Product information begins on page 2.

Lucent and Ascend have merged.

With the Lucent-Ascend merger, customers gain a broader and more powerful portfolio of next-generation data, voice, fax, and video services and products. To access up-to-the-minute information about our products, see page 2.

We also invite you to contact us with your questions directly at: info@ascend.com

Ascend

White Paper

IP Security

Table of Contents

1. Executive Summary

Today's enterprise networks are far more "open" than ever before. They support dial-in access for mobile workers and telecommuters. They provide communications with customers and suppliers. They interface to the vast capabilities of the worldwide Internet. All of this helps promote employee interaction, increase productivity, gain a competitive edge, and more. But along with the many benefits of open networking comes a responsibility to secure all network-attached resources from unwanted intrusion.

The purpose of this resource guide is to position industry standard IP Security's role in the overall framework of network security. The material is suitable for network specialists—from the department manager to the network manager—who desire a high-level overview of network security in general, and IP Security in particular.

Network security, in general, is embodied in three related areas:

- Protection of resources
- Proof of identity
- Privacy of information

IP Security, or IPSec, is a series of standards that provide general-purpose security for any IP-based network, including intranets, extranets and the Internet itself. IPSec has a vital role to play in two of the three areas: proof of identity and privacy of information.

IPSec's Authentication Header (AH) positively identifies the source of an incoming packet using a special digital signature for each and every packet. IPSec's Encapsulating Security Payload encrypts all packets so that their content cannot be viewed by others, and is therefore kept confidential. Both elements of IPSec also provide integrity, which assures that no one has altered any packet while in transit.

A fourth area of network security is managing the individual and collective elements. For IPSec, this means key management. Keys are long strings of characters used to "lock" and "unlock" the specially-coded headers and packets. Key management involves the creation, as well as the distribution or exchange of these keys.

The white paper's content is organized into two remaining chapters. Chapter 2, *End-to-End Security,* gives the "big picture" of securing networks to establish an overall framework. Chapter 3, *IPSec's Role in Network Security* provides an introduction to IPSec technology and outlines the specific contributions made by IP Security.

2. End-to-End Network Security

When your information is compromised, so is your business. Networks are invaluable, but they are also vulnerable to industrial espionage, disgruntled employees—both current and former—undesirables, sophisticated criminals, computer viruses, and more. The potential exposure is especially profound when the enterprise network is interfaced to the public, worldwide Internet.

In a 1996 Ernst & Young survey, conducted jointly with *Information Week* magazine, Information Technology (IT) professionals were asked about the importance of security to senior management. Not surprisingly, three-fourths of the respondents felt security was important or extremely important; a mere three percent believed taking security precautions was unnecessary. The same study revealed that over half of the respondents suffered some loss related to networked information security and disaster recovery.

Extremely Important: 35%
Important: 40%
Somewhat Important: 22%
Not Important: 3%

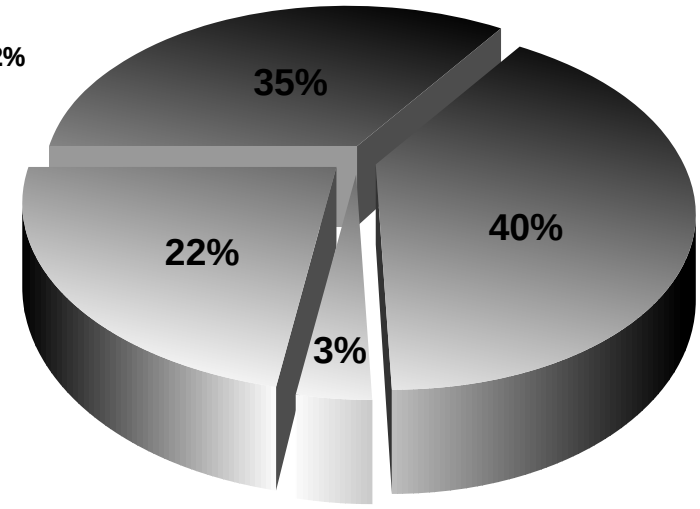

A recent Ernst & Young survey showed that nearly all IT professionals feel network security is important to their senior management.

In separate survey, conducted jointly by the Computer Security Institute of San Francisco and the FBI's International Computer Crime Squad, 42% of the respondents had experienced some unauthorized use of networked systems within the preceding 12 month period. Over half also reported their networks were being "probed" regularly from the outside, as well as from within. The need for network security is undeniable.

The Three P's of Security

Comprehensive network security is embodied in three areas that are easy to remember, as each begins with the letter ***P***:

- ***Protection of resources,*** which grants trusted users access only to those network resources they are authorized to use
- ***Proof of identity,*** which is necessary to ensure that *trusted* users—and *only* trusted users—are granted access to the network
- ***Privacy of information*** that is transmitted across the network, which assures strict confidentiality

There is a fourth area of security, and that is managing the three P's. This is equally important because when security is poorly managed, it becomes ineffective. For example, if passwords or keys are easily stolen, or firewalls are improperly configured, the result could be to render these fundamental elements practically useless.

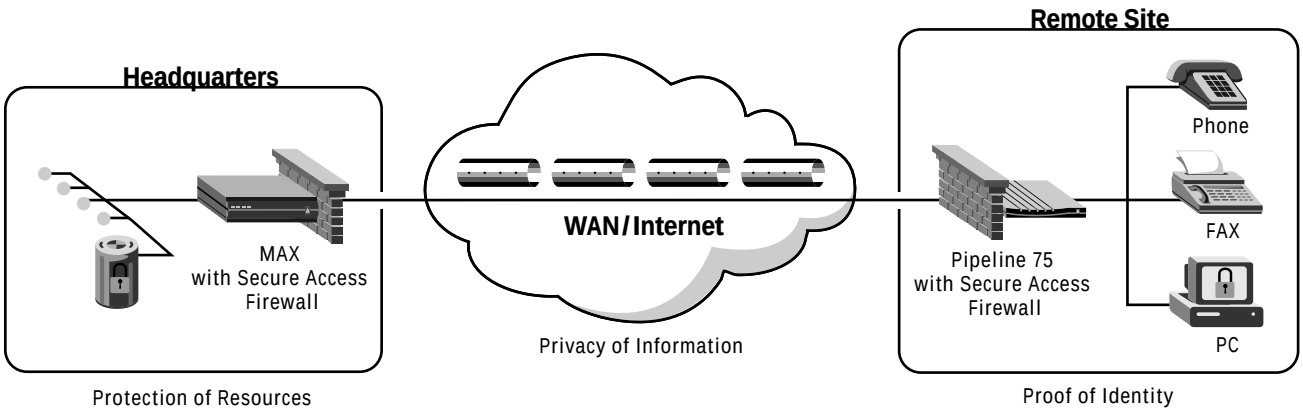

Properly implemented, the three P's of security create an impenetrable fence around any enterprise network.

Network-Enforced Security

Protection of resources is the area that comes to mind when most people think of security. Long before there were PCs and LANs, methods were available to make certain users were granted only "approved" access to system resources. Users logged on with simple passwords to mainframe and midrange computers to gain access to specific applications.

More recently, the network infrastructure itself has begun to play a vital role in securing network-attached hosts and servers. The various elements of security provided by the network itself include (relevant standards are listed where applicable):

- *Password identification* for remote access, which affords the minimum level of protection necessary for any network (Password Authorization Protocol)
- *User authentication,* which positively identifies legitimate users, provides tighter control of remote access (Challenge Handshake Authentication Protocol)
- *Token cards* offer virtually "bulletproof" remote access authentication with single-use passwords
- *Calling Line ID (CLID)* and *Callback* help ensure that only approved users, dialing from approved locations, are attempting to log on
- *Authorization* grants authenticated users permitted access only (RADIUS)
- *Packet authentication* adds integrity for transmitted data and, implicitly, its sender (IPSec's Authentication Header)
- *Encryption* keeps data strictly confidential during transit through the Internet (IPSec's Encapsulating Security Payload)
- *Network Address Translation* replaces internal IP addresses, which prevents them from being "discovered" via the Internet
- *"Demilitarized zones" and firewalls* isolate private resources—both clients and servers—from public access
- *Specialty applications-oriented solutions* for electronic commerce, Web browsing and e–mail (see sidebar)

IPSec's Encapsulating Security Payload and Authentication Header are covered at length in the next chapter. Some of the remaining network security elements, which are not self-explanatory, are addressed briefly here.

User authentication positively identifies legitimate or "trusted" users before granting access to the network. Common methods, all listed above, include PAP, CHAP, CLID, callback, and physical or software-based token cards for single-use or one-time passwords. All are handled by the network access equipment, and can be used individually or in combination.

Authorization is generally associated with hosts and servers. The user's ID or address, possibly validated with a password, determines access privileges. The level of protection can be as granular as read-only and/or modify capabilities down to individual fields in a database. But network access equipment can also play a vital role in authorization by granting authenticated users access only to approved hosts and servers. By relentlessly monitoring the source and destination addresses of network traffic, a capable network access switch detects violations and blocks packets to/from unauthorized devices. Internal users, for example, might be limited to accessing only their own department's resources. Or external users, such as customers and suppliers, might be restricted to just a few "public" servers.

Network Address Translation (NAT) was originally intended to substitute official Internet addresses for private, or unregistered, IP addresses. It turns out that this translation also offers an effective way to hide internal addresses from detection on the Internet. Without the addresses, would-be intruders have a difficult time even attempting access. A significant advantage of NAT is that many users are served with a relatively small allotment of Internet addresses, thus eliminating the need to assign every client, server and host a separate registered address. But even if all nodes already have regular Internet addresses, the process of substituting another on outbound traffic keeps a regular address from ever traversing the Internet as either a source or destination, which thwarts potential attacks.

A nearly identical address-hiding function is achieved with the combination of tunneling and encryption. Tunneling is normally used to "wrap" non-IP or private IP packets in an official IP packet for transport across the Internet, often as part of Virtual Private Network, or VPN. Popular standards include the Point-to-Point Tunneling Protocol (PPTP), created by Microsoft and Ascend, and the Layer-2 Tunneling Protocol (L2TP), which is a combination of PPTP and Cisco's Layer-2 Forwarding (L2F). If tunneling is performed by the network access equipment, then it is this equipment at each end which supplies the packet's source and destination addresses. The original packet, complete with its header containing the internal client and server addresses, becomes the tunneled packet's payload. IPSec's Encapsulating Security Payload (explained in the next chapter) encrypts or "scrambles" the original packet so that the internal addresses cannot be "discovered" while traversing the Internet. So when tunneling is combined with IPSec encryption, all internal addresses remain an internal secret.

"Demilitarized zones" (DMZs) and firewalls are used to isolate the private network from a public one, such as the Internet. The DMZ is where public resources, like Web servers, are located; DMZ resources are open to access by virtually anyone on the Internet. It is increasingly common for organizations to outsource their Web server needs to an Internet Service Provider (ISP), which locates the DMZ at the ISP's Point of Presence (POP). Between the DMZ and the private network is the firewall. The firewall stops all traffic to and from all "untrusted" users, which is everyone not explicitly authorized for access. The firewall can also be used to prevent internal users from accessing certain public resources on the Internet, such as popular "entertainment" Web sites.

There are no industry standards for firewalls, because firewalls do not need to interoperate with anything else to work. They simply function as barriers that pass or block incoming and outbound traffic. The firewall itself can be a stand-alone server or integral to the network access equipment. Integral firewalls have the twin advantages of being more effective and less expensive. With total freedom of implementation, at least four different types of firewalls are now available, shown in increasing order of sophistication:

- Packet filtering, which is a simple static means of examining traffic based on addresses and/or packet type
- Circuit-level gateways that provide "openings" for all approved sessions based on an assortment of criteria
- Proxy or application gateways that perform a more in-depth analysis of traffic, including the higher-level application
- Stateful inspection, which combines features of the other types to achieve a truly dynamic way of adapting to changing traffic patterns

Specialty Applications-Oriented Security Solutions

Several standards leverage various elements of security to satisfy the particular needs of specialized applications. These standards have two things in common. First, they are implemented in the application itself, and not in the network access equipment. Second, they complement, rather than replace, other general-purpose security mechanisms, such as IPSec or firewalls.

Two of the most well publicized are the Secure Electronic Transaction (SET) and Private Communications Technology (PCT). Both are used for electronic commerce or Internet commerce where consumers transmit their credit card numbers on-line to purchase goods and services. The user must be authenticated, the card number encrypted, and the whole end-to-end process managed thoroughly.

Special security solutions are also available for the two most popular Internet applications: the World Wide Web and e–mail. The Secure HyperText Transport Protocol (S-HTTP) and the Secure Multipurpose Internet Mail Extension (S/MIME) make Web and e–mail sessions, respectively, secure enough to transmit sensitive data with confidence. These and other Internet security provisions are built atop two common utilities: the Secure Sockets Layer (SSL) and Generic Security Services (GSS). SSL and GSS exist as application interfaces to the TCP/IP protocol stack.

Managing Network Security

Management tools, the fourth implied area of network security, are used to configure systems, monitor traffic, log and report of attempted security violations, and so on. The most popular standard is ***RADIUS,*** the Remote Authentication Dial-In User Service. Just as network security itself has three P's, RADIUS can be thought of as providing three A's of security management applications: Authentication, Authorization and Accounting.

RADIUS supports these applications with its database that maintains access profiles for all trusted users. The information in each user's profiles includes passwords (authentication), access privileges (authorization) and network usage (accounting). The network access equipment interacts with the RADIUS server securely, transparently and automatically. When a user attempts to log on remotely, the network access switch queries the RADIUS server to obtain that user's profile for authentication and authorization. Similarly, usage is logged by interacting with the Call Detail Reporting (CDR) feature of the network access switch to provide a complete accounting for billback or other purposes. Proxy RADIUS capability lets the RADIUS server at a Network Service Provider POP access an organization's RADIUS server to obtain any necessary user information, which is necessary to secure Internet-based VPNs. By making the management task easier, RADIUS makes the security measures more effective.

The only other major aspect of security management involves managing the keys used for IPSec's Authentication Header and Encapsulating Security Payload. Key management's twin task of creating and distributing keys is covered in the next chapter.

Ideally, the management tools will integrate the many interrelated aspects of network security. An increasingly popular option is to treat the firewall as the central element of security management based on its pivotal purpose. These advanced firewalls access the RADIUS database, and their management interfaces can be used to control IPSec and other security functions. For example, this approach is ideal for handling the problem of dynamically-assigned or "wildcard" addresses given by ISPs to members of an Internet-based VPN, whether mobile or tethered. The firewall control is able to match user-level authentication with network-level encryption and authentication.

3. IPSec's Role in Network Security

Long before the advent of the computer, there was a need to keep information confidential—particularly during the period of potential exposure while being exchanged from one trusted party to another. Rather than leave sensitive information in a form that could be read easily by anyone, someone devised a scheme to encode the material. Separately, the receiving trusted party was given the means to decode it. This was the beginning of the science of cryptography.

Essentially cryptography protects network-bound packets by "scrambling" the data and/or its header. Keys, which are long strings of characters, are required to lock (encrypt) and unlock (decrypt) portions of the packet. Without the keys, the content of an encrypted packet, including its source and destination addresses, is pure gibberish.

With some 2,000 years of history dating back to the Roman Empire, cryptography is now a tried and true method for preserving confidentiality. Computers, of course, have automated the laborious process of encrypting and decrypting information, and added substantial sophistication to both the algorithms and key management. These proven techniques are now readily available and quite affordable for data networking applications.

This chapter highlights IPSec technology, and examines IPSec's role in the three P's of network security.

How Cryptography Works

The basics of cryptography are the same whether used by the ancient Romans or the U.S. National Security Agency. The original information, known as ***cleartext,*** is transformed, or ***encrypted,*** at the sending end into a specially-coded equivalent, known as ***ciphertext.*** The opposite process occurs at the receiving end where the ciphertext is converted, or ***decrypted,*** back into the original cleartext message. ***Hashing*** is a similar operation that creates a special ***digital signature*** for validating the content and/or the sender of packets.

The encryption, decryption and hashing processes employ special equations called ***algorithms.*** An algorithm may be as simple as a one-to-one correspondence of letters or as elaborate as sophisticated mathematical equations requiring a computer to operate in real-time. The "unknown" in all of this is the ***key.*** The key is, essentially, an input variable that makes the outcome unique. Even when the algorithms are known, they are useless to an untrusted party who does not have the key (see sidebar on *Cracking the Code*). Herein lies the beauty of IPSec: the end-to-end procedure can be fully standardized without compromising the protection afforded.

Cracking the Code

Decrypting ciphertext without the key(s) is a trial-and-error process that is utterly impractical, even with the most powerful of computers. For example, the common 56-bit key could require as many as 72 quadrillion guesses, which is a million billion, or a one followed by 15 zeros (1,000,000,000,000,000). To crack this code, a sample packet would need to be decrypted with a trial key, then analyzed to determine if the results produced intelligible data—a task that, in itself, is quite difficult to automate with the wide variety of data formats these days. At the rate of a million tries per second, the endeavor could take over 2,000 years! Of course, the would-be decrypter could get lucky with the right guess early in the effort, but on average, the process generally takes (too) many years.

For organizations that want even more protection, key lengths up to 168 bits are available. In between is the common 128-bit key, which produces 340 trillion trillion trillion unique keys, or 340,000,000,000,000,000,000,000,000,000,000,000,000. The cost to crack such codes is so great that it exceeds any reasonable potential value, which serves as an effective deterrent to those who might contemplate the attempt.

There are two general categories of keys: public and secret. ***Public-key cryptography,*** also called asymmetric, is used for electronic commerce or other applications that must be both strictly confidential and publicly available. This form of cryptography often relies on a trusted third-party called a ***Certificate Authority*** (CA) to authenticate the principal parties—typically a merchant and a consumer—and oversee the exchange of keys. Up to three types of keys can be involved: ***public keys, private keys*** and ***session keys.*** Public keys are those published by any party seeking a secure exchange of information with another party, who also has a public key. In addition, each party has a private key, which is unpublished. The sending party needs the receiving party's public key to encrypt the packets; the receiving party's private key is needed to decrypt the packets. Sometimes the public and private keys are used to create and exchange a temporary session key for use by both parties during the exchange. These clever interactions solve the biggest dilemma of public privacy: how to have secure communications with an "open" exchange of keys. ***RSA,*** named after its inventors Rivest, Shamir and Adelman, is the dominant standard for public-key cryptography.

Secret-key cryptography, also called symmetric, is designed for use in private enterprise networks. With this form of cryptography, one of the principal parties—typically the IT group at the headquarters—creates and distributes all keys. The other party could be a remote employee for intra-organizational communications (an intranet), or a customer or supplier for inter-organizational communications (an extranet). Both ends use the same key pair, which can be assigned either for an enduring period of time or a single session. For networks that use the Internet Protocol, ***IP Security*** embodies the relevant standards. Session keys are short-lived versions of secret keys that can be negotiated at session start-up time using the ***Diffie-Hellman*** technique, also named after its inventors. When the session is over, or when the agreed-upon lifetime expires, the current session keys are abandoned and new ones are negotiated, if necessary.

IP Security

IP Security, or simply ***IPSec,*** consists of a fundamental architecture and a collection of Request for Comment (RFC) standards developed by the Internet Engineering Task Force's (IETF's) IP Security Working Group. Naturally, IPSec is not the only standard for Internet-related security; there are several special applications-oriented efforts in the works (see sidebar in the previous chapter). But IPSec is the solution when dependable, general-purpose security is needed for confidential communications via the Internet or a private IP backbone.

IPSec provides three distinct forms of protection for the transfer of private data via a public or private IP network, including the Internet (definitions courtesy of RFC 1825, the *Security Architecture for the Internet Protocol*):

- ***Authentication—*** The property of knowing that the data received is the same as the data that was sent and that the claimed sender is in fact the actual sender.
- ***Integrity*** — The property of ensuring that data is transmitted from source to destination without undetected alteration.
- ***Confidentiality*** — The property of communicating such that the intended recipients know what was being sent but unintended parties cannot determine what was sent.

To provide these three forms of protection, there are three basic elements in IPSec: the Authentication Header (AH); the Encapsulating Security Payload; (ESP); and the Internet Key Management Protocol (IKMP). AH and ESP can be used separately or in combination to achieve the desired level of protection.

An ***Authentication Header*** (AH) involves a keyed code placed in the headers of all packets. As the name implies, AH authenticates the user with a "digital signature" known only to the holders of the key(s). The signature is the unique result of "hashing" the packet through a special algorithm. AH also provides data integrity because any changes, however minor, to the payload during transmission are detected by the packet's hashed signature in the header. AH does not provide any confidentiality because it does not encrypt the packet's payload. The two most popular AH standards are the ***Message Digest*** version 5 (MD5) and the ***Secure Hash Algorithm*** version 1 (SHA-1). MD5 uses up to a 128-bit key; SHA-1 offers stronger protection with key lengths up to 160 bits. Standards for 96-bit MD5 and SHA-1 are expected soon.

The ***Encapsulating Security Payload*** keeps transmitted information strictly confidential by fully encrypting the data, or payload, in all packets. This prevents other users from "listening in" to the open exchange of information. Because only trusted users have the key(s), ESP also provides authentication and integrity. As with AH, integrity results from a mismatch between data received and the packet's checksum—here in cleartext—in the header. The dominant ESP standard is the ***Data Encryption Standard*** (DES). DES supports key lengths up to 56 bits. ***Triple DES*** (3DES) uses three sets of keys to encrypt, then decrypt and finally re-encrypt the payload, which is the equivalent of using a key of up to 168 bits long. Because ESP actually encrypts all data, it introduces more "overhead" and requires more processing time than AH, both of which can impact performance.

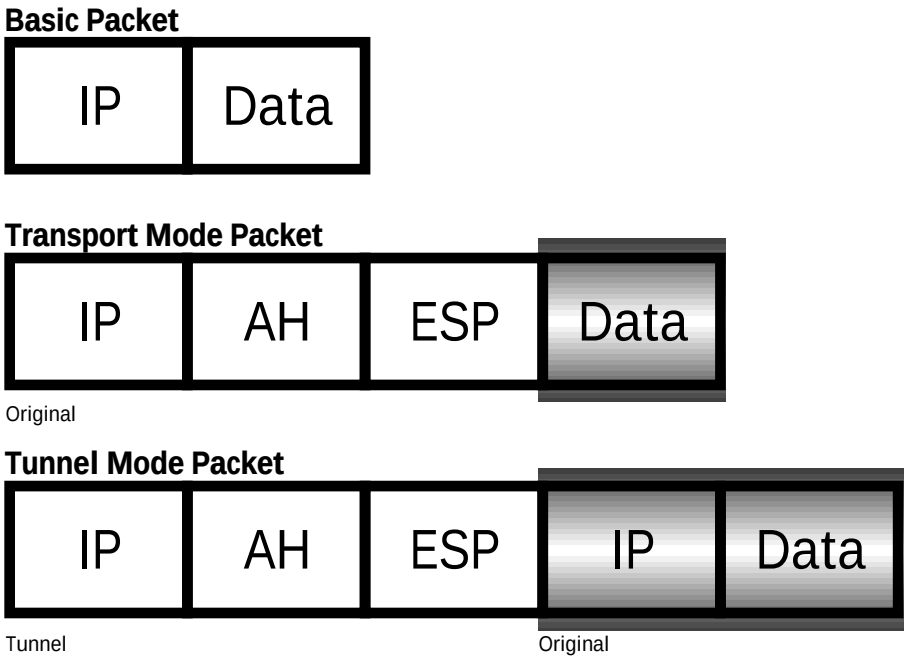

IPSec packets can be used in Tunnel mode which encapsulates original IP Addresses along with the Data.

Key management involves both the determination and distribution of keys. Up to four keys are needed: transmit and receive pairs for both AH and ESP. The key itself is a binary string of characters, normally represented in hexadecimal notation. For example, a 56-bit key might be 5F39DA752E0C25B4. Note that the total length is 64 bits (16 hexadecimal digits), which includes 8 bits of parity. A 56-bit key (DES) is adequate for most business applications.

Key management can be either manual or automated. Manual systems work fine for limited security needs, on a very small scale, where keys need only be changed every few months. Automated systems are required for all other applications.

With manual key management, keys are determined at a managing site then distributed to all remote users. The transmit/receive key or key pair is then entered, manually, into the equipment's configuration at each end. The actual keys can be calculated using a random number generator or simply "made up" arbitrarily. With manual systems, each key should be changed according to the organization's security policy.

Automated key management systems determine and distribute keys dynamically, transparently and, as the name implies, automatically. Like their manual counterpart, automated secret-key management systems have a central point of control. A centralized "key keeper" can itself be made more secure, which maximizes IPSec's effectiveness.

There is no need for a standard with manual key management, but some standardization is required for automated systems. This is because all network access equipment must interact (regularly and automatically) with the centralized key management system. ***Oakley*** specifies how keys are determined, while the Internet Security Association Key Management Protocol (***ISAKMP***) defines the method for distributing keys. Together ISAKMP/Oakley, also known as the ***Internet Key Management Protocol*** (IKMP), provide a complete and automated end-to-end key management system. Diffie-Hellman defines a way for two parties, that are not under the same secret-key control, to secretly exchange keys in a manner similar to that employed in public-key cryptography.

The IPSec Standards

IPSec is embodied in the following Request for Comment (RFC) standards, which are optional for IPv4 (the current version 4) and mandatory for IPv6 (version 6, sometimes called next-generation IP or IPng):

- *RFC 1825: Security Architecture for the Internet Protocol*
- *RFC 1826: IP Authentication Header*
- *RFC 1827: IP Encapsulating Security Payload (ESP)*
- *RFC 1828: IP Authentication Using Keyed MD5 (Message Digest)*
- *RFC 1829: The ESP DES-CBC Transform*
- *RFC 2085: HMAC-MD5 IP Authentication with Replay Prevention*
- *RFC 2104: HMAC: Keyed-Hashing for Message Authentication*

The Internet Key Management Protocol (IKMP) standards are expected to be finalized and assigned RFC numbers in mid-1997. Conformant vendor implementations should become available beginning in the second half of 1997. In the meantime, some organizations are using manual systems or rudimentary standards like Simple Keys for IP (SKIP).
RFCs related to IPSec, but not under the auspices of the IPSec Working Group include:

- *RFC 1321: The MD5 Message-Digest Algorithm*
- *RFC 1751: A Convention for Human-Readable 128-bit Keys*
- *RFC 1851: The ESP and Triple DES Transform*
- *RFC 1852: IP Authentication using Keyed SHA*

IPSec Implementations

A fundamental advantage to IPSec is that it can be implemented entirely in shared network access equipment, rather than in all hosts and servers. This approach eliminates the need to upgrade any network-attached resources. At the client end, the IPSec architecture permits use of either remote access routers or purely software-based implementations for PCs and workstations that use ordinary modems.

The shared network access switches that implement IPSec can be either on-site as Customer Premises Equipment (CPE) or off-site at the Internet Service Provider's (ISP's) Point of Presence (POP). ESP offers even greater flexibility with its two different modes of operation: transport mode and tunnel mode.

- ***Transport mode*** is normally used when ESP is implemented in a host (client or server), which serves as the tunnel end point. Transport mode uses the original cleartext IP header and encrypts only the data, including its Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) header.
- ***Tunnel mode*** is normally used when IPSec is implemented in network access equipment, which usually involves multiple hosts on the attached network. Tunnel mode treats the entire IP packet—complete with the full TCP/IP or UDP/IP header and data—as its payload. It then adds a new IP header with its own address as the source. When tunnel mode is used with CPE, it offers the additional advantage of concealing internal host/server and client addresses.

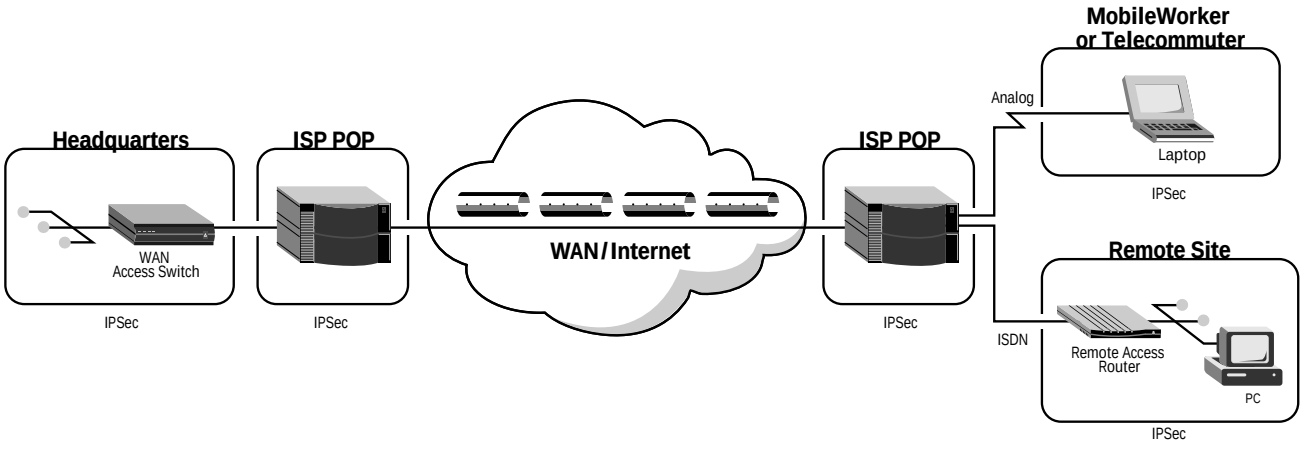

IPSec offers tremendous flexibility as to both how and where it is implemented.

With IPSec's many advantages comes the processing time and packet size "overhead" associated with encryption and hashing. Both have potential effects on system and/or network performance. Processing time is generally not an issue when IPSec is implemented in the network access equipment or client-based software. Network access switches can be sized to run the algorithms in real-time, and most client systems have ample CPU and memory resources to handle a single session with little or no degradation in performance—especially at modem speeds.

Encryption will, however, generally increase the size of packets. If the original cleartext packets are already at the maximum permitted for IP, the larger encrypted packets will become fragmented. IP routers routinely handle packet fragments, but there is a corresponding increase in the number of extra headers, thus adding to the overall traffic load. The impact can be mitigated by tightly controlling which traffic is encrypted through what means. For example, only confidential information needs to be fully encrypted with ESP. Most other traffic may require only AH protection, which adds very little overhead.

IPSec Interoperability

Standards are necessary for multivendor interoperability, but they are sometimes not sufficient. The standard itself may be incomplete or might permit various options. Incompatibilities also arise over different interpretations. To avoid such problems with IPSec, several "implementors workshops" now exist. One is the Secure Wide Area Network initiative known as S/WAN. S/WAN's objectives are:

- *Interoperability test readiness*
- *Synergy with TCP/IP vendors*
- *Education on standards*

Ad-hoc interoperability testing is also occurring on a regular basis. Vendors will either get together on their own, or are brought together by a joint customer implementing IPSec. Finally, owing to the IETF's philosophy of "build it first to prove it works—then publish the standard," IPSec already enjoys a solid foundation to guide future efforts.

Multi-national Issues

Organizations with offices throughout the world are accustomed to dealing with technology export restrictions imposed by the U.S. government. IPSec is one such technology. Currently, the use of 40-bit keys is permitted worldwide. Larger keys require the vendor to obtain permission and establish a procedure for key escrow. Naturally, different rulings are made for different countries, but most vendors can help companies export larger keys to international offices and subsidiaries. As of this writing, legislation is before the Congress to relax these restrictions. Check with your vendor for the most recent information in this area.

IPSec Applications

The need for IPSec depends on two factors: the sensitivity of the information itself; and the type of Wide Area Network (WAN) transporting it. Sensitivity ranges from "none", which is the case for publicly available information, to "strictly confidential", which constitutes an organization's most sensitive and proprietary information. The WAN might be a private network, an Internet-based virtual private network, or a combination of the two. Private network segments generally require less security than VPN segments based on public data networks.

Note that because an enterprise network consists of many separate segments, different IPSec options can be used for different types of segments. With firewall-based control of IPSec, different options can even be used on the same segment for different applications.

Private networks are not necessarily as private as the name implies, and therefore can benefit from the protection afforded by IPSec. Private networks that use, in whole or in part, the Public Switched Telephone Network (PSTN) are subject to the "openness" of this public resource. PSTN services encompass both the Integrated Services Digital Network (ISDN) and ever-popular analog modem. Leased lines too can be "tapped" physically, making them potentially vulnerable, as well. And cellular communications are quite insecure owing to the "broadcast" nature of the transmission. Government agencies and security-conscious organizations have an increasing desire to secure these so-called "private" links and the network access switch entry points to the enterprise network.

Virtual private networks are the reason IPSec exists. Indeed, IPSec is what puts the "private" in virtual private networking. VPNs use public data networks (PDNs), such as the Internet or Frame Relay, to carry private information. The Internet especially has no inherent security, so IPSec's ESP protection should be used to add confidentiality, as well as authentication and integrity for all intranet applications.

When an Internet-based VPN is used to create an extranet for private communications with customers and/or suppliers, the security formula changes. Because extranets are generally a more "open" environment, there is less need for strict confidentiality. But to keep the extranet from becoming too open, the solid user authentication provided by AH is advisable.

The Internet's worldwide presence, combined with its affordable access and ever-expanding capabilities, makes Internet-based VPNs ideal for "multimedia" applications. Multimedia involves the integration of data with voice and/or video for collaborative work sessions that employ shared whiteboards or full videoconferencing. Owing to the already-marginal performance of IP multicast and IP-based telephony services, however, IPSec should be used sparingly.

Worldwide and North American Headquarters
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, United States
Tel: 510.769.6001
Fax: 510.747.2300
E-mail: info@ascend.com
Toll Free: 800.621.9578
Fax Server: 415.688.4343
Web Site: http://www.ascend.com

European Headquarters
Rosemount House
Rosemount Avenue
West Byfleet
Surrey KT14 6NP, United Kingdom
Tel: +44 (0) 1932.350.115
Fax: +44 (0) 1932.350.199

Japan Headquarters
Level 19 Shinjuku Daiichi-Seimei Bldg.
2-7-1 Nishi-Shinjuku
Shinjuku-ku, Tokyo 163-07, Japan
Tel: +81.3.5325.7397
Fax: +81.3.5325.7399
Web Site: http://www.ascend.co.jp

Asia-Pacific Headquarters
Suite 1908, Bank of America Tower
12 Harcourt Road
Hong Kong
Tel: +852.2844.7600
Fax: +852.2810.0298

Latin, South America and the Caribbean Headquarters
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, United States
Tel: 510.769.6001
Fax: 510.747.2669

ASCEND

15-10-wp
09/97