After completing this chapter, you will be able to perform the following tasks:

- Understand network security
- Understand VPN technologies
- Use the Cisco Security Wheel
- Understand the basics of the IPSec protocol framework

CHAPTER 1

Network Security and Virtual Private Network Technologies

This opening chapter provides an overview of network security and looks at the Cisco Architecture for Voice, Video, and Integrated Data (AVVID) and the SAFE blueprint. It also covers the IP Security (IPSec) framework and identifies the main encryption and algorithm protocols. Then it looks at how IPSec works before finishing with the five steps of IPSec operation. These five steps are very important to remember and also are very useful for implementing and troubleshooting any IPSec-based virtual private network (VPN), whether firewall-, router-, or VPN Concentrator-based.

Network Security Overview

Network security is essential because the Internet is a network of interconnected networks without a boundary. Because of this fact, the organizational network becomes accessible from and vulnerable to any other computer in the world. As companies become Internet businesses, new threats arise because people no longer require physical access to a company's computer assets: They can access everything over the public network.

In a recent survey conducted by the Computer Security Institute (CSI, www.gocsi.com), 70 percent of the organizations polled stated that their network security defenses had been breached and that 60 percent of the incidents came from within the organizations themselves.

Network security faces four primary threats:

- Unstructured threats
- Structured threats
- External threats
- Internal threats

Unstructured Threats

Unstructured threats consist of mostly inexperienced individuals using easily available hacking tools from the Internet. Some of the people in this category are motivated by malicious intent, but most are motivated by the intellectual challenge and are commonly called *script kiddies*. They are not the most talented or experienced hackers, but they have the motivation, which is all that matters.

Structured Threats

Structured threats come from hackers who are more highly motivated and technically competent. They usually understand network system designs and vulnerabilities, and they can understand as well as create hacking scripts to penetrate those network systems.

External Threats

External threats are individuals or organizations working outside your company who do not have authorized access to your computer systems or network. They work their way into a network mainly from the Internet or dialup access servers.

Internal Threats

Internal threats occur when someone has authorized access to the network with either an account on a server or physical access to the wire. They are typically disgruntled former or current employees or contractors.

The three types of network attacks are

- Reconnaissance attacks
- Access attacks
- Denial of service (DoS) attacks

Reconnaissance Attacks

Reconnaissance is the unauthorized discovery and mapping of systems, services, or vulnerabilities. It is also called information gathering. In most cases, it precedes an actual access or DoS attack. The malicious intruder typically ping-sweeps the target network first to determine what IP addresses are alive. After this is accomplished, the intruder determines what services or ports are active on the live IP addresses. From this information, the intruder queries the ports to determine the application type and version as well as the type and version of the operating system running on the target host.

Reconnaissance is somewhat analogous to a thief scoping out a neighborhood for vulnerable homes he can break into, such as an unoccupied residence, an easy-to-open door or window, and so on. In many cases, an intruder goes as far as "rattling the door handle"—not to go in immediately if it is open, but to discover vulnerable services he can exploit later when there is less likelihood that anyone is looking.

Access Attacks

Access is an all-encompassing term that refers to unauthorized data manipulation, system access, or privilege escalation. Unauthorized data retrieval is simply reading, writing, copying, or moving files that are not intended to be accessible to the intruder. Sometimes this is as easy as finding shared folders in Windows 9*x* or NT, or NFS exported directories in UNIX systems with read or read-write access to everyone. The intruder has no problem getting to the files. More often than not, the easily accessible information is highly confidential and completely unprotected from prying eyes, especially if the attacker is already an internal user.

System access is an intruder's ability to gain access to a machine that he is not allowed access to (such as when the intruder does not have an account or password). Entering or accessing systems that you don't have access to usually involves running a hack, script, or tool that exploits a known vulnerability of the system or application being attacked.

Another form of access attacks involves privilege escalation. This is done by legitimate users who have a lower level of access privileges or intruders who have gained lower-privileged access. The intent is to get information or execute procedures that are unauthorized at the user's current level of access. In many cases this involves gaining root access in a UNIX system to install a sniffer to record network traffic, such as usernames and passwords, that can be used to access another target.

In some cases, intruders only want to gain access, not steal information—especially when the motive is intellectual challenge, curiosity, or ignorance.

DoS Attacks

DoS is when an attacker disables or corrupts networks, systems, or services with the intent to deny the service to intended users. It usually involves either crashing the system or slowing it down to the point where it is unusable. But DoS can also be as simple as wiping out or corrupting information necessary for business. In most cases, performing the attack simply involves running a hack, script, or tool. The attacker does not need prior access to the target, because usually all that is required is a way to get to it. For these reasons and because of the great damaging potential, DoS attacks are the most feared—especially by e-commerce website operators.

Network Security as a Continuous Process

Network security should be a continuous process built around a security policy. A continuous security policy is most effective, because it promotes retesting and reapplying updated security measures on a continuous basis. The Security Wheel, shown in Figure 1-1, represents this continuous security process.

Figure 1-1 *Security Wheel*

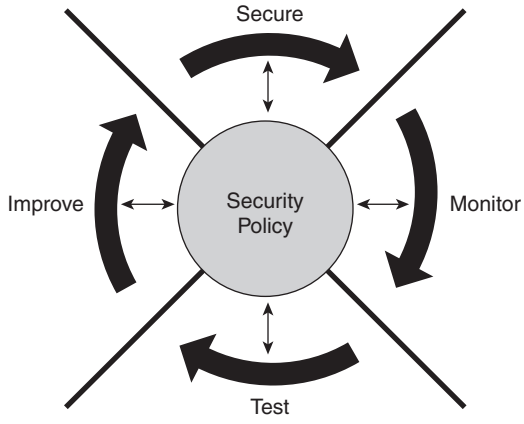

To begin this continuous process known as the Security Wheel, you need to create a security policy that enables the application of security measures. A security policy needs to accomplish the following tasks:

- Identify the organization's security objectives
- Document the resources to be protected
- Identify the network infrastructure with current maps and inventories

To create or implement an effective security policy, you need to determine what it is you want to protect and in what manner you will protect it. You should know and understand your network's weak points and how they can be exploited. You should also understand how your system normally functions so that you know what to expect and are familiar with how the devices are normally used. Finally, consider your network's physical security and how to protect it. Physical access to a computer, router, or firewall can give a user total control over that device.

After the security policy is developed, it becomes the hub on which the next four steps of the Security Wheel are based:

Step 1 Secure the system. This involves implementing security devices—firewalls, identification authentication systems, encryption, and so on—with the intent to prevent unauthorized access to network systems. This is where the Cisco PIX Firewall is effective.

Step 2 Monitor the network for violations and attacks against the corporate security policy. These attacks can occur within the network's secured perimeter—from a disgruntled employee or contractor—or from a source outside your trusted network. You should monitor the network with a real-time intrusion detection device such as the Cisco Intrusion

Detection System (IDS). This helps you discover unauthorized entries. It also serves as a system of checks and balances to ensure that devices implemented in Step 1 of the Security Wheel have been configured and are working properly.

Step 3 Test the effectiveness of the security safeguards that are in place. Use the Cisco Secure Scanner to identify the network's security posture with respect to the security procedures that form the hub of the Security Wheel. Validation is a must. You can have the most sophisticated network security system, but if it is not working, your network can be compromised. This is why you need to test the devices you implemented in Steps 1 and 2 to make sure they are functioning properly. The Cisco Secure Scanner is designed to validate your network security.

Step 4 Improve corporate security. The improvement phase of the Security Wheel involves analyzing the data collected during the monitoring and testing phases and developing and implementing improvement mechanisms that feed into your security policy and the securing phase in Step 1. If you want to keep your network as secure as possible, you must keep repeating the cycle of the Security Wheel, because new network vulnerabilities and risks are created every day.

All four steps—secure, monitor, test, and improve—should be repeated on a continuous basis and should be incorporated into updated versions of the corporate security policy.

Cisco AVVID

The Internet is creating tremendous business opportunities for Cisco and Cisco customers. Internet business solutions such as e-commerce, supply chain management, e-learning, and customer care are dramatically increasing productivity and efficiency.

Cisco AVVID is the one enterprise architecture that provides the intelligent network infrastructure for today's Internet business solutions. As the industry's only enterprise-wide, standards-based network architecture, Cisco AVVID provides the road map for combining customers' business and technology strategies into one cohesive model.

With Cisco AVVID, customers have a comprehensive road map for enabling Internet business solutions and creating a competitive advantage. Cisco AVVID has four benefits:

- **Integration**—By leveraging the Cisco AVVID architecture and applying the network intelligence that is inherent in IP, companies can develop comprehensive tools to improve productivity.
- **Intelligence**—Traffic prioritization and intelligent networking services maximize network efficiency for optimized application performance.
- **Innovation**—Customers can adapt quickly in a changing business environment.

- **Interoperability**—Standards-based application programming interfaces (APIs) enable open integration with third-party developers, providing customers with choice and flexibility.

Combining the network infrastructure and services with new-world applications, Cisco AVVID accelerates the integration of technology strategy with business vision.

The following sections discuss the different parts of the Cisco AVVID architecture. They are shown in Figure 1-2.

Figure 1-2 *Cisco AVVID Architecture*

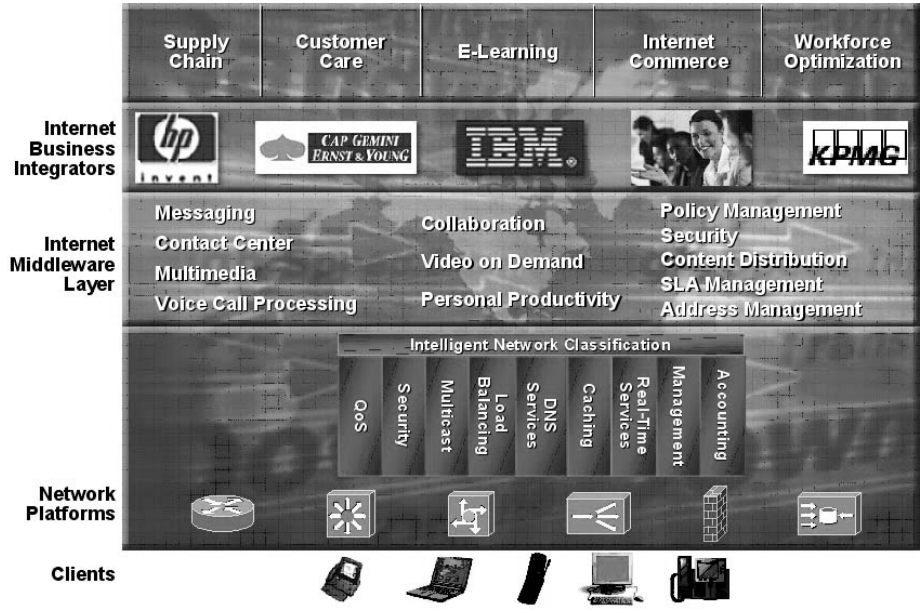

Clients

Clients are the wide variety of devices that can be used to access the Internet business solutions through the network. These can include phones, PCs, personal digital assistants (PDAs), and any other mobile Internet device. One key difference from traditional proprietary architectures is that the Cisco AVVID standards-based solution lets a wide variety of devices be connected—even some not yet in broad use. Unlike traditional telephony and video solutions, proprietary access devices are unnecessary. Instead, functionality is added through the intelligent network services provided in the infrastructure.

Network Platforms

The network infrastructure provides the physical and logical connection for devices, bringing them into the network. Network platforms are the LAN switches, routers, gateways, and other equipment that interconnect users and servers. Cisco network platforms are competitive as far as features, performance, and price, but their key capabilities are the integration and interaction with other elements of the Cisco AVVID framework. This layer of Cisco AVVID is the foundation for all applications that are integrated to solve business problems.

Intelligent Network Services

The intelligent network services, provided through software that operates on network platforms, are a major benefit of an end-to-end architecture for deploying Internet business solutions. From quality of service (QoS) (prioritization) through security, accounting, and management, intelligent network services reflect the enterprise's business rules and policies in network performance. A consistent set of the services end-to-end through the network is vital if the infrastructure is to be relied on as a network utility. These consistent services allow new Internet business applications and e-business initiatives to roll out very quickly without a major reengineering of the network each time. In contrast, networks built on best-of-breed strategies might promise higher performance in a specific device, but they cannot be counted on to deliver these sophisticated features end-to-end in a multivendor environment. Cisco AVVID supports standards to provide for migration and the incorporation of Internet business integrators, but the added intelligent network services offered by an end-to-end Cisco AVVID solution go far beyond what can be achieved in a best-of-breed environment.

Internet Middleware Layer

The next section, including service control and communication services, is a key part of any networking architecture, providing the software and tools to break down the barriers of complexity arising from new technology. These combined layers provide the tools for integrators and customers to tailor their network infrastructure and customize intelligent network services to meet application needs. These layers manage access, call setup and teardown, perimeter security, prioritization and bandwidth allocation, and user privileges. Software, such as distributed customer contact suites, messaging solutions, and multimedia and collaboration provide capabilities and a communication foundation that enable interaction between users and a variety of application platforms. In a best-of-breed strategy, many of these capabilities must be individually configured or managed. In traditional proprietary schemes, vendors dictated these layers, limiting innovation and responsiveness.

Rapid deployment of Internet business solutions depends on consistent service control and communication services capabilities throughout the network. These capabilities are often delivered by Cisco from servers distributed throughout the network. The service control and

communication services layers are the glue that joins the Internet technology layers of the Cisco AVVID framework with the Internet business solutions. In effect, this tunes the network infrastructure and intelligent network services to the needs of the Internet business solutions. In turn, the Internet business solutions are adapted for the best performance and availability on the network infrastructure by exploiting the end-to-end services available through the Cisco AVVID framework.

Internet Business Integrators

As part of the open ecosystem, it is imperative to enable partners with Cisco AVVID. Cisco realizes the crucial requirement to team with integrators, strategic partners, and customers to deliver complete Internet business. Cisco AVVID offers a guide for these interactions by describing a consistent set of services and capabilities that form the basis of many types of partner relationships.

Internet Business Solutions

Enterprise customers are deploying Internet business solutions to reengineer their organizations. The applications associated with Internet business solutions are not provided by Cisco, but they are enabled, accelerated, and delivered through Cisco AVVID. Being able to move their traditional business models to Internet business models and to deploy Internet business solutions is key to companies' survival. Cisco AVVID is the architecture on which e-businesses build Internet business solutions that can be easily deployed and managed. Ultimately, the more Internet business solutions that are delivered, the more efficiently and effectively companies will increase productivity and added value.

Cisco SAFE Blueprint

SAFE is a flexible, dynamic security blueprint for networks that is based on Cisco AVVID. SAFE lets businesses securely and successfully take advantage of e-business economies and compete in the Internet economy.

As the leader in networking for the Internet, Cisco is ideally positioned to help companies secure their networks. The SAFE blueprint, in conjunction with an ecosystem of best-of-breed, complementary products, partners, and services, ensures that businesses can deploy robust, secure networks in the Internet age.

Implementing the SAFE blueprint for secure e-business has several major benefits:

- It provides the foundation for migrating to secure, affordable, converged networks.
- It lets companies cost-effectively deploy a modular, scalable security framework in stages.

- It delivers integrated network protection via high-level security products and services.

The SAFE blueprint provides a robust security blueprint that builds on Cisco AVVID. SAFE layers are incorporated throughout the Cisco AVVID infrastructure:

- **Infrastructure layer**—Intelligent, scalable security services in Cisco platforms, such as routers, switches, firewalls, IDSs, and other devices
- **Appliances layer**—Incorporates key security functionality in mobile handheld devices and remote PC clients
- **Service control layer**—Critical security protocols and APIs that let security solutions work together cohesively
- **Applications layer**—Host and application-based security elements that ensure the integrity of critical e-business applications

To facilitate rapidly deployable, consistent security throughout the enterprise, SAFE consists of modules that address the distinct requirements of each network area. By adopting a SAFE blueprint, security managers do not need to redesign the entire security architecture each time a new service is added to the network. With modular templates, it is easier and more cost-effective to secure each new service as it is needed and to integrate it with the overall security architecture. Figure 1-3 shows an example of the module approach.

Figure 1-3 *Cisco SAFE Modular Blueprint*

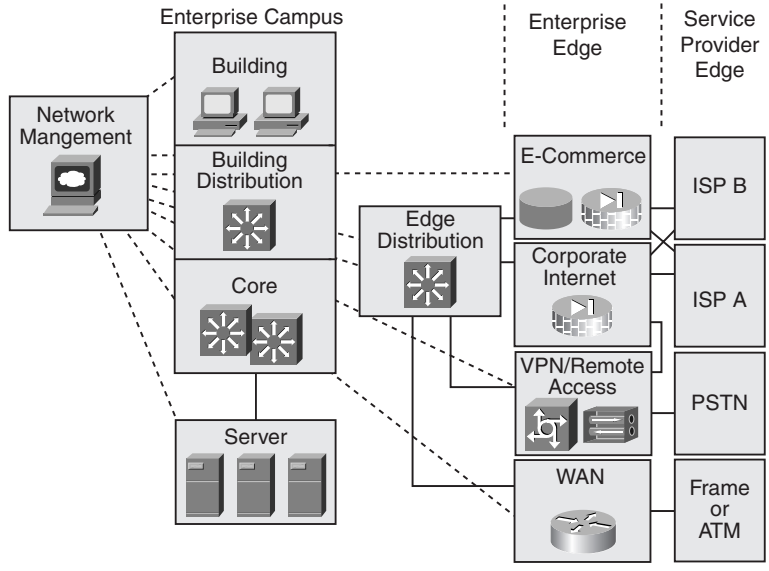

One of the unique characteristics of the SAFE blueprint is that it is the first industry blueprint that recommends exactly which security solutions should be included in which sections of the network, and why they should be deployed. Each module in the SAFE blueprint is designed specifically to provide maximum performance for e-business while allowing enterprises to maintain security and integrity.

Cisco has opened its Cisco AVVID architecture and SAFE blueprint to key third-party vendors to create a security solutions ecosystem to spur development of best-in-class multi-service applications and products. The Cisco AVVID architecture and SAFE blueprint provide interoperability for third-party hardware and software using standards-based media interfaces, APIs, and protocols. This ecosystem is offered through the Security and Virtual Private Network Associate Program, an interoperability solutions program that provides Cisco customers with complimentary tested and certified products for securing their businesses. The ecosystem lets businesses design and roll out secure networks that best fit their business model and enable maximum agility.

Overview of VPNs and IPSec Technologies

A VPN is a service offering secure, reliable connectivity over a shared public network infrastructure, such as the Internet. Cisco products support the latest in VPN technology.

Cisco defines a VPN as an encrypted connection between private networks over a public network, such as the Internet. The V and N stand for virtual network. The information from a private network is transported over a public network, an internet, to form a virtual network. The P stands for private. To remain private, the traffic is encrypted to keep the data confidential. Therefore, a VPN is a private virtual network.

The three types of implementation scenarios for VPNs are

- Remote-access
- Site-to-site
- Firewall-based

Remote-Access VPNs

The first VPN solution is remote access. Remote access is targeted at mobile users and home telecommuters. In the past, corporations supported remote users via dial-in networks. This typically necessitated a toll or toll-free call to access the corporation. With the advent of VPNs, a mobile user can make a local call to his ISP to access the corporation via the Internet wherever he is. This is an evolution of dial networks. Remote-access VPNs can support the needs of telecommuters, mobile users, extranet consumer-to-business, and so on. Figure 1-4 shows a remote-access VPN. The black dotted line shows VPN traffic across the Internet.

Figure 1-4 *Remote-Access VPN*

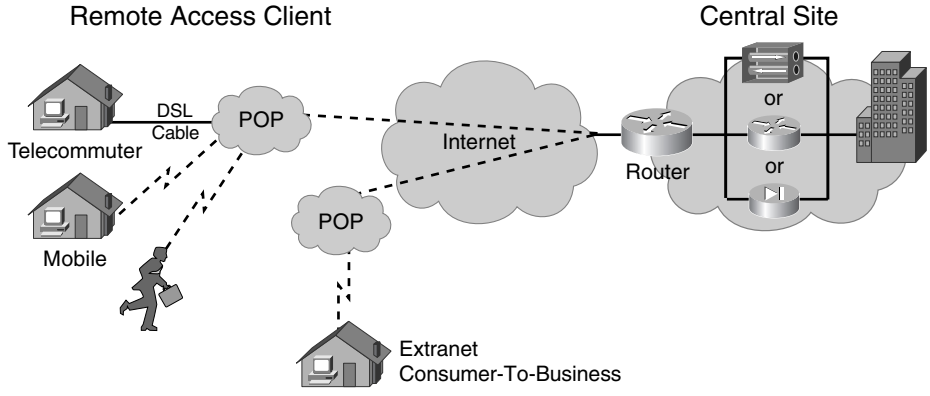

Remote Access VPN—Extension/Evolution of Dial

Site-to-Site VPNs

The next VPN solution is site-to-site (see Figure 1-5). A site-to-site VPN can be used to connect corporate sites. In the past, a leased line or Frame Relay connection was required to connect sites, but now most corporations have Internet access. With Internet access, you can replace leased lines and Frame Relay lines by implementing a site-to-site VPN. Site-to-site VPNs can be used to provide the network connection between the corporation's sites. Site-to-site VPNs can support company intranets and business partner extranets. A site-to-site VPN is an extension of the classic WAN network.

Figure 1-5 *Site-to-Site VPN*

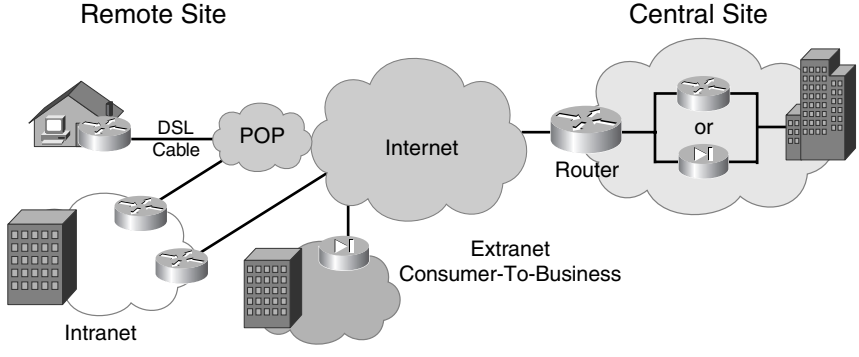

Site-to-Site VPN—Extension of Classic WAN

Firewall-Based VPNs

The last solution is firewall-based VPNs. A firewall-based VPN is inherently a site-to-site solution. Firewall-based VPN solutions are not a technical issue but a security issue. The question is who manages the VPN. If corporate security manages the VPN, a firewall-based VPN might be the VPN solution of choice. Corporations can enhance their existing firewall systems to support VPN services.

The Need for VPNs

The introduction of a VPN into your corporate network infrastructure can provide many benefits:

- **Reduced costs**—Businesses vastly reduce their costs by using the Internet to provide the site-to-site and remote-access infrastructure. Before VPNs, businesses connected using expensive leased lines and telephone systems.
- **Improved communications**—With a VPN, remote-access and home-based users can connect to the central office network from anywhere at any time.
- **Flexibility and scalability**—The introduction of a VPN simplifies and centralizes network administration. The VPN infrastructure can be easily adapted to the company's changing needs, both now and in the future.
- **Security and reliability**—Security is inherent within a VPN, provided through tunneling protocols and encryption software. The reduced number of entry points and the inherent resilience of the Internet mean that the solution is considerably more reliable.
- **Wireless networking**—VPN technology is increasingly combined with wireless connectivity to ensure complete privacy of the data transmitted in environments where data privacy is mandated, such as financial institutions. This ensures that an organization is not vulnerable to inherently weak standard wireless security features.

It is important to note that VPNs can also bring you increased business benefits. They let you develop trust relationships with your suppliers and partners and give your employees round-the-clock access to vital information. Any intranets and extranets that are developed can promote knowledge sharing among partners and employees, and the ease with which information can be accessed and communicated can boost employee morale. These types of benefits cannot be easily measured but can add real value to how you do business and ultimately have a positive impact on turnover and profit.

Implementing a VPN can bring you the benefits outlined here and are a cost-effective, flexible, secure method of managing your digital communications. However, it is important that you work with a partner who ensures that the technology is implemented effectively and forms part of a competent security infrastructure. Many organizations undermine their technical investments, because a lack of detailed knowledge during implementation can leave gaps in security infrastructures that can be exploited and give open access to business-critical information.

IPSec

IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices (peers), such as PIX Firewalls, Cisco routers, Cisco VPN 3000 Concentrators, Cisco VPN Clients, and other IPSec-compliant products. IPSec is not bound to any specific encryption or authentication algorithms, keying technology, or security algorithms. IPSec is a framework of open standards. Because it isn't bound to specific algorithms, IPSec allows newer and better algorithms to be implemented without patching the existing IPSec standards. IPSec provides data confidentiality, data integrity, and data origin authentication between participating peers at the IP layer. IPSec is used to secure a path between a pair of gateways, a pair of hosts, or a gateway and a host.

IPSec spells out the rules for secure communications. IPSec, in turn, relies on existing algorithms to implement the encryption, authentication, and key exchange.

Some of the standard algorithms are as follows:

- **Data Encryption Standard (DES) algorithm**—Used to encrypt and decrypt packet data.
- **3DES algorithm**—Effectively doubles encryption strength over 56-bit DES.
- **Advanced Encryption Standard (AES)**—A newer cipher algorithm designed to replace DES. Has a variable key length between 128 and 256 bits. Cisco is the first industry vendor to implement AES on all its VPN-capable platforms.
- **Message Digest 5 (MD5) algorithm**—Used to authenticate packet data.
- **Secure Hash Algorithm 1 (SHA-1)**—Used to authenticate packet data.
- **Diffie-Hellman (DH)**—A public-key cryptography protocol that allows two parties to establish a shared secret key used by encryption and hash algorithms (for example, DES and MD5) over an insecure communications channel.

Figure 1-6 shows four IPSec framework squares to be filled. When configuring security services to be provided by an IPSec gateway, you first must choose an IPSec protocol. The choices are ESP or ESP with AH. The second square is an encryption algorithm. Choose the encryption algorithm appropriate for the level of security desired: DES or 3DES. The third square is authentication. Choose an authentication algorithm to provide data integrity: MD5 or SHA. The last square is the DH algorithm group. Choose which group to use: DH1, DH2, or DH5. IPSec provides the framework, and the administrator chooses the algorithms used to implement the security services within that framework.

Figure 1-6 *IPSec Protocol Framework*

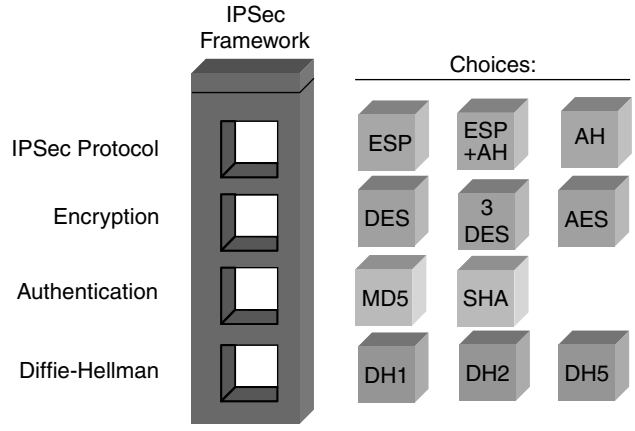

IPSec security services provide four critical functions:

- **Confidentiality (encryption)**—The sender can encrypt the packets before transmitting them across a network. By doing so, no one can eavesdrop on the communication. If intercepted, the communications cannot be read.
- **Data integrity**—The receiver can verify that the data was transmitted through the Internet without being changed or altered in any way.
- **Origin authentication**—The receiver can authenticate the packet's source, guaranteeing and certifying the source of the information.
- **Anti-replay protection**—Anti-replay protection verifies that each packet is unique, not duplicated. IPSec packets are protected by comparing the sequence number of the received packets and a sliding window on the destination host, or security gateway. A packet whose sequence number is before the sliding window is considered late, or a duplicate. Late and duplicate packets are dropped.

Confidentiality

The good news is that the Internet is a public network. The bad news is that the Internet is a public network. Clear-text data transported over the public Internet can be intercepted and read. To keep the data private, it can be encrypted. Through digital scrambling, the data is rendered unreadable.

For encryption to work, both the sender and receiver need to know the rules used to transform the original message into its coded form. Rules are based on an algorithm and a key. An algorithm is a mathematical function that combines a message, text, characters, or all three with a string of characters called a key. The output is an unreadable cipher string. Decryption is extremely difficult or impossible without the correct key.

In Figure 1-7, someone wants to send a financial document across the Internet. At the local end, the document is combined with a key and is run through an encryption algorithm. The output is cipher text. The cipher text is then sent through the Internet. At the remote end, the message is recombined with a key and is sent back through the encryption algorithm. The output is the original financial document.

Figure 1-7 *Encryption*

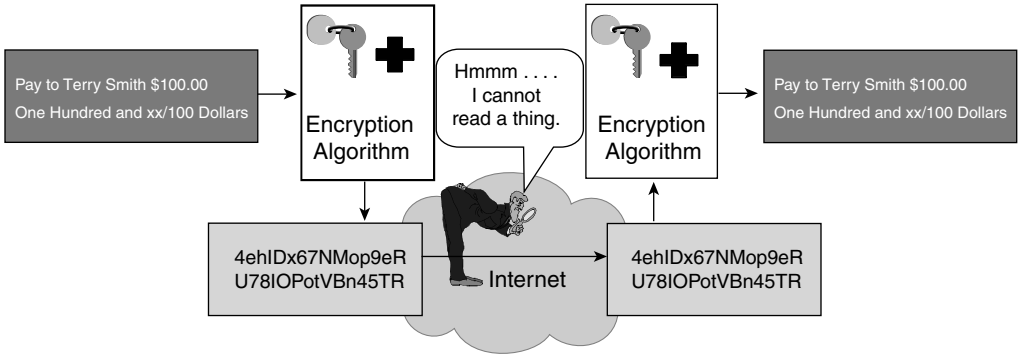

Two types of encryption keys exist: symmetric and asymmetric. With symmetric key encryption, each peer uses the same key to encrypt and decrypt the data. With asymmetric key encryption, the local end uses one key to encrypt the traffic, and the remote end uses another key to decrypt it. Both are discussed in further detail later in this chapter.

Encryption Algorithms

The degree of security depends on the key's length. If someone tries to hack the key through a brute-force attack, guessing every possible combination, the number of possibilities is a function of the key's length. The time to process all the possibilities is a function of the computer's computing power. Therefore, the shorter the key, the easier it is to break. A 64-bit key with a relatively sophisticated computer can take approximately 1 year to break. A 128-bit key with the same machine can take roughly 10^{19} years to decrypt.

Some of the encryption algorithms are as follows:

- **DES**—DES was developed by IBM. It uses a 56-bit key, ensuring high-performance encryption. DES is a symmetric key algorithm.
- **3DES**—The 3DES algorithm is a variant of the 56-bit DES. 3DES operates similarly to DES, in that data is broken into 64-bit blocks. 3DES then processes each block three times, each time with an independent 56-bit key. 3DES effectively doubles encryption strength over 56-bit DES. 3DES is a symmetric key algorithm.

- **AES**—The National Institute of Standards and Technology (NIST) recently adopted AES to replace the existing DES encryption in cryptographic devices. AES provides stronger security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys. Cisco now supports VPN encryption from version 4.0 of the VPN Concentrator software and the addition of a SEP-E module. The older SEP modules do not perform hardware encryption on AES, only on DES and 3DES.
- **RSA**—Rivest, Shamir, and Adelman (RSA) encryption, shown in Figure 1-8, uses asymmetric keys for encryption and decryption. Each end, local and remote, generates two encryption keys: a private key and a public key. It keeps its private key and exchanges its public key with people with whom it wants to communicate. To send an encrypted message to the remote end, the local end encrypts the message using the remote's public key and the RSA encryption algorithm. The result is an unreadable cipher text. This message is sent through the insecure network. The remote end uses its private key and the RSA algorithm to decrypt the cipher text. The result is the original message. The only one who can decrypt the message is the destination that owns the private key. With RSA encryption, the opposite also holds true. The remote end can encrypt a message using its own private key. The receiver can decrypt the message using the sender's public key. This RSA encryption technique is used for digital signatures.

Figure 1-8 *RSA Encryption*

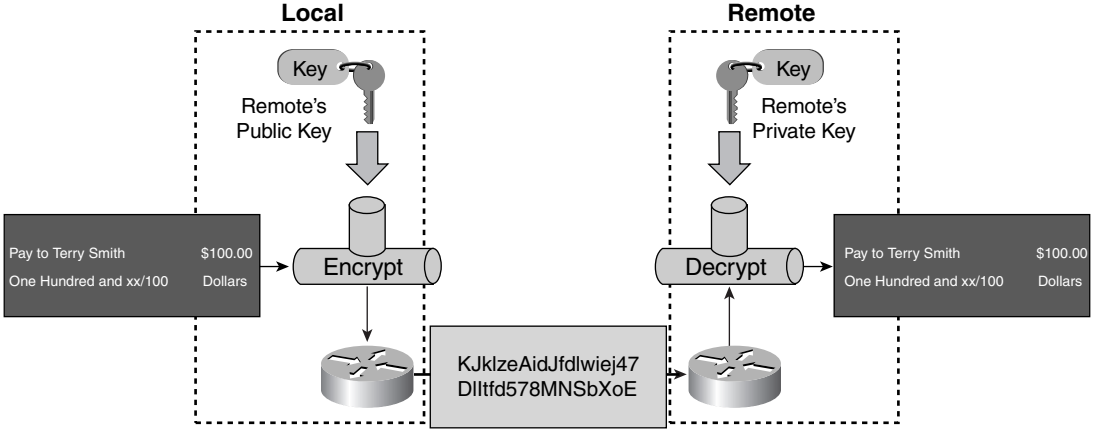

Key Exchange

DES, 3DES, AES, and also the two authentication algorithms, MD5 and SHA-1, all require a symmetric shared secret key to perform encryption and decryption. The question is, how do the encrypting and decrypting devices get the shared secret key?

The keys can be sent by e-mail, courier, overnight express, or public key exchange. The easiest method is DH public key exchange. The DH key agreement is a public key exchange method that provides a way for two peers to establish a shared secret key that only they know, although they are communicating over an insecure channel.

Public key cryptosystems rely on a two-key system: a public key, which is exchanged between end users, and a private key, which is kept secret by the original owners. The DH public key algorithm states that if user A and user B exchange public keys, and a calculation is performed on their individual private key and one another's public key, the end result of the process is an identical shared key. The shared key is used to derive encryption and authentication keys.

Variations of the DH key exchange algorithm are known as DH groups 1 through 7. DH groups 1, 2, and 5 support exponentiation over a prime modulus with a key size of 768 bits, 1024 bits, and 1536 bits, respectively. Cisco 3000 Clients support DH groups 1, 2, and 5. DES and 3DES encryption support DH groups 1 and 2. AES encryption supports DH groups 2 and 5. In addition to these, the Certicom movianVPN Client supports group 7. Group 7 supports Elliptical Curve Cryptography (ECC), which reduces the time needed to generate keys. During tunnel setup, VPN peers negotiate which DH group to use.

Security is not an issue with DH key exchange. Although someone might know a user's public key, the shared secret cannot be generated, because the private key never becomes public.

DH Key Exchange

DH key exchange is a public key exchange method that provides a way for two IPSec peers to establish a shared secret key that only they know, although they are communicating over an insecure channel.

With DH, each peer generates a public/private key pair. The private key generated by each peer is kept secret and never shared. The public key is calculated from the private key by each peer and is exchanged over the insecure channel. Each peer combines the other's public key with its own private key and computes the same shared secret number. The shared secret number is then converted into a shared secret key. The shared secret key is never exchanged over the insecure channel.

The following steps are used to implement the DH process (see Figure 1-9):

1. Each peer generates a large prime integer, p and q. Each peer sends the other its prime integer over the insecure channel. For example, peer A sends p to peer B. Each peer then uses the p and q values to generate g, a primitive root of p.
2. Each peer generates a private DH key (peer A: X_a, peer B: X_b).

3. Each peer generates a public DH key. The local private key is combined with the prime number p and the primitive root g in each peer to generate a public key, Y_a for peer A and Y_b for peer B. The formula for peer A is $Y_a = g\text{^}X_a \bmod p$. The formula for peer B is $Y_b = g\text{^}X_b \bmod p$. The exponentiation is computationally expensive. The ^ character denotes exponentiation (g to the X_a power); mod denotes modulus.
4. The public keys Y_a and Y_b are exchanged in public.
5. Each peer generates a shared secret number (ZZ) by combining the public key received from the opposite peer with its own private key. The formula for peer A is $ZZ = (Y_bX_a) \bmod p$. The formula for peer B is $ZZ = (Y_aX_b) \bmod p$. The ZZ values are identical in each peer. Anyone who knows p or g, or the DH public keys, cannot guess or easily calculate the shared secret value—largely because of the difficulty in factoring large prime numbers.
6. Shared secret number ZZ is used to derive the encryption and authentication symmetric keys.

Figure 1-9 *DH Key Exchange*

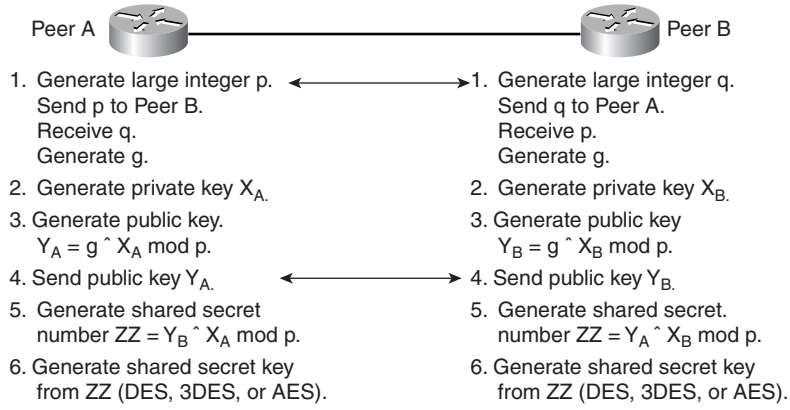

Data Integrity

The next VPN-critical function is data integrity. VPN data is transported over some form of insecure network, such as the Internet. Potentially, this data could be intercepted and modified. To guard against this, each message has a hash attached to it. This is called a Hash-based Message Authentication Code (HMAC). A hash guarantees the integrity of the original message. If the transmitted hash matches the received hash, the message has not been tampered with. However, if there is no match, the message was altered.

In Figure 1-10, someone is trying to send Terry Smith a check for $100. At the remote end, Alex Jones is trying to cash the check for $1000. As the check progressed through the

Internet, it was altered. Both the recipient and the dollar amount were changed. In this case, the hashes do not match, so the transaction is no longer valid.

Figure 1-10 *Data Integrity*

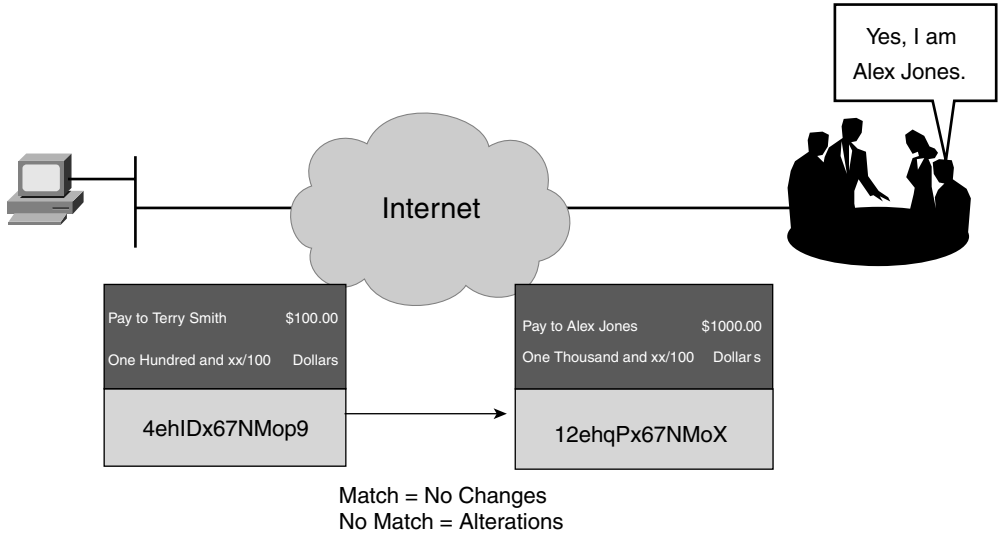

Data integrity is commonly known and talked about as authentication. The packets are authenticated using the hash that is attached to each packet. Two main algorithms facilitate data integrity within the IPSec framework—MD5 and SHA-1.

HMAC guarantees the message's integrity. At the local end, the message and a shared secret key are sent through a hash algorithm, which produces a hash value. Basically, a hash algorithm is a formula used to convert a variable-length message into a single string of digits of a fixed length. It is a one-way algorithm. A message can produce a hash, but a hash cannot produce the original message. It is analogous to dropping a plate on the floor. The plate can produce a multitude of pieces, but the pieces cannot be recombined to reproduce the plate in its original form. The message and hash are sent over the network.

At the remote end, a two-step process occurs, as shown in Figure 1-11. First, the received message and shared secret key are sent through the hash algorithm, resulting in a recalculated hash value. Second, the receiver compares the recalculated hash with the hash that was attached to the message. If the original hash and recalculated hash match, the message's integrity is guaranteed. If any of the original message is changed while in transit, the hash values are different.

Figure 1-11 *HMAC*

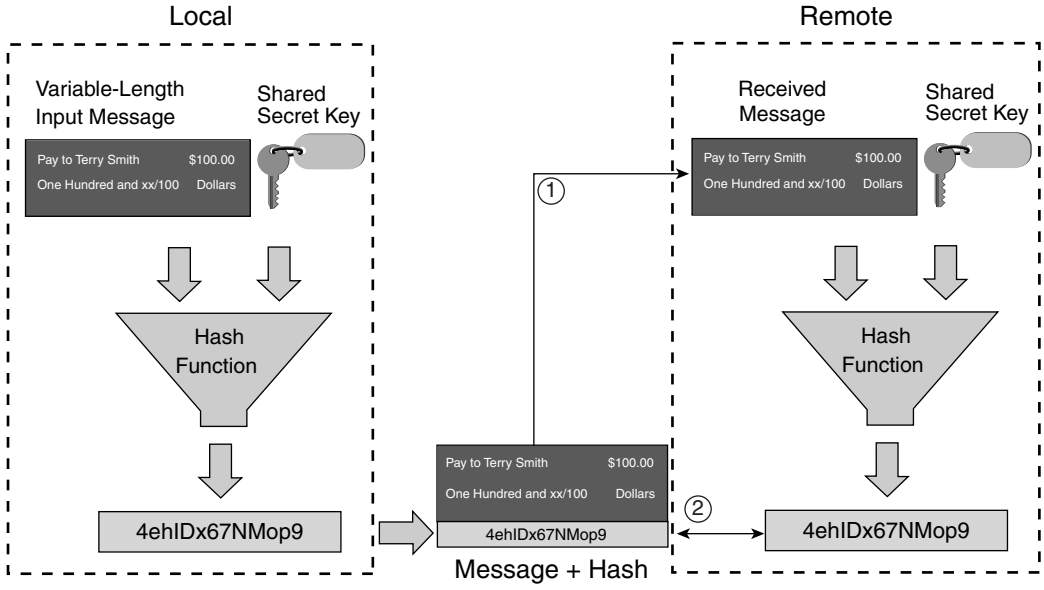

The two common HMAC algorithms are as follows:

- **HMAC-MD5**—HMAC-MD5 uses a 128-bit shared secret key. The variable-length message and the 128-bit shared secret key are combined and run through the HMAC-MD5 hash algorithm. The output is a 128-bit hash. The hash is appended to the original message and forwarded to the remote end.
- **HMAC-SHA-1**—HMAC-SHA-1 uses a 160-bit secret key. The variable-length message and the 160-bit shared secret key are combined and run through the HMAC-SHA-1 hash algorithm. The output is a 160-bit hash. The hash is appended to the original message and forwarded to the remote end.

  HMAC-SHA-1 is considered cryptographically stronger than HMAC-MD5. HMAC-SHA-1 is recommended when its slightly superior security is important.

Origin Authentication

In the middle ages, a seal guaranteed the authenticity of an edict. In modern times, a signed document is notarized with a seal and a signature. In the electronic era, a document is signed using the sender's private encryption key—a digital signature. A signature is authenticated by decrypting the signature with the sender's public key.

In Figure 1-12, the local device derives a hash and encrypts it with its private key. The encrypted hash (digital signature) is attached to the message and is forwarded to the remote

end. At the remote end, the encrypted hash is decrypted using the local end's public key. If the decrypted hash matches the recomputed hash, the signature is genuine. A digital signature ties a message to a sender. The sender is authenticated. It is used during the initial establishment of a VPN tunnel to authenticate both ends to the tunnel.

Figure 1-12 *Digital Signature*

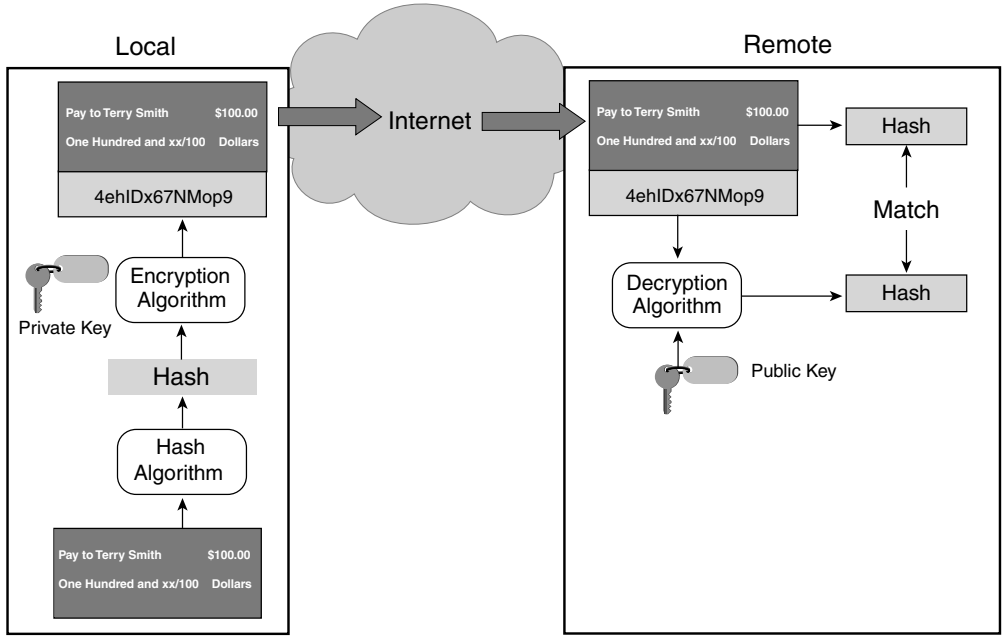

The two common digital signature algorithms are RSA and Directory System Agent (DSA). RSA is used commercially and is the most common. DSA is used by U.S. Government agencies and is not as common.

When conducting business long distance, it is necessary to know who is on the other end of the phone, e-mail, or fax. The same is true of VPNs. The device on the other end of the VPN tunnel must be authenticated before the communication path is considered secure.

The three peer authentication methods are as follows:

- **Preshared keys**—A secret key value entered into each peer manually that is used to authenticate the peer.
- **RSA signatures**—Uses the exchange of digital certificates to authenticate the peers.
- **RSA encrypted nonces**—Nonces (random numbers generated by each peer) are encrypted and then exchanged between peers. The two nonces are used during the peer authentication process.

Preshared Keys

With preshared keys, the same preshared key is configured on each IPSec peer. At each end, the preshared key is combined with other information to form the authentication key. Starting at the local end, the authentication key and the identity information (device-specific information) are sent through a hash algorithm to form hash_I. The local Internet Key Exchange (IKE) peer provides one-way authentication by sending hash_I to the remote peer. If the remote peer can independently create the same hash, the local peer is authenticated, as shown in Figure 1-13.

Figure 1-13 *Preshared Keys*

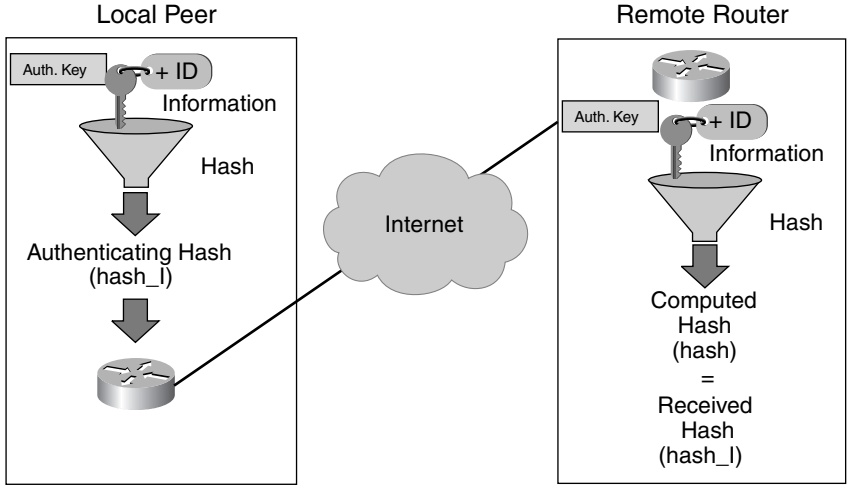

The authentication process continues in the opposite direction. The remote peer combines its identity information with the preshared-based authentication key and sends them through a hash algorithm to form hash_R. hash_R is sent to the local peer. If the local peer can independently create the same hash from its stored information and preshared-based authentication key, the remote peer is authenticated. Each peer must authenticate its opposite peer before the tunnel is considered secure. Preshared keys are easy to configure manually but do not scale well. Each IPSec peer must be configured with the preshared key of every other peer with which it communicates.

RSA Signatures

With RSA signatures (see Figure 1-14), hash_I and hash_R not only are authenticated but also are digitally signed. Starting at the local end, the authentication key and identity information (device-specific information) are sent through a hash algorithm to form hash_I. hash_I is then encrypted using the local peer's private encryption key. The result is a digital signature. The digital signature and a digital certificate are forwarded to the remote peer.

(The public encryption key for decrypting the signature is included in the digital certificate exchanged between peers.)

Figure 1-14 *RSA Signatures*

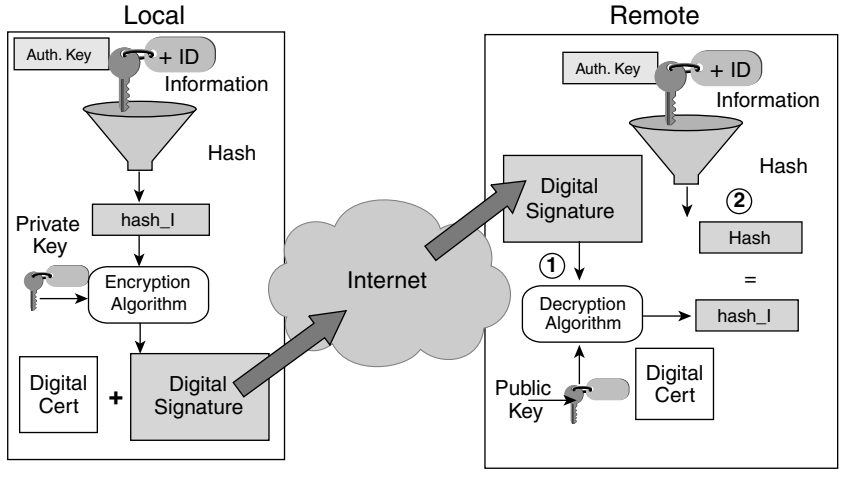

At the remote peer, local peer authentication is a two-step process. First, the remote peer verifies the digital signature by decrypting it using the public encryption key enclosed in the digital certificate. The result is hash_I. Next, the remote peer independently creates hash_I from stored information. If the calculated hash_I equals the decrypted hash_I, the local peer is authenticated (as shown in the figure). Digital signatures and certificates are discussed in more detail in Chapter 5, "Configuring the Cisco VPN 3000 for Remote Access Using Digital Certificates."

After the remote peer authenticates the local peer, the authentication process begins in the opposite direction. The remote peer combines its identity information with the authentication key and sends them through a hash algorithm to form hash_R. hash_R is encrypted using the remote peer's private encryption key—a digital signature. The digital signature and certificate are sent to the local peer. The local peer performs two tasks: It creates hash_R from stored information, and it decrypts the digital signature. If the calculated hash_R and the decrypted hash_R match, the remote peer is authenticated. Each peer must authenticate its opposite peer before the tunnel is considered secure.

RSA Encrypted Nonces

RSA encrypted nonces require that each party generate a nonce—a pseudorandom number. The nonces are then encrypted and exchanged. Upon receipt of the nonce, each end formulates an authentication key made up of the initiator and responder nonces, the DH key, and the initiator and responder cookies. The nonce-based authentication key is combined with

device-specific information and run through a hash algorithm, where the output becomes hash_I. The local IKE peer provides one-way authentication by sending hash_I to the remote peer. If the remote peer can independently create the same hash from stored information and its nonce-based authentication key, the local peer is authenticated.

After the remote end authenticates the local peer, the authentication process begins in the opposite direction. The remote peer combines its identity information with the nonce-based authentication key and sends them through a hash algorithm to form hash_R. Hash_R is sent to the local peer. If the local peer can independently create the same hash from stored information and the nonce-based key, the remote peer is authenticated. Each peer must authenticate its opposite peer before the tunnel is considered secure.

Anti-Replay Protection

IPSec uses anti-replay mechanisms to ensure that IP packets cannot be intercepted by a third party or man in the middle and then be changed and reinserted into the data stream. This is implemented in IPSec by the Authentication Header (AH) protocol and the Encapsulating Security Payload (ESP) protocol. The anti-replay mechanism works by keeping track of the sequence number allocated to each packet as it arrives at the VPN endpoint. When a security association is established between two VPN endpoints, the sequence counter is set to 0. The packets that are encrypted and transmitted over the VPN are sequenced starting from 1. Each time a packet is sent, the receiver of the packet verifies that the sequence number is not that of a previously sent packet. If the receiver receives a packet with a duplicate sequence number, the packet is discarded, and an error message is sent back to the transmitting VPN endpoint to log this event.

NOTE AH implements anti-replay by default, although ESP does it only when data authentication is turned on (for example, MD5 or SHA-1) in the IPSec transform-set.

IPSec Protocol Framework

The preceding section discussed encryption, authentication, and integrity. This section explains how encryption, integrity, and authentication are applied to the IPSec protocol suite.

As mentioned, IPSec is a framework of open standards. IPSec spells out the messaging to secure the communications but relies on existing algorithms, such as DES and 3DES, to implement the encryption and authentication. The two main IPSec framework protocols are as follows:

- **AH**—AH, shown in Figure 1-15, is the appropriate protocol when confidentiality is not required or permitted. It provides data authentication and integrity for IP packets passed between two systems. It is a means of verifying that any message passed from

Router A to Router B was not modified during transit. It verifies that the data's origin was either Router A or Router B. AH does not provide data confidentiality (encryption) of packets. It does the following:

- Ensures data integrity
- Provides origin authentication (ensures that packets definitely came from the peer router)
- Uses a keyed-hash mechanism
- Does not provide confidentiality (no encryption)
- Provides anti-replay protection

All text is transported in the clear.

Figure 1-15 *AH*

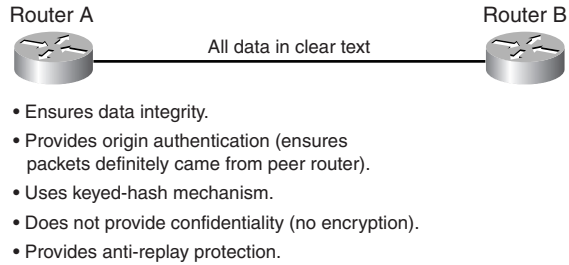

- **ESP**—A security protocol may be used to provide confidentiality (encryption) and authentication. ESP, shown in Figure 1-16, provides confidentiality by performing encryption at the IP packet layer. IP packet encryption conceals the data payload and the identities of the ultimate source and destination. ESP provides authentication for the inner IP packet and ESP header. Authentication provides data origin authentication and data integrity. Although both encryption and authentication are optional in ESP, at a minimum, one of them must be selected. ESP provides
  - Data confidentiality (encryption)
  - Data integrity
  - Data origin authentication
  - Anti-replay protection

Figure 1-16 *ESP*

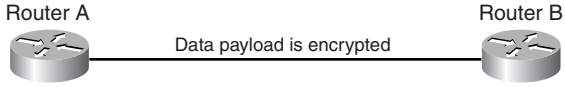

- Data confidentiality (encryption).
- Data integrity.
- Data origin authentication.
- Anti-replay protection.

AH

Authentication is achieved by applying a keyed one-way hash function to the packet to create a hash or message digest. The hash is combined with the text and is transmitted. Changes in any part of the packet that occur during transit are detected by the receiver when it performs the same one-way hash function on the received packet and compares the value of the message digest the sender has supplied. The fact that the one-way hash also involves the use of a symmetric key between the two systems means that authenticity is guaranteed.

The AH function is applied to the entire datagram, except for any mutable IP header fields that change in transit (such as Time To Live [TTL] fields that are modified by the routers along the transmission path). AH works as follows:

1. The IP header and data payload are hashed.
2. The hash is used to build an AH header, which is inserted into the original packet.
3. The new packet is transmitted to the IPSec peer router.
4. The peer router hashes the IP header and data payload.
5. The peer router extracts the transmitted hash from the AH header.
6. The peer router compares the two hashes. The hashes must exactly match. Even if one bit is changed in the transmitted packet, the hash output on the received packet changes, and the AH header does not match.

AH supports the HMAC-MD5 and HMAC-SHA-1 algorithms. AH authentication and integrity are shown in Figure 1-17.

Figure 1-17 *AH Authentication and Integrity*

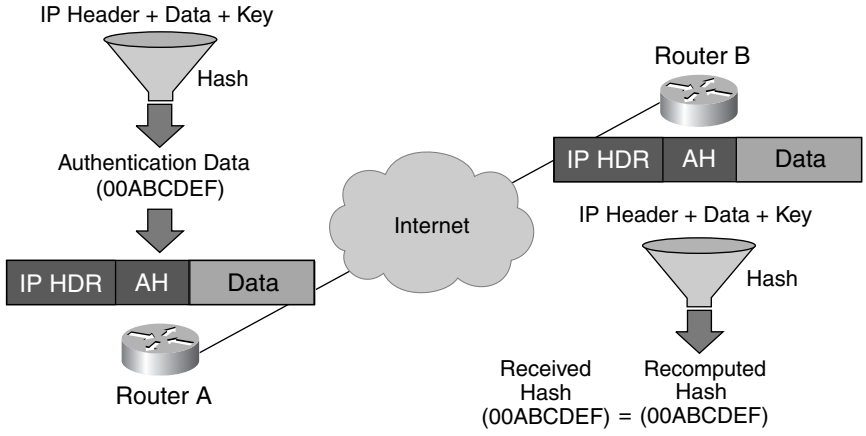

NOTE AH is IP protocol 51. If you are using IPSec, ensure that any perimeter routers will pass IP protocol 51 as well as Internet Security Association and Key Management Protocol (ISAKMP)—that is, UDP port 500.

ESP

ESP, shown in Figure 1-18, provides confidentiality by encrypting the payload. It supports a variety of symmetric encryption algorithms. The default algorithm for IPSec is 56-bit DES. Cisco products also support the use of 3DES and AES for stronger encryption.

Figure 1-18 *ESP Protocol*

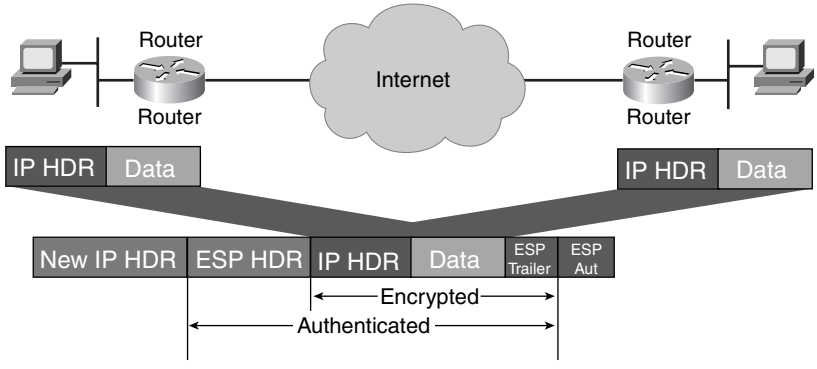

ESP can be used alone or in combination with AH. ESP with AH also provides integrity and authentication of datagrams. First, the payload is encrypted. Next, the encrypted payload is sent through a hash algorithm—HMAC-MD5 or HMAC-SHA-1. The hash provides origin authentication and data integrity for the data payload.

Alternatively, ESP may also enforce anti-replay protection by requiring that a receiving host set the replay bit in the header to indicate that the packet has been seen.

Between two security gateways, the original payload is well protected, because the entire original IP datagram is encrypted. An ESP header and trailer are added to the encrypted payload. With ESP authentication, the encrypted IP datagram and the ESP header or trailer are included in the hashing process. Last, a new IP header is appended to the front of the authenticated payload. The new IP address is used to route the packet through the Internet.

When both ESP authentication and encryption are selected, encryption is performed before authentication. One reason for this order of processing is that it facilitates rapid detection and rejection of replayed or bogus packets by the receiving node. Before decrypting the packet, the receiver can authenticate inbound packets. By doing this, it can detect the problems and potentially reduce the impact of DoS attacks.

NOTE ESP is IP protocol 50. If you are using IPSec, ensure that any perimeter routers pass IP protocol 50 as well as ISAKMP—that is, UDP port 500.

Modes of Operation

ESP and AH can be applied to IP packets in two different ways, or modes:

- Transport mode
- Tunnel mode

Transport Mode

Transport mode, shown in Figure 1-19, protects the packet's payload, higher-layer protocols, but leaves the original IP address in the clear. The original IP address is used to route the packet through the Internet. ESP transport mode is used between two hosts. Transport mode provides security to the higher-layer protocols only.

Figure 1-19 *Transport Mode*

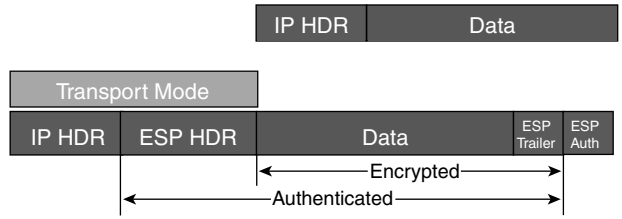

Tunnel Mode

ESP tunnel mode, shown in Figure 1-20, is used when either end of the tunnel is a security gateway, a Concentrator, a VPN optimized router, or a PIX Firewall. Tunnel mode is used when the final destination is not a host, but a VPN gateway. The security gateway encrypts and authenticates the original IP packet. Next, a new IP header is appended to the front of the encrypted packet. The new outside IP address is used to route the packet through the Internet to the remote end security gateway. Tunnel mode provides security for the whole original IP packet.

Figure 1-20 *Tunnel Mode*

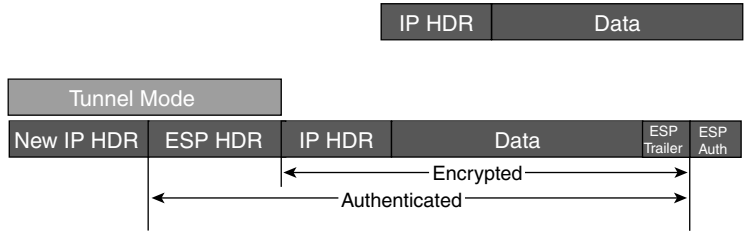

How IPSec Works

The goal of IPSec is to protect the desired data with the needed security services. IPSec's operation can be broken into five primary steps:

Step 1 **Define interesting traffic**—Traffic is deemed interesting when the VPN device recognizes that the traffic you want to send needs to be protected.

Step 2 **IKE Phase 1**—Between peers, a basic set of security services is negotiated and agreed on. This basic set of security services protects all subsequent communications between the peers. IKE Phase 1 sets up a secure communications channel between peers.

Step 3 **IKE Phase 2**—IKE negotiates IPSec security association (SA) parameters and sets up matching IPSec SAs in the peers. These security parameters are used to protect data and messages exchanged between endpoints.

Step 4 **Data transfer**—Data is transferred between IPSec peers based on the IPSec parameters and keys stored in the SA database.

Step 5 **IPSec tunnel termination**—IPSec SAs terminate through deletion or by timing out.

Step 1: Define Interesting Traffic

Determining what traffic needs to be protected is done as part of formulating a security policy for use of a VPN. The policy is used to determine what traffic needs to be protected and what traffic can be sent in the clear. For every inbound and outbound packet, you have three choices:

- Apply IPSec
- Bypass IPSec
- Discard the packet

For every packet protected by IPSec, the system administrator must specify the security services applied to the packet. The security policy database specifies the IPSec protocols, modes, and algorithms applied to the traffic. The services are then applied to traffic destined for each particular IPSec peer. With the VPN Client, you use menu windows to select connections you want secured by IPSec. When interesting traffic transits the IPSec client, the client initiates the next step in the process: negotiating an IKE Phase 1 exchange. Figure 1-21 shows two routers with Host A and Host B at either end. You have to decide whether to encrypt, not encrypt, or drop the packets.

Figure 1-21 *Step 1: Define Interesting Traffic*

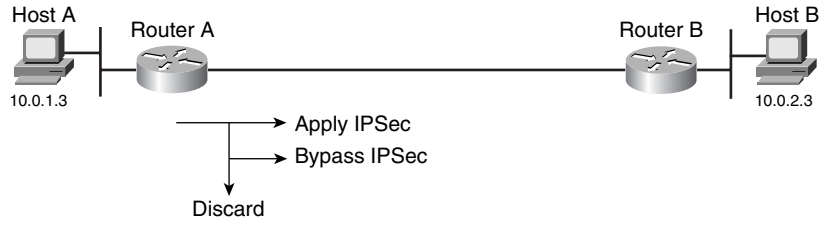

Step 2: IKE Phase 1

The basic purpose of IKE Phase 1, shown in Figure 1-22, is to negotiate IKE policy sets, authenticate the peers, and set up a secure channel between the peers. IKE Phase 1 occurs in two modes: main mode and aggressive mode.

Figure 1-22 *Step 2: IKE Phase 1*

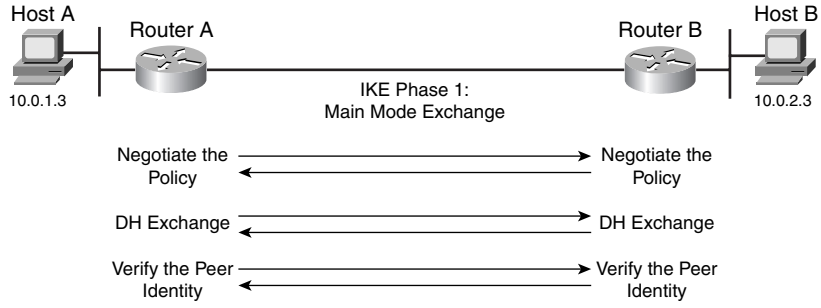

Main mode has three two-way exchanges between the initiator and receiver:

- **First exchange**—The algorithms and hashes used to secure the IKE communications are negotiated and agreed on between peers.
- **Second exchange**—Uses a DH exchange to generate shared secret keys and to pass nonces, which are random numbers sent to the other party, signed, and returned to prove their identity. The shared secret key is used to generate all the other encryption and authentication keys.
- **Third exchange**—Verifies the other side's identity. It is used to authenticate the remote peer. The main outcome of main mode is a secure communication path for subsequent exchanges between the peers. Without proper authentication, it is possible to establish a secure communication channel with a hacker who is now stealing all your sensitive material.

In aggressive mode, fewer exchanges are done and with fewer packets. On the first exchange, almost everything is squeezed in: the IKE policy set negotiation; the DH public key generation; a nonce, which the other party signs; and an identity packet, which can be used to verify the identity via a third party. The receiver sends everything back that is needed to complete the exchange. The only thing left is for the initiator to confirm the exchange.

When trying to make a secure connection between Host A and B through the Internet, IKE security proposals are exchanged between Routers A and B. The proposals identify the IPSec protocol being negotiated (for example, ESP). Under each proposal, the originator must delineate which algorithms are employed in the proposal (for example, DES with MD5). Rather than negotiate each algorithm individually, the algorithms are grouped into IKE transform sets. A transform set delineates which encryption algorithm, authentication

algorithm, mode, and key length are proposed. These IKE proposals and transform sets are exchanged during the IKE main mode first exchange phase. If a transform set match is found between peers, the main mode continues. If no match is found, the tunnel is torn down.

In Figure 1-23, Router A sends IKE transform sets 10 and 20 to Router B. Router B compares its set, transform set 15, with those received from Router A. In this instance, a match occurs: Router A's transform set 10 matches Router B's transform set 15.

Figure 1-23 *Step 2: IKE Transform Sets*

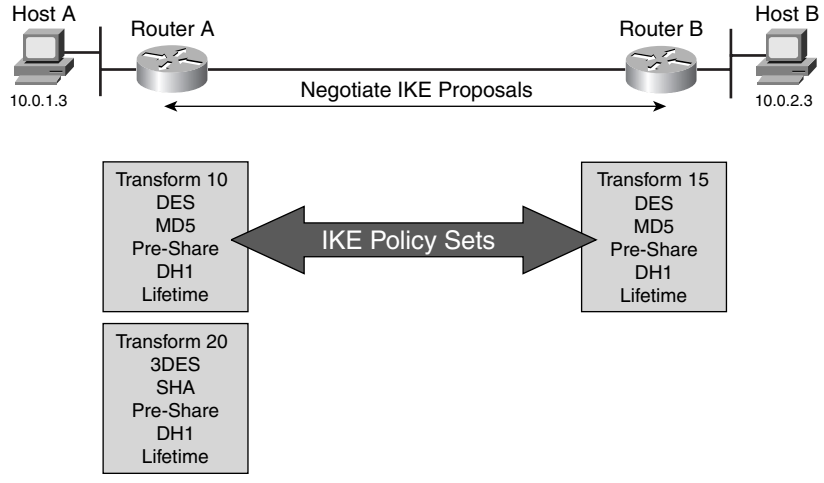

In a point-to-point application, each end might need only a single IKE policy set defined. However, in a hub-and-spoke environment, the central site might require multiple IKE policy sets to satisfy all the remote peers.

Step 3: IKE Phase 2

The purpose of IKE Phase 2 is to negotiate the IPSec security parameters that are applied to the interesting traffic traversing the tunnel negotiated during Phase 1. IKE Phase 2 performs the following functions:

- Negotiates IPSec security parameters and IPSec transform sets
- Establishes IPSec SAs
- Periodically renegotiates IPSec SAs to ensure security
- Optionally performs an additional DH exchange

IKE Phase 2 has one mode—quick mode. Quick mode occurs after IKE has established the secure tunnel in Phase 1. It negotiates a shared IPSec transform, derives shared secret keying material used for the IPSec security algorithms, and establishes IPSec SAs. Quick mode exchanges nonces that are used to generate new shared secret key material and to prevent replay attacks from generating bogus SAs.

Quick mode is used to renegotiate a new IPSec SA when the IPSec SA lifetime expires. It's also used to refresh the keying material used to create the shared secret key based on the keying material derived from the DH exchange in Phase 1. Figure 1-24 shows the negotiation of IPSec parameters between Router A and Router B.

Figure 1-24 *Step 3: IKE Phase 2*

The ultimate goal of IKE Phase 2 is to establish a secure IPSec session between endpoints. Before that can happen, each pair of endpoints negotiates the level of security required (for example, encryption and authentication algorithms for the session). Rather than negotiate each protocol individually, the protocols are grouped into IPSec transform sets. IPSec transform sets are exchanged between peers during quick mode. If a match is found between sets, IPSec session establishment continues. If no match is found, the session is torn down.

In Figure 1-25, Router A sends IPSec transform sets 30 and 40 to Router B. Router B compares its set, transform set 55, with those received from Router A. In this instance, a match occurs. Router A's transform set 30 matches Router B's transform set 55. These encryption and authentication algorithms form an SA.

When the peers agree on the security services, each VPN peer device enters the information in a security policy database (SPD). The information includes the encryption and authentication algorithm, destination IP address, transport mode, key lifetime, and so on. This information is the SA—a one-way logical connection that provides security to all traffic traversing the connection. Because most traffic is bidirectional, two SAs are required: one for inbound traffic, and one for outbound traffic. The VPN device indexes the SA with a number, a Security Parameter Index (SPI). Rather than send the SA's individual parameters across the tunnel, the source gateway, or host, inserts the SPI into the ESP header. When the IPSec peer receives the packet, it looks up the destination IP address, IPSec protocol, and SPI in its SA database (SAD) and then processes the packet according to the algorithms listed under the SPD.

Figure 1-25 *Step 3: IPSec Transform Sets*

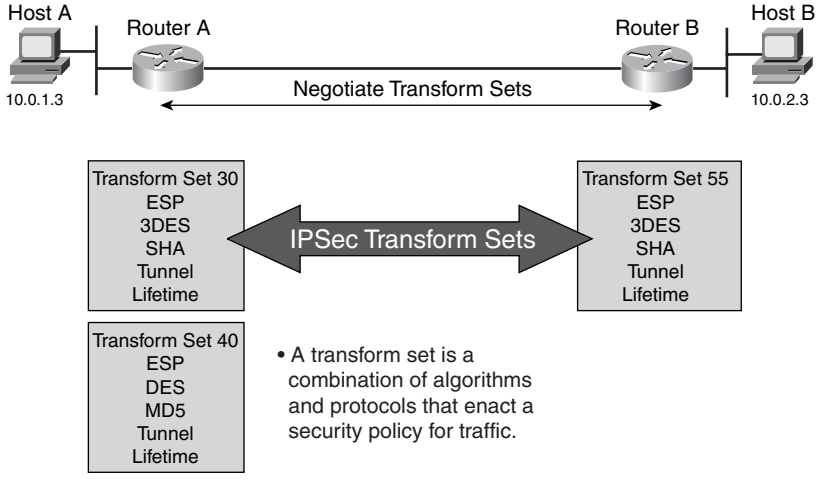

The IPSec SA is a compilation of the SAD and SPD. The SAD identifies the SA destination IP address, IPSec protocol, and SPI number. The SPD defines the security services applied to the SA, encryption and authentication algorithms, and mode and key lifetime. For example, in the corporate-to-bank connection shown in Figure 1-26, the security policy provides a very secure tunnel using 3DES, SHA, tunnel mode, and a key lifetime of 28,800. The SAD value is 192.168.2.1, ESP, and SPI-12. For the remote user accessing e-mails, a less secure policy is negotiated using DES, MD5, tunnel mode, and a key lifetime of 28,800. The SAD values are a destination IP address of 192.169.12.1, ESP, and SPI-39.

With a password on your company PC, the longer you keep it, the more vulnerable it becomes. The same thing is true of keys and SAs. For good security, the SA and keys should be changed periodically. There are two parameters: lifetime type and duration. How is the lifetime measured? Is it measured by the number of bytes transmitted or the amount of time transpired? The second parameter is the unit of measure: kilobytes of data or seconds of time. An example is a lifetime based on 10,000 KB of data transmitted or 28,800 seconds of time expired. The keys and SAs remain active until their lifetime expires or until an external event—such as the client dropping the tunnel—causes them to be deleted.

Figure 1-26 *Step 3: SA*

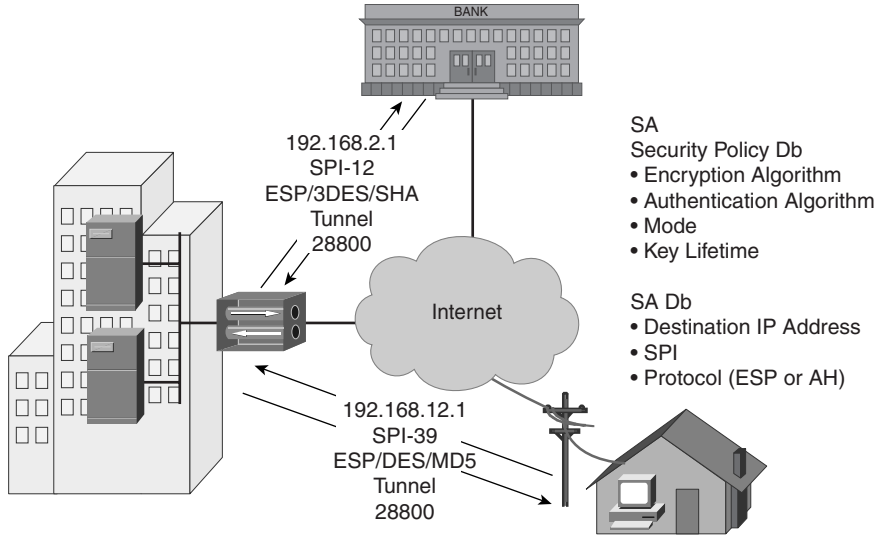

Step 4: Data Transfer

After IKE Phase 2 is complete and quick mode has established IPSec SAs, traffic is exchanged between Hosts A and B via a secure tunnel, as shown in Figure 1-27. Interesting traffic is encrypted and decrypted according to the security services specified in the IPSec SA.

Figure 1-27 *Step 4: IPSec Data Transfer*

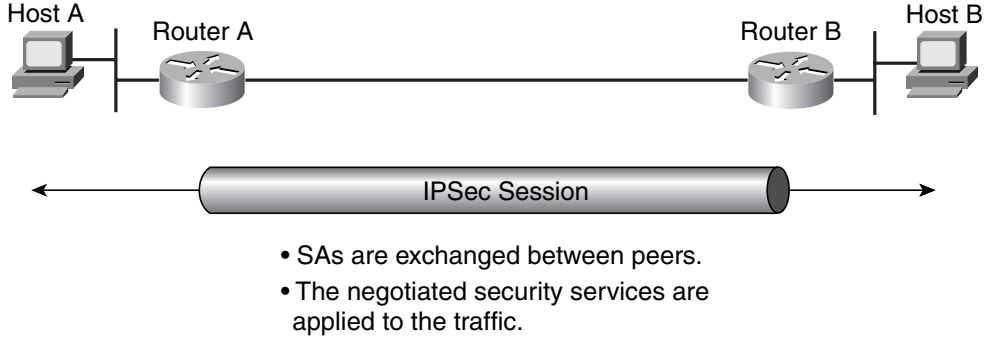

Step 5: IPSec Tunnel Termination

IPSec SAs terminate through deletion or by timing out. An SA can time out when a specified number of seconds has elapsed or when a specified number of bytes has passed through the tunnel. When the SAs terminate, the keys are also discarded. When subsequent IPSec SAs are needed for a flow, IKE performs a new Phase 2 (and, if necessary, a new Phase 1) negotiation. A successful negotiation results in new SAs and new keys. New SAs are usually established before the existing SAs expire so that a given flow can continue uninterrupted. This final step is shown in Figure 1-28.

Figure 1-28 *Step 5: IPSec Tunnel Termination*

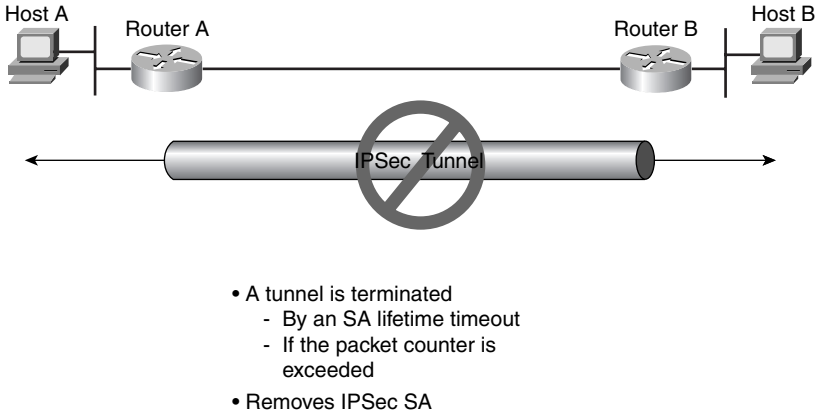

Summary

This chapter provided a technology overview that is the foundation of the rest of this book. It started by looking at network security before providing a brief overview of the AVVID program and the Cisco SAFE security blueprint. It then looked at IPSec and the components that make up this framework. There is quite a lot to IPSec, and quite a few permutations can be used, such as encryption and authentication algorithms. This chapter covered each of these options and finished by looking at the five important steps in how IPSec is implemented in all Cisco devices. The next chapters build on this information and delve more into the configuration aspects of VPNs for the Cisco VPN 3000 Concentrator range of products.

Review Questions

The following questions test your retention of the material presented in this chapter. The answers appear in Appendix A, "Answers to Chapter Review Questions."

1. What two main protocols make up the IPSec framework?
2. What IP protocol does ESP use?
3. What are the two modes of IKE Phase 1?
4. What three key lengths can AES currently use?
5. What type of VPN would you be using if you are a user based at home connecting to the central site over a VPN by using a VPN Software Client installed on your laptop computer?
6. What are the four steps of the Security Wheel?
7. What is the normal method of key exchange for the encryption algorithms used in IPSec, such as DES, 3DES, and AES?
8. What is the main issue with firewall-based VPNs?
9. What are the two modes of IPSec operation?
10. What three authentication methods are used in the IPSec protocol's origin identification feature?