VPN MAIN

Home

What is a VPN?

How do VPNs work?

Applications

News

Support Services

Product Information

White Papers

FAQ

VPN SOLUTIONS

Enterprise

Small Business

Service Provider

Authentication and Privacy in IPv4 and IPv6

Steve Martin
3Com Corporation

Abstract

This paper describes the security mechanisms for IP version 4 (IPv4) and IP version 6 (IPv6) and the services that they provide. An overview of key management requirements for systems implementing those security mechanisms will also be discussed.

There are two specific headers that are used to provide security services in IPv4 and IPv6. These headers are the IP Authentication Header (AH) and the IP Encapsulating Security Payload (ESP) header. The Authentication Header provides support for data integrity and authentication of IP packets as well as protection against replay attacks. The Encapsulating Security Payload header, like the Authentication Header, provides for data integrity, authentication, and replay protection but also provides for confidentiality through packet encryption.

Depending on the application, both AH and ESP can be used to either protect either a transport layer segment or the entire IP packet (transport-mode versus tunnel-mode).

Example applications of both the AH and the ESP will be discussed using the HMAC-MD5 algorithm to illustrate the AH and DES-CBC transform to illustrate the ESP. Descriptions will cover both AH and ESP used individually as well as in tandem.

A key concept that appears in both the authentication and privacy mechanisms for IP is the security association. This security association is uniquely identified by the internet destination address, the security protocol and a Security Parameter Index (SPI). The SPI is enclosed in both the AH and the ESP

header and is the method by which a key management mechanism is linked to the authentication and privacy mechanisms.

This loose coupling of key management systems allows for the use of existing systems while allowing for the development of future key systems without modification of the security mechanisms. Current key management systems mandated by the Internet Engineering Task Force (IETF) draft standards include manual management as well as the ISAKMP/OAKLEY key management and exchange protocol.

Note that many of the concepts described here have yet to be standardized and are to be considered as work in progress. Much of the available information upon which this paper is based is drawn from Internet Draft documents which are working documents of the IETF.

Background

The origins of todayís Internet are rooted in the ARPANET project of the late 1960's and early 1970's which was undertaken by a number of universities and corporations with funding by the Department of Defense's Advanced Research Projects Agency (DARPA). The primary purpose of this project was to develop a wide area data communications system that was highly resilient to nuclear attack. Unlike virtually all of the communications links of the day which were circuit switched, it was thought that a packet switched network would allow data traffic to be dynamically routed around damaged or broken links. The network started with four nodes in December of 1969 and in the ensuing years has grown to encompass millions of nodes around the world.

As the fledgling ARPANET grew and hosts were added, it soon became apparent that the existing protocols were inadequate for internetworking dissimiliar hosts separated by thousands of miles. In response to this need, the TCP/IP Reference Model was defined in 1974. The 1983 release of University of California Berkeleyís BSD UNIX 4.2 provided the first straightforward and reliable release of IP which became the standard communications protocol for the Internet. [2]

The original designers of the ARPANET and TCP/IP were primarily concerned about maintaining reliable communications in the face of hostile external attacks. The underlying packet switching architecture of IP, coupled with the end-to-end connectivity of TCP, provided this reliability. The designers did not, however, anticipate the security measures needed to protect the computers and Internet infrastructure from covert internal attack. As use of the Internet has grown, so has the incidence of internal attack. With the increasing use of the Internet as a vehicle for commerce via Web-based shopping and Virtual Private Networks, secure communications has become a topic of significant research. The work of standardizing this research into a form that has practical application to Internet security is being carried out by the IPSEC (IP Security) Working Group of the IETF.

General Requirements for Network Security

The requirements for computer and network security can be generalized into four categories: **secrecy, authentication and integrity**, and **availability** [1]. And each of these security requirements has associated with it one or more general classifications of threats against it.

Secrecy requires that information in a computer system or transmitted over a network can only be accessed by authorized parties. The requirement for secrecy is undermined by the threat of interception by unauthorized parties who can either release the otherwise private information or can engage in *traffic analysis* attacks.

Authentication requires that only authorized parties utilize computer and network facilities. The requirement for authentication is undermined by the threat of unauthorized parties breaking into computer and network facilities. Typical threats to authentication are *dictionary attacks on passwords*, *replay attacks* and *man-in-the-middle* attacks against session key distribution protocols.

Integrity requires that information in a computer system or transmitted over a network can only be modified by authorized parties. The requirement for integrity is undermined by the threat of modification

or fabrication of data. Typical threats to integrity are *replay* attacks and *man-in-the-middle* attacks.

Availability requires that access to computer and network facilities are never denied to authorized parties. The requirement for availability is undermined by the threat of interruption of service. Typical of this threat is *denial of service* attacks.

Note that IPv4 and IPv6 security mechanisms are primarily designed to address the requirements for secrecy, authentication and integrity and do little to prevent attacks on availability.

IPv4 and IPv6 Security Mechanisms

Security Associations

The concept of a Security Association (SA) is fundamental to both the IP Encapsulating Security Payload and the IP Authentication Header. The combination of a given Security Parameter Index (SPI), Security protocol and Destination Address uniquely identifies a particular SA. An implementation of the Authentication Header or the Encapsulating Security Payload must support this concept of a Security Association. A Security Association normally includes the parameters listed below [3], but might include additional parameters as well:

- Authentication algorithm and algorithm mode being used with the IP Authentication Header [required for AH implementations].
- Key(s) used with the authentication algorithm in use with the Authentication Header [required for AH implementations].
- Encryption algorithm, algorithm mode, and transform being used with the IP Encapsulating Security Payload [required for ESP implementations].
- Key(s) used with the encryption algorithm in use with the Encapsulating Security Payload [required for ESP implementations].
- Presence/absence and size of a cryptographic synchronization or initialization vector field for the encryption algorithm [required for ESP implementations].
- Authentication algorithm and mode used with the ESP transform (if any is in use) [recommended for ESP implementations].
- Authentication key(s) used with the authentication algorithm that is part of the ESP transform (if any) [recommended for ESP implementations].
- Lifetime of this Security Association [recommended for all implementations].
- Source Address(es) of the Security Association, might be a wildcard address if more than one sending system shares the same Security Association with the destination [recommended for all implementations].

The sending host uses the sending userID and Destination Address to select an appropriate Security Association (and hence SPI value). The receiving host uses the combination of SPI value, Security protocol and Destination Address to distinguish the correct association. Therefore, an AH implementation will always be able to use the SPI in combination with the Destination Address to determine the security association and related security configuration data for all valid incoming packets.

A security association is normally one-way. An authenticated communications session between two hosts will normally have two Security Parameter Indexes in use (one in each direction). The combination of a particular Security Parameter Index and a particular Destination Address uniquely identifies the Security

Association.

The receiver-orientation of the Security Association implies that the destination system will normally select the SPI value. By having the destination select the SPI value, there is no potential for manually configured Security Associations that conflict with automatically configured (e.g., via a key management protocol) Security Associations.[3]

Authentication Header (AH)

The IP Authentication Header is designed to provide integrity and authentication without confidentiality to IP datagrams. The lack of confidentiality ensures that implementations of the Authentication Header will be widely available on the Internet, even in locations where the export, import, or use of encryption to provide confidentiality is regulated. The Authentication Header supports security between two or more hosts implementing AH, between two or more gateways implementing AH, and between a host or gateway implementing AH and a set of hosts or gateways (Figure 1). A security gateway is a system which acts as the communications gateway between external untrusted systems and trusted hosts on their own subnetwork. It also provides security services for the trusted hosts when they communicate with the external untrusted systems. A trusted subnetwork contains hosts and routers that trust each other not to engage in active or passive attacks and trust that the underlying communications channel (e.g. Ethernet LAN) isn't being attacked. Trusted systems always should be trustworthy, but in practice they often are not.[3]

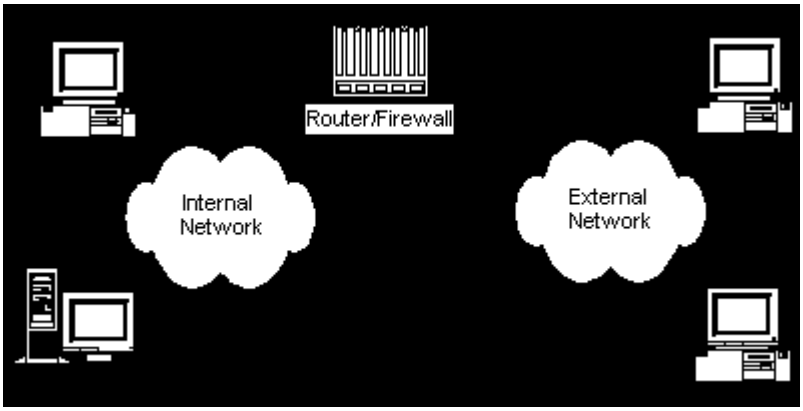

In the case where a security gateway is providing services on behalf of one or more hosts on a trusted subnet, the security gateway is responsible for establishing the security association on behalf of its trusted host and for providing security services between the security gateway and the external system(s). In this case, only the gateway need implement AH, while all of the systems behind the gateway on the trusted subnet may take advantage of AH services between the gateway and external systems.[3]

The IP Authentication Header holds authentication information for its IP datagram. It does this by computing a cryptographic authentication function over the IP datagram and using a secret authentication key in the computation. The sender computes the authentication data prior to sending the authenticated IP packet. Fragmentation occurs after the Authentication Header processing for outbound packets and re-assembly occurs prior to Authentication Header processing for inbound packets. The receiver verifies the correctness of the authentication data upon reception. Certain fields which must change in transit are omitted from the authentication calculation. Some authentication algorithms (e.g., asymmetric algorithms that use secret information known only to the sender in authentication calculations) might provide non-repudiation, but all authentication algorithms that might be used with the Authentication Header do not necessarily provide it. A compliant AH implementation must support the following algorithms: HMAC-MD5 (Hash Message Authentication Code - Message Digest 5), and HMAC-SHA-1 (Secure Hash Algorithm).[4] Confidentiality and traffic analysis protection are not provided by the Authentication Header.

Use of the Authentication Header will increase the IP protocol processing costs in participating systems and will also increase the communications latency. The increased latency is primarily due to the

calculation of the authentication data by the sender and the calculation and comparison of the authentication data by each receiver for each IP datagram containing an Authentication Header (AH).

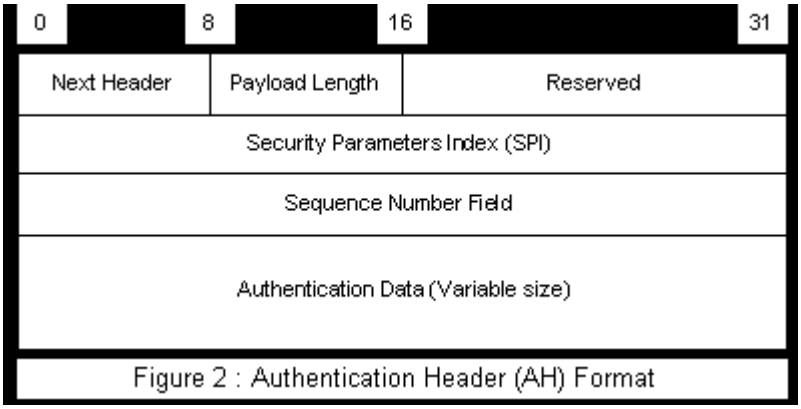

Figure 2 : Authentication Header (AH) Format

The fields of the Authentication Headers are as follows:

- **Next Header** - 8 bit field which identifies the type of the next payload after the AH
- **Payload Length** - 8 bit field which specifies the length of the AH in 32 bit words
- **Reserved** - 16 bit field which must be set to zero
- **Security Parameters Index (SPI)** - Arbitrary 32 bit value that in combination with the destination address identifies the Security Association for the datagram
- **Sequence Number Field** - Unsigned 32 bit field contains a monotonically increasing counter value for defense against replay attacks
- **Authentication Data** - Variable length field that contains the Integrity Check Value (ICV) for the packet

The Authentication Header may be used in two ways: transport-mode or tunnel-mode (Figure 1). Transport-mode is applicable only for host-to-host implementations and provides protection for upper layer protocols and selected IP header fields.

In transport mode, AH is inserted after the IP header and before an upper layer protocol (i.e. TCP) or other IPSEC headers have been inserted (Figure 3):

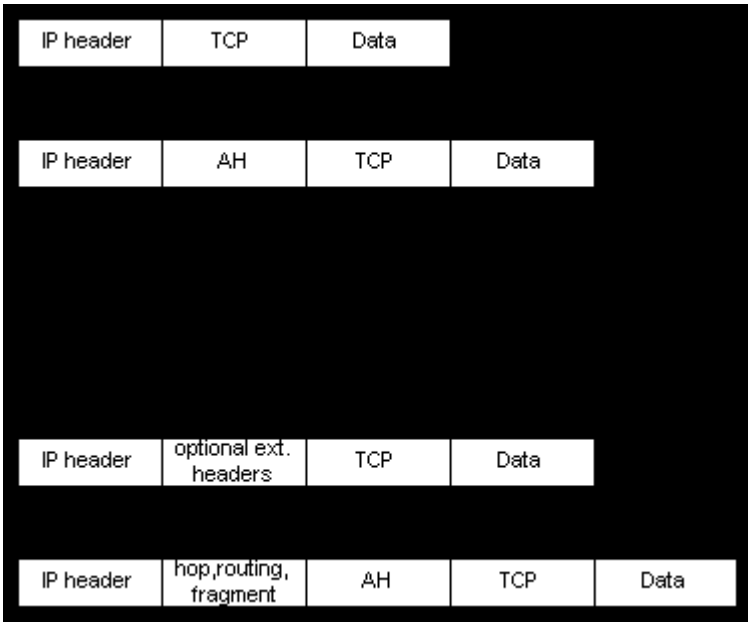

Tunnel-mode AH may be employed in either hosts or security gateways. When AH is used in security gateways to protect transit traffic, tunnel-mode must be used (Figure 1). In tunnel-mode, the inner IP header carries the ultimate source and destination address while an outer IP header contains the address of the security gateway (Figure 4). This allows the AH processing burden to placed on the gateway thereby relieving the hosts on the network.

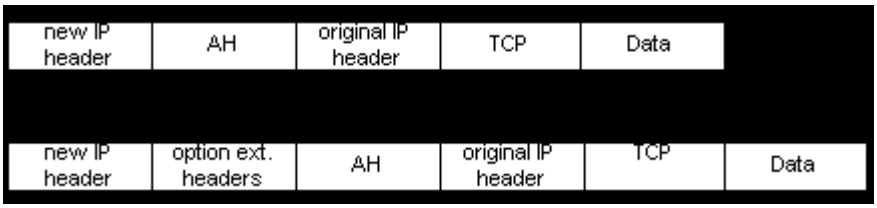

Encapsulating Security Payload (ESP)

The IP Encapsulating Security Payload (ESP) is designed to provide confidentiality and optional integrity and authentication to IP datagrams. The ESP supports security between two or more hosts implementing ESP, between two or more gateways implementing ESP, and between a host or gateway implementing ESP and a set of hosts and/or gateways.[3]

Gateway-to-gateway encryption is most valuable for building private virtual networks across an untrusted backbone such as the Internet. It does this by excluding outsiders. As such, it is often not a substitute for host-to-host encryption since it provides no protection from local threats on the local trusted network, and therefore the two types of encryption can be and often should be used together.

In the case where a security gateway is providing services on behalf of one or more hosts on a trusted subnet, the security gateway is responsible for establishing the security association on behalf of its trusted host and for providing security services between the security gateway and the external system(s). In this case, only the gateway need implement ESP, while all of the systems behind the gateway on the trusted subnet may take advantage of ESP services between the gateway and external systems.[3]

The Encapsulating Security Payload header holds encryption, replay, and authentication information for its IP datagram. If authentication is selected as part of the SA, encryption is done first followed by authentication. The encryption algorithm used is specified by the SA. ESP is designed to be used with symmetric encryption algorithms. A compliant ESP implementation must support DES-CBC for

encryption. The authentication function is computed over the IP datagram using a secret authentication key. Certain fields which must change in transit are omitted from the authentication calculation. Some authentication algorithms (e.g., asymmetric algorithms that use secret information known only to the sender in authentication calculations) might provide non-repudiation, but all authentication algorithms that might be used with ESP do not necessarily provide it. A compliant ESP implementation must support the following algorithms for authentication: HMAC-MD5 and HMAC-SHA-1. [5]

The encapsulating security approach used by ESP can noticeably impact network performance in participating systems, but use of ESP should not adversely impact routers or other intermediate systems that are not participating in the particular ESP association. Protocol processing in participating systems will be more complex when encapsulating security is used, requiring both more time and more processing power. Use of encryption will also increase the communications latency. The increased latency is primarily due to the encryption and decryption required for each IP datagram containing an Encapsulating Security Payload. The precise cost of ESP will vary with the specifics of the implementation, including the encryption algorithm, key size, and other factors. Hardware implementations of the encryption algorithm are recommended when high throughput is desired. [5]

The fields of the Encapsulating Security Payload header are as follows (Figure 5):

- **Security Parameters Index (SPI)** - Arbitrary 32 bit value that in combination with the destination address identifies the Security Association for the datagram
- **Sequence Number Field** - Unsigned 32 bit field contains a monotonically increasing counter value for defense against replay attacks
- **Payload Data** - Variable length field containing data described by the Next Header field. If the encryption algorithm requires an Initialization Vector then that would be contained here
- **Padding (0-255 bytes)** - May be required to satisfy requirements for encryption algorithm
- **Pad Length** - Indicates the number of pad bytes immediately preceding
- **Next Header** - Identifies the type of data contained in the Payload field.
- **Authentication Data** - Variable length field that contains the Integrity Check Value (ICV) for the packet

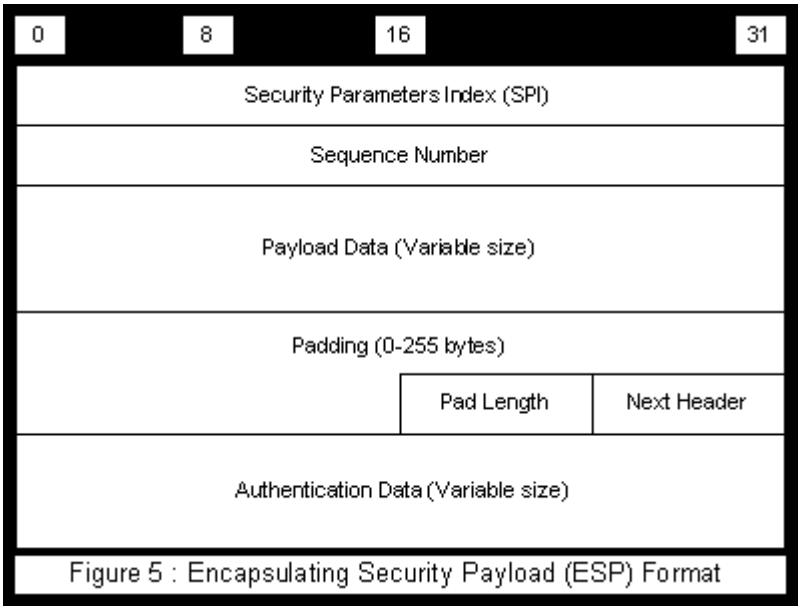

Figure 5 : Encapsulating Security Payload (ESP) Format

The ESP, like the AH, may be used in two ways: transport-mode or tunnel-mode (Figure 1).

Transport-mode is applicable only for host-to-host implementations and provides protection for upper layer protocols. Transport-mode encrypts the data carried by IP. Typically, this data is the transport layer segment such as TCP or UDP. For this mode, the ESP header is inserted into the IP packet immediately prior to the transport layer header (Figure 6).

Transport-mode operation provides privacy for any application that uses it, thus avoiding the need to implement privacy in every individual application. This mode of operation is also reasonably efficient, adding little to the total length of the IP packet.

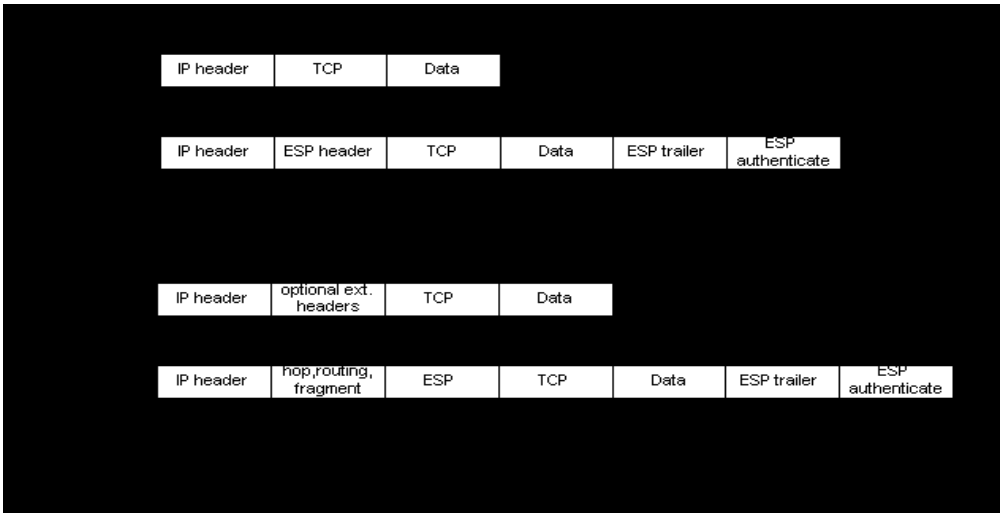

While transport-mode is suited for protecting connections between hosts that support the ESP feature, the tunnel-mode is useful in a configuration that includes a firewall or security gateway that protects a trusted network from external networks. In this case, encryption only occurs between an external host and the security gateway relieving the host(s) on the internal network of the processing burden of encryption. This also simplifies the key distribution task by reducing the number of needed keys and it discourages traffic analysis since the final destination is unknown.

Tunnel-mode is used to encrypt the entire IP packet. For this mode, the ESP is prefixed to the packet, and then the packet plus a trailing portion of the ESP header is encrypted (Figure 7). This method can be used to counter traffic analysis.

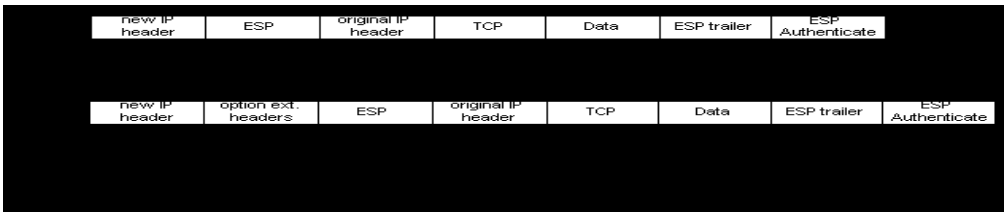

Authentication Header plus Encapsulating Security Payload

The Authentication Header and the Encapsulating Security Payload may be employed separately as shown in the previous sections, or they may be used in combination with one another.

The Authentication Header may be applied alone to provide services for data integrity and authentication of IP packets as well as for protection against replay attacks - even in locations where the export, import, or use of encryption to provide confidentiality is regulated.

The Encapsulating Security Payload may be applied alone to provide the same services as that of the Authentication Header plus encryption service. However, the authentication service provided by the ESP does not protect any IP header fields unless those fields are encapsulated by ESP in tunnel-mode. Therefore, it may be desirable to apply both the ESP and the AH for encryption with the highest level of authentication protection.

Manual Key Management and Exchange

The simplest form of key management is manual key management, where a person manually configures each system with its own key and also with the keys of other communicating systems. This is practical in small, static environments but does not scale and is prone to the keys being compromised since they are manually distributed. For example, within a small LAN it is practical to manually configure keys for each system. Another case is where an organization has an encrypting firewall between the internal network and the Internet at each of its sites and it connects two or more sites via the Internet. In this case, the encrypting firewall might selectively encrypt traffic for other sites within the organization using a manually configured key, while not encrypting traffic for other destinations. [3]

Another disadvantage of manual key management is that it renders the replay attack protection mechanism impractical to apply and therefore it must be disabled as part of the Security Association. This is particularly true for high speed links where the number of packets per unit time is significantly higher. The replay protection, when enabled, allows the receiver to verify that the sequence number contained in either the AH or the ESP is monotonically increasing as each packet belonging to a specific SA is received. When the sequence number reaches maximum and must rollover, a new SA must be established and the sequence number restarted at zero. The establishment of a new SA requires the exchange of a new set of keys. This becomes impractical in an environment which depends on manual key exchanges and has a high packet throughput.

ISAKMP/OAKLEY Key Management and Exchange

In addition to manual key management, the ISAKMP/OAKLEY key management and exchange protocol must be supported by standards compliant implementations.

ISAKMP is the Internet Security Association and Key Management Protocol. ISAKMP defines the procedures for creation and management of Security Associations. It also defines message exchanges that allow for peer authentication and threat mitigation. ISAKMP is designed as a common framework for agreeing to the format of SA attributes, and for negotiating, modifying, and deleting Security Associations.

It is distinct from any particular key exchange protocol and is designed so that different protocols can be specified depending on security requirements or improvements in key exchange protocols. [7]

OAKLEY is the key determination protocol currently mandated for compliant implementations by the IETF. The OAKLEY protocol is related to the Station-To-Station (STS) protocol. It shares the similarity of authenticating the Diffie-Hellman exponent exchange and their subsequent use in computing a shared key with STS. [8]

The authentication of the DH exponentials is required due to the vulnerability of the basic Diffie-Hellman mechanism to man-in-the-middle attacks. OAKLEY uses public key cryptography [6] or out of band, pre-shared symmetric keys to authenticate the participants in the DH exponential exchange.

The "Resolution of ISAKMP with OAKLEY" document presents a protocol that describes methods to obtain authenticated keying material for use with ISAKMP and the Security protocols, for the IETF IP Domain of Interpretation. It is a hybrid protocol based on ISAKMP, OAKLEY and SKEME. [9]

A typical ISAKMP/OAKLEY exchange is divided into two phases. The first phase establishes a security association and a session key between the ISAKMP peers. The second phase establishes SAs and session keys for the security protocols (AH and ESP).

Conclusion

The primary objective of the draft standards currently under development in the IPSEC Working Group of the IETF is to ensure that IPv4 and IPv6 will have solid cryptographic security mechanisms available to users who desire security. These mechanisms are designed to avoid adverse impacts on users who do not employ these security mechanisms for their traffic while allowing flexibility and standards-based implementations for applications that do have a requirement. These mechanisms are intended to be algorithm-independent so that the cryptographic algorithms can be altered without affecting the other parts of the implementation so that forward compatibility can be realized. The IPSEC drafts are continuing to move through the standardization process and should be completed in 1998.

While providing significant solutions to many of the security problems facing Internet access, these IP-layer mechanisms are not a complete panacea. For instance, they do not provide security against a number of traffic analysis attacks. Nor do they provide protection for attacks on data above the IP layer. Finally, there are serious performance issues in terms of hardware and software support for the algorithms that must be addressed.

References

[1] W. Stallings, *Data and Computer Communications,* Upper Saddle River, NJ: Prentice Hall, 1997.

[2] S. Garfinkel, G. Spafford, *Practical UNIX & Internet Security,* Sebastopol, CA: O'Reilly & Associates, Inc., 1996.

[3] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", IETF IPSEC Internet Draft, Nov. 1997.

[4] S. Kent, R. Atkinson, "IP Authentication Header", IETF IPSEC Internet Draft, Oct. 2, 1997

[5] S. Kent, R. Atkinson, "IP Encapsulating Security Payload (ESP)", IETF IPSEC Internet Draft, Oct. 2, 1997.

[6] B. Schneier, *Applied Cryptography,* New York, NY: John Wiley & Sons, 1996.

[7] D. Maughan, M. Schertler, M. Schneider, J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", IETF IPSEC Internet Draft, July 26, 1997.

[8] H.K. Orman, "The OAKLEY Key Determination Protocol", IETF IPSEC Internet Draft, ???.

[9] H. Krawczyk, "SKEME: A Versatile Secure Key Exchange Mechanism for Internet", from IEEE Proceedings of the 1996 Symposium on Network and Distributed Systems Security.

Additional Reading

A. Tanenbaum, *Computer Networks,* Upper Saddle River, NJ: Prentice Hall, 1996.

R. Atkinson, "Security Architecture for the Internet Protocol", IETF RFC 1825, 1995.

R. Atkinson, "IP Authentication Header", IETF RFC 1826, 1995.

R. Atkinson, "IP Encapsulating Security Payload (ESP)", IETF RFC 1827, 1995.

Biography

Steve Martin (smartin@usr.com) received an A.S. Communication Electronics from Cincinnati State Technical College (Cincinnati, OH) in 1979 and a B.S. Computer Engineering from San Jose State University (San Jose, CA) in 1990. He is currently working on a M.S. Computer Science at DePaul University (Chicago, IL) while employed as a Hardware Development Manager for the Carrier Systems Division of 3Com Corporation (Mt. Prospect, IL).