509 Public Key Infrastructure Operational Protocols: FTP and HTTP," 04/16/1998, <draft-ietf-pkix-opp-ftp-http-03.txt>

Howes, Tim, S. Boeyen, and P. Richard. "Internet X.509 Public Key Infrastructure LDAPv2 Schema," 03/16/1998, <draft-ietf-pkix-ldapv2-schema-00.txt>

Howes, Tim, S. Boeyen, and P. Richard. "Internet X.509 Public Key Infrastructure Operational Protocols—LDAPv2," 03/16/1998, <draft-ietf-pkix-ipki2opp-07.txt>

Solo, D., C. Adams, D. Kemp, and M. Myers. "Certificate Request Message Format," 02/25/1998,

<draft-ietf-pkix-crmf-00.txt>

Solo, D., Russ Housley, Warwick Ford, and T. Polk. "Internet Public Key Infrastructure X.509 Certificate and CRL Profile," 04/06/1998, <draft-ietf-pkix-ipki-part1-07.txt>

Remote Authentication Dial-In User Service

Calhoun, Pat, and Sumit Vakil. "RADIUS IP Security Extensions," 11/10/1997, <draft-ietf-radius-ipsec-00.txt>

Rubens, Allan, B. Aboba, and Pat Calhoun. "Extensible Authentication Protocol Support in RADIUS," 05/13/1998, <draft-ietf-radius-eap-05.txt>

Shriver, J., D. Leifer, Allan Rubens, and Glen Zorn, "RADIUS Attributes for Tunnel Protocol Support," 04/08/1998, <draft-ietf-radius-tunnel-auth-05.txt>

Zorn, Glen. "RADIUS Attributes for MS-CHAP Support," 11/06/1997, <draft-rfced-info-zorn-01.txt>

Zorn, Glen. "RADIUS Attributes for MS-CHAP Support," 11/18/1997, <draft-ietf-radius-mschap-attr-01.txt>

Zorn, Glen, and David Mitton. "RADIUS Accounting Modifications for Tunnel Protocol Support," 12/03/1997, <draft-ietf-radius-tunnel-acct-00.txt>

RSVP Admission Policy

Anon. "RSVP Extensions for Policy Control," 03/13/1998, <draft-ietf-rap-rsvp-ext-00.txt>

Herzog, Shai, A. Sastry, R. Rajan, Ron Cohen, J. Boyle, and David Durham. "The COPS (Common Open Policy Service) Protocol," 03/16/1998, <draft-ietf-rap-cops-01.txt>

Yavatkar, R., R. Guerin, and D. Pendarakis. "A Framework for Policy-Based Admission Control," 11/26/1997, <draft-ietf-rap-framework-00.txt>

Individual Submissions

Aboba, B. "Lightweight Directory Access Protocol (v3): Dynamic Attributes for the Remote Access Dial-in User Service (RADIUS)," 11/21/1997, <draft-aboba-dynradius-01.txt>

Aboba, B. "Lightweight Directory Access Protocol (v3): Extension for PPP Authentication," 11/21/1997, <draft-aboba-ppp-01.txt>

Aboba, B. "Lightweight Directory Access Protocol (v3): Schema for the Remote Access Dial-in User Service (RADIUS)," 02/05/1998, <draft-aboba-radius-02.txt>

Bartz, Larry. "LDAP Schema for Role-Based Access Control," 10/14/1997, <draft-bartz-hyperdrive-ldap-rbac-schema-00.txt>

Berkowitz, H. "To Be Multihomed: Requirements & Definitions," 03/11/1998, <draft-berkowitz-multirqmt-01.txt>

Carrel, D., and L. Grant. "The TACACS+ Protocol Version 1.78," 01/06/1998, <draft-grant-tacacs-02.txt>

Demizu, Noritoshi, and H. Izumiyama. "Dynamic Tunnel Configuration Protocol," 12/04/1997, <draft-demizu-udlr-dtcp-00.txt>

Ellesson, Ed, Dinesh Verma, R. Rajan, and S. Kamat. "Schema for Service Level Administration of Differentiated Services and Integrated Services in Networks," 02/25/1998, <draft-ellesson-sla-schema-00.txt>

Ferguson, P. "Simple Differential Services: IP TOS and Precedence, Delay Indication, and Drop Preference," 03/12/1998, <draft-ferguson-delay-drop-02.txt>

Heinanen, Juha, and E. Rosen. "VPN support for MPLS," 03/09/1998, <draft-heinanen-mpls-vpn-01.txt>

Jamieson, Dwight, Scott Pegrum, and Matthew Yuen. "VPN Multipoint to Multipoint Tunnel Protocol (VMMT)," 03/16/1998, <draft-pegrum-vmmt-00.txt>

Karn, P., and W. Simpson. "Photuris: Extended Schemes and Attributes," 03/06/1998, <draft-simpson-photuris-schemes-05.txt>

Karn, P., and W. Simpson. "Photuris: Session Key Management Protocol," 02/27/1998, <draft-simpson-photuris-18.txt>

Li, Tony, and Y Rekhter, "Provider Architecture for Differentiated Services and Traffic Engineering," 01/14/1998, <draft-li-paste-00.txt>

McDonald, D., B. Phan, and C. Metz. "PF_KEY Key Management API, Version 2," 02/27/1998, <draft-mcdonald-pf-key-v2-05.txt>

Moskowitz, Robert. "Network Address Translation Issues with IPsec," 02/12/1998, <draft-moskowitz-net66-vpn-00.txt>

O'Hara, John. "Configuration of Tunnel Mode IPSec Endpoint Parameters," 11/26/1997, <draft-ohara-ipsecparam-00.txt>

Ravikanth, Ravadurgam, and Pasi Vaananen. "Framework for Traffic Management in MPLS Networks," 03/19/1998, <draft-vaananen-mpls-tm-framework-00.txt>

Simpson, W. "Photuris: Secret Exchange," 05/06/1998, <draft-simpson-photuris-secret-00.txt>

Simpson, W., and Perry Metzger. "IP Authentication Using Keyed SHA1 with Data Padding," 05/01/1996, <draft-simpson-ah-sha-kdp-00.txt>

Srisuresh, Pyda, and Kjeld Egevang. "The IP Network Address Translator (NAT)," 03/05/1998, <draft-rfced-info-srisuresh-05.txt>

Suzuki, M. "Architecture of the Resource Reservation Service for the Commercial Internet," 02/09/1998, <draft-rfced-info-suzuki-00.txt>

Web Sites

Internet and ISP Information

CIX (Commercial Internet Exchange) www.cix.org

IETF www.ietf.org/

Internet Infrastructure, Service Provider Links, Topology Maps www.clark.net/pub/rbenn/isp.html

Mapnet www.caida.org/Tools/Mapnet/Backbones/

Russ Haynal's ISP Page navigators.com/isp.html

Security

Bruce Schneir www.counterpane.com

Cryptographic Protocols and Standards www.cs.hut.fi/ssh/crypto/protocols.html

International Computer Security Association www.icsa.net/

Koops' Crypto Law Survey cwis.kub.nl/~frw/people/koops/bertjaap.htm

NIST Computer Security Resource Clearinghouse csrc.nist.gov/publications/welcome.html

VPN-Related Information

IPSec Papers, RFCs www.ietf.cnri.reston.va.us/ids.by.wg/ipsec.html

L2F vs PPTP www.nortel.com/rapport/product/faqpptp.html

L2TP www.totalb.com/~12tp/

Point-to-Point Tunneling Protocol www.microsoft.com/communications/PPTP.htm

Automotive Network eXchange www.aiag.org/anx/

ANX Ottawa Test Results www.aiag.org/anx/ottawa.html

Dynamic Virtual Private Networks esp.tradewave.com/papers/securevpn.html

NewOak ROI Calculator m80.environs.com/newoak/roi/advanced.asp

VPN ROI Calculator www.baynetworks.com/products/switches/roi_calculator/index.html

Security Working Group News www.cs.arizona.edu/xkernel/www/ipsec/ipsec.html

S/WAN www.rsa.com/rsa/SWAN/home.html

SKIP IP-level Encryption skip.incog.com/

VPN Information Center www.checkpoint.com/vpn/index.html

VPN Source Page techweb.cmp.com/internetwk/VPN/default.html

Previous | Table of Contents | Next

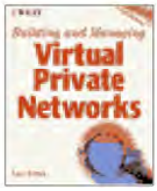

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

APPENDIX B
VPN Vendors and Products

For more information on new product and industry updates, please visit the book's companion Web site at www.wiley.com/compbooks/kosiur

3Com Corporation, 5400 Bayfront Plaza, Santa Clara, CA 95052, (408) 764-5000, www.3com.com

NETBuilder routers
Dual Processing Engine (DPE) for NETBuilder II routers
Dial Access Outsourcing (VPN bundle)
Virtual Leased Lines (VPN bundle)

Ascend Communications, Inc., 1701 Harbor Bay Parkway, Alameda, CA 94502, (510) 769-6001 or (800) 621-9578, info@ascend.com, www.ascend.com/

MultiVPN Architecture (IP Navigator, Customer Network Management Platform)
Pipeline 220 router
SecureConnect (family of VPN software products)

Assured Digital, Inc., P.O. Box 248, 9–11 Goldsmith Street, Littleton, MA 01460, (978) 486-0555, www.assured-digital.com/

ADI VPN-100 (remote dial-up software)
ADI VPN-500 (PC-based encryptor)
ADI VPN-1000 (VPN hardware, 10-Mbps Ethernet)
ADI VPN-2000 (VPN hardware, 10-Mbps Ethernet)
ADI VPN-4500 (VPN hardware, 100-Mbps Ethernet)
ADI Management System

Aventail Corporation, 117 South Main Street, Fourth Floor, Seattle, WA 98104, (206) 215-1111 or (888) 762-5785, sales@aventail.com, www.aventail.com/

Mobile VPN (server and client software)

Axent Technologies, Inc., 2400 Research Blvd., Rockville, MD 20850, (301) 258-5043, (800) 298-2620

ext. 801, fax: (301) 330-5756, info@axent.com, www.axent.com

PowerVPN software

Bay Networks, Inc., 4401 Great America Pkwy., Santa Clara, CA 95054, (800)-8-BAYNET, www.baynetworks.com/

Contivity (formerly Extranet Switch) -1000, -4000

VPN Secure Manager

VPN Secure Client software

VPN 500n, 550n (integrated VPN hardware)

BayStream Dial VPN Service (BayDVS)

CheckPoint Software Technologies Ltd., Three Lagoon Drive, Suite 400, Redwood City, CA 94065, (650) 628-2000, info@checkpoint.com, www.checkpoint.com/

Firewall-1

Firewall-1 SecuRemote (dial-in client software)

Floodgate-1 (bandwidth-management software)

Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134, (408) 526-4000 or (800) 553-6387, www.cisco.com/

PIX Firewall

IOS 11.2 (operating software for routers, switches, etc.)

Compatible Systems, 4730 Walnut Street, Suite 102, Boulder, CO 80301, (303) 444-9532 or (800) 356-0283, info@compatible.com, www.compatible.com/

IntraPort (self-contained VPN hardware)

Cylink Corporation, 910 Hermosa Court, Sunnyvale, CA 94086, (408) 735-5800 or (800) 533-3958, info@cylink.com, www.cylink.com/

SecureDomain (self-contained VPN hardware)

Data Fellows Inc., 675 N. First Street, 8th Floor, San Jose, CA 95112, (408) 938-6700, info@DataFellows.com, www.datafellows.com/

F-Secure VPN (VPN software for Intel PCs)

Digital Equipment Corp., Littleton, MA (978) 493-5111, fax: (978) 506-2017, altavista.software.digital.com/

AltaVista Tunnel (VPN software)

Entrust Technologies, 2323 North Central Expressway, Suite 360, Richardson, TX 95080, (972) 994-8000, entrust@entrust.com, www.entrust.com/

Entrust/Directory (LDAP-compatible certificate server)

Entrust/Manager

Extended Systems, Inc., 5777 N. Meeker Avenue, Boise, ID 83713, (208) 322-7800 or (800) 235-7576,

info@extendsys.com, www.extendsys.com/

ExtendNet VPN (remote access VPN server)

Fortress Technologies, Inc., 2701 N. Rocky Point Drive, Suite 650, Tampa, FL 33607, (813) 288-7388, info@fortresstech.com, www.fortresstech.com/

NetFortress VPN-1, VPN-3 (VPN hardware)

FreeGate Corporation, 1208 East Arques Avenue, Sunnyvale, CA 94086, (408) 617-1000, sales@freegate.com, www.freegate.com/

Multiservices Internet Gateway

Frontier Technologies, 10201 N. Port Washington Rd., Mequon, WI 53092, (414) 241-4555, (800) 929-3054, fax: 414-241-7084, Info@FrontierTech.Com

E-Lock Desktop (PKI and VPN client software)

E-Lock Director (PKI and VPN server software)

Indus River Networks, Inc., 31 Nagog Park, Acton, MA 01720, (978) 266-8100, fax: (978) 266-8111, www.indusriver.com/

Riverworks Enterprise VPN (hardware for dial-in VPNs)

RTS-5000 Tunnel Server

InfoExpress, Inc., 1270 Payne Dr., Los Altos, CA 94024, (650) 969-9609, fax: (650) 969-6924, info@infoexpress.com, sales@infoexpress.com for sales, www.infoexpress.com

VTPC/Secure

Information Resource Engineering, Inc. (IRE), 8029 Corporate Drive, Baltimore, MD 21236, (410) 931-7500, www.ire.com/

SafeNet/LAN (encrypting firewall for VPN)

SafeNet/Soft, SafeNet/Soft-PK (client software)

SafeNet/Smart (smart card-based VPN software)

SafeNet/Security Center (management workstation)

SafeNet/Dial (encrypting modem with smart card and software)

SafeNet/Firewall (proxy firewall)

SafeNet/Trusted Service (managed VPN service)

Intel Corp., Santa Clara, CA, www.intel.com/network/doc/1460/INDEX.COM

IntelExpress Router

IBM, www.software.ibm.com/enetwork/technology/vpn/

2210 Nways Multiprotocol Routers

Nways Multiprotocol Access Services

AIX Firewall

Internet Devices Inc., 1287 Anvilwood Avenue, Sunnyvale, CA 94089, (408) 541-1400 or (888) 237-2244, sales@InternetDevices.com, www.InternetDevices.com/

Fort Knox (integrated VPN hardware)

Internet Dynamics, Lombard, IL, (630) 953-7700, fax: (630) 953-7701, sales@interdyn.com, www.interdyn.com/

Conclave (integrated VPN software)

Microsoft Corp., Redmond, WA, (425) 882-8080, fax: (425) 936-7329, www.microsoft.com/

NT Server with Routing and Remote Access Services (software)

Proxy Server

Milkyway Networks Corp., 2650 Queensview Dr., Suite 150, Ottawa, ON, CANADA, K2B 8H6, (613) 596-5549, fax: (613) 596-5615, www.milkyway.com/

SecurIT Firewall (encrypting firewall)

SecurIT Access (remote access software)

NEC Systems Laboratory, Inc., 110 Rio Robles Drive, San Jose, CA 95134, prod-info@socks5.nec.com, www.socks5.nec.com

SOCKS5 E2 Client

SOCKS5 Internet Access Management Framework

NetScreen Technologies Inc., 4699 Old Ironsides Drive, Suite 300, Santa Clara, CA 95054, (408) 970-8889 or (877) NETSCREEN, info@netscreen.com, www.netscreen.com/

NS-10, NS-100 (integrated VPN hardware)

Novell Inc., Provo, UT (801) 861-5588 or (800) 638-9273, fax: (801) 861-5155, www.novell.com/bordermanager/

Bordermanager (integrated software including VPN functions)

RADGUARD, 24 Raoul Wallenberg Street, Tel Aviv 69719 ISRAEL, +972 3 645 5444; fax: +972 3 648 0859, www.radguard.com/

cIPro-VPN (VPN hardware)

NetCryptor (hardware encryptor)

CryptoCA (certificate server)

CryptoManage

Raptor Systems, Inc., 266 Second Avenue, Waltham MA, 02154, (800) 9-EAGLE-6, (781) 530-2200, fax: 781-487-6755, info@raptor.com, www.raptor.com

Eagle (firewall)

RedCreek Communications, Inc., 3900 Newpark Mall Rd., Newark, CA 94560, (510) 745-3900, fax: (510) 745-3999, www.redcreek.com/

Ravlin-4, -10, -100 (VPN hardware)

Secure Computing Corporation, One Almaden Blvd., Suite 400, San Jose, CA 95113, (408) 918-6100 or (800) 379-4944, www.securecomputing.com/

BorderWare (firewall)

SideWinder Security Server

SecureZone (firewall)

Security Dynamics Technologies, Inc., 20 Crosby Drive, Bedford, MA 01730, (800) SECURID, www.securid.com/index.html

SecurID (token-based authentication system)

Shiva Corporation, 28 Crosby Drive, Bedford, MA 01730-1437, (781) 687-1000, fax: (781) 687-1001, sales@shiva.com, www.shiva.com/

LANRover VPN Gateway

VPN client software

Certificate Authority software

Storage Technology Corporation, 2270 South 88th Street, Louisville, CO 80028, (800) STORTEK or (800) 786-7835, www.network.com/

NetSentry BorderGuard (VPN software)

Sun Microsystems, Inc., Palo Alto, CA, www.sun.com/

Sunscreen EFS (firewall)

TimeStep Corporation, 362 Terry Fox Drive, Kanata, Ontario, Canada K2K-2P5, (613) 599-3610 ext. 4532, info@timestep.com, www.timestep.com

PERMIT/Connect (VPN hardware and Entrust certificate server system)

PERMIT/Gateway 4520 (VPN hardware)

PERMIT/Client (remote access software)

PERMIT/Config

PERMIT/2505, 4505 (hardware-based firewalls)

Trusted Information Systems, 1-888-TISFIRST or 1-888-FIREWALL, sales@tis.com, www.tis.com/

Gauntlet (firewall)

UAC, 200 Lincoln Street, Suite 201, Boston, MA 02111, (617) 695-0137 ext. 19, www.uac.com/uacpn7.htm

PN7 (firewall)

V-ONE Corporation, 20250 Century Blvd., Suite 300, Germantown, MD 20874, (301) 515-5200, sales@v-one.com, www.v-one.com/

SmartGate Enterprise (VPN server)

SmartPass (token-based remote access)

SmartAdmin (management software)

SmartWall (firewall)

VPNet Technologies, Inc., 1530 Meridian Avenue, San Jose, CA 95125, (408) 445-6600 or (888) VPNET-88, sales@vpnet.com, www.vpnet.com/

VSU-10, -1010 (integrated VPN hardware)

VPNmanager (management software)

VPNywhere (VPN hardware with roaming services)

Watchguard Technologies, Inc., 316 Occidental Avenue S., Suite 200, Seattle, WA 98104, (206) 521-8340, fax: (206) 521-8342, www.watchguard.com/

Firebox II

Global Security Manager

Watchguard Security System (Firebox II, management and authentication tools)

Commercial VPN Providers

ANS, www.ans.net/

AT&ampT WorldNet VPN Services, (800) 831-5259, www.att.com/worldnet/wmis/virtual.html

Compuserve Network Services, Columbus Center, 6550 Metro Place South, Suite 560, Dublin, OH 43017, (614) 792-1901, networkinfo@csi.compuserve.com

Concentric Network, www.concentric.net/business/vpns/

GRIC Communications (roaming service), 1421 McCarthy Blvd., Milpitas, CA 95035, (408) 955-1920, gricinfo@gric.com, www.gric.com/

GTE Internetworking, 150 Cambridge Park Dr., Cambridge, MA 02140, (617) 873-2000, fax: (617) 873-5011, www.bbn.com, www.gte.net

iPass Inc. (roaming service), 650 Castro Street, Suite 500, Mountain View, CA 94041, (650) 237-7300, fax: (650) 237-7321, www.ipass.com

Netcom, www.netcom.com/

networkMCI, www.networkmci.com/Contactus/index.html

TCG CERFnet, P.O. Box 919014, San Diego, CA 92191-9014, (800) 876-CERF (2373), (619) 812-5000, fax: (619) 812-3990, www.cerfnet.net

UUNET (Extralink), Fairfax, VA, www.uunet.net/

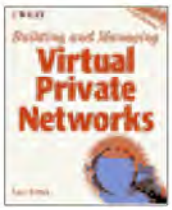

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

APPENDIX C
What's on the Web Site?

The companion Web site to "Building and Managing Virtual Private Networks" is located at www.wiley.com/compbooks/kosiur. The site aims to be your central source of information on VPNs and includes the following sections:

- Information about new VPN products
- Pointers to published reviews of VPN products and services
- Pointers to information on VPN protocols (IETF RFCs and Internet-drafts)
- Information on the latest VPN-related tests and reports from ANX and ICSA
- Pointers to Web sites on security, network management, and VPNs
- Feature checklists of available VPN products
- A directory of VPN vendors
- Corrections and addenda to the book

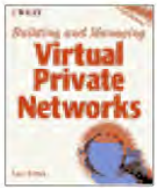

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

Glossary

aggressive mode

In Oakley, the name of a mechanism used in the first phase of establishing a security association. In aggressive mode, the sender and receiver negotiate the basic algorithms and hashes for remaining SA exchanges. Unlike main-mode exchanges, aggressive mode accomplishes the exchange in three packets rather than six.

application gateway (or application proxy)

A type of firewall that controls external access to applications within a network.

ARPANET

The predecessor of the Internet, ARPANET was the first wide-area data communications network. It was originally funded by the Department of Defense's *Advanced Research Project Agency* (ARPA) to provide nationwide connectivity among military, educational, and research sites.

authentication

In cryptography, the process of ensuring that the data is coming from the source it claims to come from.

Authentication Header (AH)

In IPSec, the IP header used to verify that the contents of a packet haven't been altered.

automatic rekeying

The process of changing a key used for encryption periodically without any manual intervention. If rekeying intervals are kept short, automatic rekeying can help defeat attacks because the amount of data subject to the attack is relatively small, and the attacker has less time in which to crack the key.

biometrics

Using a unique physical trait to identify the user. Biometric technologies measure human characteristics such as fingerprints, voice recordings, iris and retinal scans, heat patterns, facial images, and even keystroke patterns.

block cipher

A crypto algorithm that encrypts data in blocks of a fixed size.

blowfish

A 64-bit block cipher with a variable-length key designed by Bruce Schneier for implementation

on large microprocessors. It's optimized for applications in which the key does not change often.

brute force attack

The process of trying to recover a cryptographic key by trying all reasonable possibilities.

Certificate Authority (CA)

A trusted company or organization that will accept your public key, along with some proof of your identity, and serve as a repository of digital certificates. Others then can request verification of your public key from the certificate authority.

Certificate Revocation List (CRL)

Certificate authorities must maintain a list of digital certificates that are no longer valid (not including those expired).

Challenge Handshake Authentication Protocol (CHAP)

A protocol for authenticating remote users. CHAP incorporates three steps to produce a verified link after the link is first initiated. Instead of a simple two-step password/approval process, CHAP uses a one-way hashing function.

challenge-response

An authentication mechanism in which the authentication process sends a challenge to a process that requests authentication; the latter is authenticated only if it sends the correct response to the authentication process.

Channel Service Unit/Data Service Unit (CSU/DSU)

An interface between a digital line and a communication device used to provide the interface for circuit data services, which includes the physical framing, clocking, and channelization of the circuit.

Cipher Block Chaining (CBC)

A block cipher mode that combines the previous block of cipher text with the current block of plain text before encrypting it.

circuit gateway (circuit proxy)

A type of firewall that forms circuit-level connections between an external computer and a computer inside the network.

compulsory tunnel

Tunnels created without the user's consent, which may be transparent to the end user. The client-side endpoint of a compulsory tunnel typically resides on a *remote access server* (RAS). 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.

confidentiality

Preventing anyone from reading or copying your data as it travels across the Internet.

Data Encryption Standard (DES)

A block-cipher algorithm created by IBM and endorsed by the U.S. government in 1977; uses a 56-bit key and operates on blocks of 64 bits; relatively fast and used to encrypt large amounts of data at one time.

Demilitarized Zone (DMZ)

A portion of a network in which the traffic is not yet screened or regulated; commonly delineated by two firewalls: one forming the boundary between the public network and the DMZ and the other forming the boundary between the DMZ and the internal network.

Diffie-Hellman

A system designed to allow two individuals to agree on a shared key, even though they only exchange messages in public. This oldest public-key cryptosystem is still in use but does not support either encryption or digital signatures.

Digital Certificate

An electronic document, issued by a certificate authority, that's used to establish a company's identity by verifying its public key. (See Public Key Certificate.)

Digital Data Service (DDS)

DDS was the first digital service for private line applications, offering 56-Kbps connections to corporate customers. (Also called Dataphone Digital Service.)

Digital Signature Algorithm (DSA)

Developed by NSA and based on what's called the El Gamal algorithm, the signature scheme uses the same sort of keys as Diffie-Hellman and can create signatures faster than RSA. DSA is being pushed by NIST as DSS, the Digital Signature Standard.

Domain Name Service (DNS)

The network service responsible for converting numeric IP addresses into text-based names.

Domain of Interpretation (DOI)

In IPSec, the specification of which protocols and parameters are required for negotiation of security association.

Encapsulating Security Payload (ESP)

In IPSec, an IP header that contains the encrypted contents of an IP packet.

encapsulation

Placing the contents of one network's packet into that of another network's packet. (The protocols for the two networks can be identical.)

encryption

Conversion of human-readable cleartext (or plaintext) to ciphertext, using cryptographic algorithms.

firewall

A device acting as a network filter to restrict access to a private network from the outside, implementing access controls based on the contents of the packets of data that are transmitted between two parties or devices on the network.

frame relay

A high-speed, connection-oriented, public data packet switching technology that provides a very reliable and efficient packet delivery over *virtual circuits* (VCs). It supports access speeds up to 1.544 Mbps (T1) or 2.048 Mbps (E1) in Europe. The basic transport unit, which is called frame, can be up to 4,096 bytes and carries both routing and user information.

HMAC-MD5

A message authentication algorithm coupled with the MD5 hash function; operates on 64-byte blocks of data and produces a 128-bit authentication value.

HMAC-SHA-1

A message authentication algorithm coupled with the SHA-1 hash function; operates on 64-byte blocks of data and produces a 160-bit authentication value.

International Data Encryption Algorithm (IDEA)

A cryptographic algorithm using a 128-bit key for strong encryption and designed to be efficient to compute in software.

Internet Architecture Board (IAB)

The primary decision-making body regarding the Internet. The IAB sets research and engineering directions and oversees the IETF.

Internet Assigned Number Authority (IANA)

The organization responsible for assigning Internet address blocks, protocol identifiers, and TCP/UDP port numbers.

Internet Engineering Task Force (IETF)

A worldwide organization that develops new technology and standards for the Internet.

Internet Key Exchange (IKE)

The key-management protocol used in conjunction with IPSec.

Internet Service Provider (ISP)

A company that provides Internet access services to individual users and businesses.

IPSec

The network cryptographic protocols for protecting IP packets.

ISAKMP

A key-management protocol accepted for use with IPSec; now combined with Oakley to form the Internet Key Exchange (IKE) protocol.

jitter

The variation in latency. In traditional terms, it is the variation between voice samples generating distortion in the delivered voice signal. The distortion of a signal as it is propagated through a network, in which the signal varies from its original reference timing. In packet-switched networks, jitter is the distortion of the interpacket arrival times as compared to the interpacket times of the original transmission.

key

A string of digits, which when used with a cryptographic algorithm, produces cipher text.

LAN-to-LAN tunnel

A VPN tunnel created between two security gateways, each of which serves as the interface between the LAN it's protecting and the public network.

latency

Network delay; the minimum time that elapses between requesting and receiving data.

Layer2 Forwarding (L2F)

A tunneling protocol originally developed by Cisco.

Layer2 Tunneling Protocol (L2TP)

A tunneling protocol that combines many of the features of L2F and PPTP; also uses IPSec for encryption; supports encapsulation of packets other than IP.

leased line

A dedicated, private line provided by a carrier or local telephone company for the exclusive use of the customer.

Lightweight Directory Access Protocol (LDAP)

An IP-based protocol that governs how information within X.500-format directories can be obtained.

main mode

In Oakley, the name of the mechanism used in the first phase of establishing a security association. In main mode, the sender and receiver negotiate the basic algorithms and hashes for remaining SA exchanges.

Network Access Point (NAP)

On-ramp to the high-speed Internet backbone maintained by Sprint, Ameritech, Worldcom, and others.

Network Address Translation (NAT)

A procedure for translating private IP addresses used on an internal network to a special reserved block of IP addresses for communications on the public Internet.

nonce

A random value sent in a communications protocol exchange; often used to detect replay attacks.

Oakley

A key exchange protocol used in IPSec as part of the Internet Key Exchange protocol.

one-way hash function

A formula used to convert a message of any length into a string of digits called a message digest. The length of the function determines the length of the digest, and no key is required.

packet filter

Hardware or software that discards packets based on the contents of the packet; used in firewalls.

Password Authentication Protocol (PAP)

A simple authentication protocol that uses passwords using a two-way handshake; passwords are sent in the clear and, therefore, are not secure.

Permanent Virtual Circuit (PVC)

A *virtual connection* (VC) established by the network management between a source and a destination, which can be left up permanently (used in X.25 and frame relay protocols).

Point-of-Presence (POP)

Local access point to a national or international communications network. Users dial into their networks by calling a local phone number, rather than a toll-free number or a long-distance call to a centralized location.

Point-to-Point Protocol (PPP)

An Internet data-link protocol used to frame data packets on point-to-point links, such as modem links.

Point to Point Tunneling Protocol (PPTP)

A tunneling protocol originally developed by Ascend and Microsoft that operates at the Link layer of a network. It depends on PPP for its basic functionality and uses *Generic Routing Encapsulation* (GRE) for encapsulating packets; supports encapsulation of packets other than IP.

proxy agent

A proxy server software module that's programmed to handle one specific type of data transfer (e.g., FTP or TCP).

proxy server

A type of firewall that employs a store-and-forward approach to protecting crucial data and applications. The proxy server terminates the incoming connection from the source and initiates a second connection to the destination, ensuring that the incoming user has appropriate access rights to use data requested from the destination before passing that data on to the user.

Public Key Certificate

Specially-formatted data blocks that tell us the value of a public key, the name of the key's owner, and a digital signature of the issuing organization. These certificates are used to identify the owner of a particular public key.

public-key cryptography

An encryption method that uses a pair of keys: one public and one private. Messages encoded with either key can be decoded by the other. Also called *asymmetric encryption.*

Public Key Infrastructure (PKI)

The organization of certificate issuers and certificate management processes.

Public Switched Telephone Network (PSTN)

A generic name for the worldwide public telephone network.

Quality of Service (QoS)

A term used to describe a set of performance parameters that characterize the transmission quality over a given connection.

quick mode

In Oakley, the name of the mechanism used after a security association has been established to negotiate changes in security services, such as new keys.

RC2

Designed by Ron Rivest. A variable key-size cipher for very fast bulk encryption. RC2 is a block cipher and can be used in place of DES.

RC4

Another variable key-size cipher designed by Ron Rivest for very fast bulk encryption. RC4 is a stream cipher and is as much as 10 times faster than DES.

Remote Authentication Dial-In User Service (RADIUS)

A protocol that uses a client/server model to securely authenticate and administer remote network connection users and sessions. It can support other types of user authentication, including PAP and CHAP.

Resource reSerVation Protocol (RSVP)

A control protocol developed for supporting different QoS classes for IP applications (such as videoconference and multimedia) and to reserve resources in an IP-based network. RSVP uses a soft state mechanism to maintain path and reservation state in each node along the reservation path and can be changed dynamically by the requesting host.

roaming service

A service for remote users that enables them to use a local ISP instead of their corporate-selected ISP. A broker service manages the roaming service, handling settlement charges between ISPs and distribution of client software.

RSA

Named after Rivest, Shamir, and Adelman, its designers. Public-key algorithm supports a variable key length as well as variable blocksize of the text to be encrypted. The plaintext block must be smaller than the key length. Common key length is 512 bits.

Secure Sockets Layer (SSL)

A protocol that provides authentication for servers and browsers as well as confidentiality and data integrity for communications between a Web server and a browser. SSL can be used for transactions other than those on the Web, but it's not designed to handle security decisions based on authentication at the application or document level.

Security Association (SA)

In IPSec, an agreement between two communicating parties on which authentication and encryption algorithms will be used, along with related data, such as key lifetimes.

security gateway

Security gateways sit between public and private networks, preventing unauthorized intrusions into the private network. They also may provide tunneling capabilities and encrypt private data before it's transmitted on the public network.

Security Parameter Index (SPI)

In IPSec, specifies to the device receiving the packet what group of security protocols the sender is using for communications.

Service Level Agreement (SLA)

A contract between a service provider and a customer that specifies the parameters defining acceptable network performance to be provided and the type of remuneration for failing to meet the guaranteed performance.

session key

A cryptographic key intended to encrypt data for a limited period of time, typically only for a single communications session between a pair of correspondents. When the session is over, the key is discarded and a new one established when a new session takes place.

Skipjack

The NSA-developed encryption algorithm designed for the Clipper, Capstone, and Fortezza

systems. The algorithm is an iterative 64-bit block cipher with an 80-bit key.

smart card

A credit card-sized plastic card with a special type of integrated circuit embedded in it. The integrated circuit holds information in electronic form and controls who uses this information and how.

SOCKS

An application proxy protocol that passes only traffic processed from specific (i.e., SOCKS) clients.

Stateful Multi-Layer Inspection (SMLI)

A type of firewall mechanism that uses smart packet filters that compare each packet with bit patterns of similar friendly packets.

Symmetric Encryption

An encryption method in which both the sender and the receiver possess the same cryptographic key, which means that both parties can encrypt and decrypt data with that key.

T1

A WAN transmission circuit that carries DS1-formatted data at a rate of 1.544 Mbps over two twisted pair wiring.

T3

A WAN transmission circuit that carries DS3-formatted data at a rate of 44.736 Mbps.

Terminal Access Controller Access Control System (TACACS)

A protocol that uses a client/server model to securely authenticate and administer remote network connection users and sessions. It can support other types of user authentication, including PAP and CHAP.

token-based authentication

A system using a hardware device that generates a one-time password to authenticate its owner; occasionally used to describe software programs that generate one-time passwords.

Transport-Mode IPSec

An IPSec mode, used either in AH or ESP, that leaves the original IP addresses in plaintext.

Triple DES

Based on DES, this cryptographic algorithm encrypts a block of data three times with two (or three) different keys.

Tunnel-Mode IPSec

An IPSec mode, used either in AH or ESP, that encrypts the original IP addresses of the source and destination and uses the IP addresses of the security gateways to route the packet through an IPSec tunnel.

tunnel or tunneling

The process of encapsulating one type of packet in another packet type so that the data can be transferred across paths that otherwise would not transmit the data. To avoid any confusion with the media-dependent virtual circuits, the paths that the encapsulated packets follow in Internet VPNs are called tunnels, not virtual circuits.

Virtual Private Network (VPN)

A private network built atop a public network (in this book, the Internet), in which secure connections are set up dynamically between a sender and a receiver.

voluntary tunnel

Voluntary tunnels are set up at the request of the end-user. When using a voluntary tunnel, it is possible for the end-user to simultaneously 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.

Wide Area Network (WAN)

A network environment in which the elements of the network are located at significant distances from each other, and the communications facilities typically use carrier facilities rather than private wiring. Typically, a routing protocol is required to support communications between two distant host systems on a WAN.

X.509

A specification for public-key certificates, originally developed as part of the CCITT's X.500 directory specification.

Previous | Table of Contents | Next

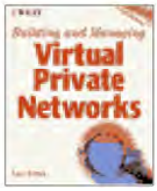

Building and Managing Virtual Private Networks
by Dave Kosiur
Wiley Computer Publishing, John Wiley & Sons, Inc.
ISBN: 0471295264 **Pub Date:** 09/01/98

INDEX

Page references in italic type indicate illustrations. Numbers are treated as if spelled out. Thus "T1 lines" would be found as if spelled "Tone lines," and "L2F" and "L2TP" as if spelled LtwoF and LtwoTP respectively.

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X