

IPsec

Trevor R.H. Clarke*
Rochester Institute of Technology
Rochester, NY, USA

November 4, 1998

*Tel: (716)475-6198 — retrev@csh.rit.edu

Abstract

Secure IP, or IPsec is a standard that provides authentication, verification, and encryption at the IP networking layer. This powerful technology has many uses, including virtual private networks (VPNs) which stretch across global networks. This paper will discuss the protocols and standards which apply to IPsec. It will present sample scenarios that utilize IPsec.

<i>CONTENTS</i>	3
-----------------	---

Contents

1 Introduction	4
2 Security Associations	4
3 The Authentication Header Protocol	5
4 The Encapsulation Security Payload Protocol	6
5 Key Management Protocols	7
6 Sample Applications	8
7 Additional Information	10

1 Introduction

IPsec is a set of protocols being developed by the Internet Engineering Task Force (IETF). It allows secure transmission of IP packets across physical networks. IPsec is divided into two main protocols, the authentication header (AH) protocol and the encapsulating security payload (ESP) protocol. The IETF makes the following definitions in [1].

- **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.
- **Encryption** – A mechanism commonly used to provide confidentiality.
- **Non-repudiation** – The property of a receiver being able to prove that the sender of some data did in fact send the data even though the sender might later desire to deny ever having sent the data.
- **Traffic Analysis** – The analysis of network traffic flow for the purpose of deducting information that is useful to an

adversary. Examples of such information are frequency of transmission, the identities of the conversing parties, sizes of packets, Flow Identifiers used, etc.

- **SPI** – Acronym for “Security Parameters Index.” An unstructured opaque index which is used in conjunction with the Destination Address to identify a particular Security Association.

2 Security Associations

A fundamental part of the IPsec protocols is the “Security Association.” An SPI and Destination Address uniquely identify a Security Association. A security association must provide the following but may provide more information.

- In an AH packet, the authentication algorithm and algorithm mode being used
- The key or keys used in the above mentioned algorithm
- In an ESP packet, the encryption algorithm, algorithm mode, and transform
- The key or keys for the above algorithm
- Any cryptographic synchronization or initialization vectors being used
- Other pertinent key information such as key lifetime
- The sensitivity level of the packets using a Security Association must be provided under certain circumstances. This

sensitivity level is often “unclassified,” “classified,” “secret,” or “top-secret” although any identifier may be used.

3 The Authentication Header Protocol

The Authentication Header (AH) provides authentication and integrity to IP datagrams. The Authentication Header is generated using all the IP datagram fields. Certain fields; most notably *time to live*, *hop count*, *ident*, and *fragment offset*; are assumed to be zero (0) for all Authentication Header calculations.

“The Authentication Header (AH) may appear after any other headers which are examined at each hop, and before any other headers which are not examined at an intermediate hop.” [2] When using an Authentication Header within an IP datagram the header preceding the Authentication Header must contain the value 51 in its Next Header field.

Next Header	Length	RESERVED
Security Parameters Index		
Authentication Data		

- **NEXT HEADER** – This identifies the next payload in the IP datagram. This

field is 8 bits wide. Legal values are defined in the most recent IANA RFC describing “*Assigned Numbers*.”

- **PAYLOAD LENGTH** – This is the length of the Authentication Data fields in 32 bit words. This field is 8 bits wide. The minimum legal value is 0.
- **RESERVED** – This field must be set to all zeros. This field is 16 bits wide and is reserved for future use.
- **SECURITY PARAMETERS INDEX (SPI)** – This field contains the SPI for the datagram. This field is 32 bits wide. For a description of the SPI see the section on Security Associations. A value of 0 indicates “no security association exists.”
- **AUTHENTICATION DATA** – This field contains the hash value of the datagram. This field contains a variable number of 32 bits words as defined in the **PAYLOAD LENGTH** field.

The Authentication Header is calculated using a message digest algorithm, usually MD5. Conventional checksums such as CRC-32 are not appropriate as they are considered cryptographically weak.

When processing an outgoing packet, the Security Parameter Index is determined. SPIs are unidirectional, therefore there are often two SPIs for each connection. SPIs are based on the userid of the sending process and the Destination Address field of

the IP packet. Other information may be encoded within the SPI if appropriate for the implementation.

There are two methods for keying SPIs: host oriented and user oriented. In host oriented keying, all processes on a given host will use the same Security Association for a given Destination Address. In this case, the userid used is often the Source Address. In user oriented keying, each user on a given host gets a unique Security Association for each Destination Address. Each process owned by a user may get a separate Security Association, or they may all share one Security Association. The userid is usually the login ID of the user, possibly modified to include a process identification.

Next, all fields modified during transit, such as Hop Count, Time to Live, and Header Checksum are set to zero. The entire IP packet is then hashed using MD5 or another secure hashing function. This information, along with the associated SPI, is used to generate the Authentication Header. All transit modifiable fields are reset to their appropriate values and the Authentication Header is added to the datagram.

Upon receipt of the datagram, the destination zeros the transit modifiable fields and deconstructs the Authentication Header. The datagram is hashed and the result is compared to the data section of the Authentication Header. If they match, the datagram is considered valid and authentic. If not, the packet is thrown out.

Currently, care should be taken when using Authentication Headers with ICMP packets. Due to MTU restrictions, authentication of these packets may be unreliable at best. For further information on using Authentication Headers and ICMP, see [2] [3] [4]

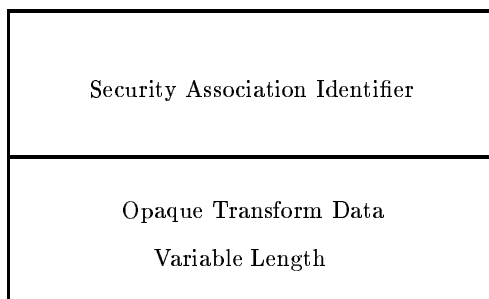
4 The Encapsulation Security Payload Protocol

The Encapsulation Security Payload (ESP) provides integrity and confidentiality to IP datagrams. ESP may be used to encrypt an entire IP datagram or just the transport-layer segment. Examples of transport-layer protocols are: TCP, UDP, ICMP, and IGMP. There are two modes ESP can operate in: tunnel-mode and transport mode.

“In Tunnel-mode ESP, the original IP datagram is placed in the encrypted portion of the Encapsulating Security Payload and that entire ESP frame is placed within a datagram having unencrypted IP headers.” [5] The unencrypted portion of the final datagram contains routing information for the IP tunnel. A description of various tunneling techniques is outside the scope of this document. A single, application specific, technique is discussed in the *Sample Applications* section of this document.

Transport-mode ESP encrypts the transport-layer protocol and inserts an ESP header immediately before the encrypted

protocol header. This mode conserves bandwidth since a new IP datagram is not generated.



“The Encapsulating Security Payload may appear anywhere after the IP headers and before the final transport-layer protocol.” [5] The header immediately preceding the ESP header should contain the value 50 in its Next Header or Protocol field.

The ESP header consists of an unencrypted section followed by an encrypted section. The encrypted section contains protected ESP header fields and the encrypted datagram or protocol frame.

The transform can be any cryptographically secure encryption function. All IPsec implementations must support the DES-CBC transform. [6]

The SPI is a 32 bit pseudorandom value as defined in [1]. If no SPI has been established, the value should be 0x00000000. SPI values in the range [0x00000001, 0x000000ff] inclusive are reserved to the Internet Assigned Numbers Authority (IANA) for future use.

5 Key Management Protocols

All IPsec key management protocols are based on the Internet Security Association and Key Management Protocol (ISAKMP). ISAKMP creates and manages Security Associations. The protocol is designed to be scalable and long lasting. Fields and definitions are generic enough to allow for any hashing or encryption function. ISAKMP uses two “phases” of negotiation. The first sets up a protected association between two ISAKMP servers. This assures the validity and privacy of negotiated Security Associations. Once an association is established between the ISAKMP servers, additional Security Associations are created and distributed. These are not limited to IPsec Security Associations, the protocol is designed to be expandable and broad based.

The OAKLEY key exchange protocol is less generic than ISAKMP but affords considerable security with little overhead. It uses the Diffie-Hellman key exchange algorithm to facilitate exchange of cryptographic secrets. The STS protocol is used to authenticate parties. OAKLEY can distribute keys using the traditional Diffie-Hellman exchange. It can also derive a new key from an existing key and distribute an externally derived key by encrypting it. OAKLEY is derived from ISAKMP and the Photuris exchange protocols. An implementation of the Photuris protocol is shipped with the OpenBSD operating system.

This document provides an overview of ISAKMP and OAKLEY. An in depth discussion of ISAKMP is presented in [7] and [8].

6 Sample Applications

Two sample applications based on the same theme will be presented. Both applications implement a point of presence (POP) within a LAN. This is a variation on a virtual private network (VPN).

6.1 Background and Situation

There are many active alumni of the Computer Science House at the Rochester Institute of Technology who routinely interact socially and technically with the current Computer Science House members. In order to facilitate this interaction I have developed a system that allows alumni to “mirror” one or more computers within the Computer Science House IP space. Two situations are supported: direct dial into the Computer Science House computer room and remote network access via a remote modem link or cable modem.

6.2 Implementation Introduction

A machine running OpenBSD is set up within the Computer Science House computer room. This machine will act as the point of presence for all participating members. It will also serve as a PPP gateway for modem

service. OpenBSD was chosen for a number of reasons. It is a free implementation of 4.4BSD with support for approximately 21 computing platforms. OpenBSD contains an implementation of the IPsec protocols which contains the required encryption systems as well as other, more efficient systems such as Blowfish. OpenBSD is based in Canada which has less strict export restrictions on strong cryptography than the United States. More information on OpenBSD is available at <http://www.openbsd.org>.

A remote LAN is set up using an RFC designated unforwardable subnet such as 192.168.x.x or 10.x.x.x. The third octet of the IP number is set to 60. This optional restriction makes the numbering scheme easier to understand since the CSH subnet is 129.21.60.x. For my examples, I will use 192.168.60.x¹ as the remote LAN subnet. The final octet for each machine in the 192.168.60.x network should be the same as the corresponding 129.21.60.x host number. This avoids an additional look up table for host number translation. For example:

DNS Name	CSH IP	Remote IP
gate1	129.21.60.3	192.168.60.3
win1	129.21.60.42	192.168.60.42
rmach1	129.21.60.43	192.168.60.43
rmach2	129.21.60.44	192.168.60.44
rmach3	129.21.60.45	192.186.60.45

Each host in the remote LAN also receives an IP number in the 129.21.60.x subnet. This is the IP number returned in DNS

¹An x in the final octet refers to a host address in the set (3,42,43,44,45)

lookups. The 192.168.60.x address is only used by the machines in the remote LAN.

The remote LAN must have a gateway to the CSH network. This may be a point to point connection via PPP or it may be a remote connection via cable modem or other high speed access system. This machine handles all authentication for the network, IP translation, and packet forwarding. This machine establishes an authenticated and possibly encrypted connection to the CSH network using IPsec.

6.3 Remote LAN Implementation Details

A sample LAN consists of four computers running OpenBSD or Windows 95. Each is connected to an Asante eight port hub via category-5 UTP. The host names and IP addresses are listed above. One of those machines acts as a gateway. Its hostname, for convenience, is gate1.csh.rit.edu. Each machine uses a netmask of 255.255.255.0. Each host's gateway (except for gate1) is set to gate1's IP address. All connections to local machines are ignored by gate1 and host name resolution is handled by a local DNS server. This server forwards unknown requests to 129.21.60.9, the CSH DNS server. A routing daemon, gated, is run on gate1 and is set to ignore requests to 192.168.60.x and forward all other requests. The gated daemon sends all forwarded requests through a pseudo-network driver called /dev/translate.

This driver decomposes the IP datagram and performs IP translations on the Source Address, Destination Address, and Hop Count (or Time to Live) fields. The address fields are modified in the following way. If the Destination Address is 129.21.60.x (an incoming packet) it is translated to 192.168.60.x and dropped onto the local network. If the Source Address is 192.168.60.x (an outgoing packet) it is translated to 129.21.60.x and passed to the IPsec network interface. In outgoing packets, the number of hops to the CSH point of presence is added to the time to live field. This makes the interface transparent and prevents packets from timing out. This value can be determined by occasionally sending a series of ICMP ECHO_REQUEST packets to the point of presence. This technique is similar to the technique used by the ping(1) and traceroute(1) UNIX commands.

The IPsec interface generates an Authentication Header (and possibly an ESP packet) for the datagram. This packet is forwarded to the CSH point of presence via a point to point connection such as PPP, or through a tunnel interface over the internet.

6.4 CSH Point of Presence Implementation Details

The CSH point of presence machine has, for convenience, the host name pop1.csh.rit.edu. This machine handles point to point (PPP) and tunnel connections to remote LANs. When a packet is received from a PPP or

tunnel interface, it is decrypted (if it is an ESP packet) and authenticated using the Authentication Header. If the packet is from a registered (I.E. valid) gateway, the packet is placed on the CSH network. If the packet is invalid, it is thrown out and the event is logged for future reference.

The pop1 machine also answers all traffic to 129.21.60.x. Any packets on the CSH network whose destination is in the above IP range are received by pop1. The number of hops to the remote LAN is added to the time to live field as in the above section. Authentication Headers are generated for the datagrams and they are optionally encrypted into an ESP packet. These new IPsec datagrams are forwarded to the appropriate LAN which authenticates as above.

6.5 Conclusion

This VPN system is scalable and secure. Appropriate security levels for each datagram can be specified allowing for authentication or authentication and encryption. This scheme prevents third party hosts from “spoofing” remote LAN IP addresses. This scheme is useful for remote LANs with no other network connection. Also, many cable modem providers only allocate one IP address per modem. This scheme allows a LAN located behind a cable modem to receive separate IP addresses and host names for each machine on the LAN.

The major downside to this scheme is efficiency. IPsec transforms can place a no-

ticeable load on gateways and network bandwidth, but this can be minimized if the minimal amount of security required is used. Also, with the cable modem variant, additional hops are added to the packet’s trip. This can significantly slow down network connections. To solve this, intelligent routing should be used. Any connection requiring a unique IP address should be routed through the above system. If a unique connection is not required, an alternate method should be employed. The most well known method is network address translation (NAT)². Using this technique, gate1 takes all outgoing packets and translates the source address to the address of the cable modem. All packets appear to come from a single address. This bypasses the CSH network and prevents additional hops from slowing down connections.

7 Additional Information

Additional information on required message hashing and encryption transforms can be found in [6] [9] [10] [11]. Information on using IPsec with multicast can be found in [12]. Information on using IPsec with ICMP datagrams can be found in [3] [4]. Miscellaneous other information can be found at [13] [14] [15] [16].

²Under Linux, NAT is usually referred to as IP Masquerading

References

- [1] R. Atkinson. *Security Architecture for the Internet Protocol*, August 1995. Category: Standards Track.
- [2] R. Atkinson. *IP Authentication Header (AH)*, August 1995. Category: Standards Track.
- [3] Michael Richardson. *IPv4 ICMP message and IPsec security gateways*, September 1998. Internet Draft: Expires in six months.
- [4] Michael Richardson. *Options for handling ICMP messages that must be forwarded*, September 1998. Internet Draft: Expires in six months.
- [5] R. Atkinson. *IP Encapsulating Security Payload (ESP)*, August 1995. Category: Standards Track.
- [6] P. Karn, P. Metzger, and W. Simpson. *The ESP DES-CBC Transform*, August 1995. Category: Standards Track.
- [7] Douglas Maughan, Mark Schertler, Mark Schneider, and Jeff Turner. *Internet Security Association and Key Management Protocol (ISAKMP)*, July 1998. Internet Draft: Expires in six months.
- [8] H. K. Orman. *The OAKLEY Key Determination Protocol*, 1998. Internet Draft: Expires in six months.
- [9] P. Metzger and W. Simpson. *IP Authentication using Keyed MD5*, August 1995. Category: Standards Track.
- [10] M. Oehler and R. Glenn. *HMAC-MD5 IP Authentication with Replay Prevention*, February 1997. Category: Standards Track.
- [11] H.L. Krawczyk, M. Bellare, and R. Canetti. *HMAC: Keyed-Hashing for Message Authentication*, February 1997. Category: Informational.
- [12] T. Hardjono, B. Cain, and N. Doraswamy. *A Framework for Group Key Management for Multicast Security*, July 1998. Internet Draft: Expires in six months.
- [13] Todd Spangler. Tech abc: Ipsec builds virtual bridges for security. *WebWeek*, November 10 1997. <http://www.internetworld.com/print/1997/11/10/infrastructure/19971110-bridges.html>.

- [14] Technical tips: Ip security. Web Page, May 17 1996.
<http://www.cisco.com/public/library/isakmp/ipsec.html>.
- [15] Naganand Doraswamy. *Implementation of Virtual Private Networks (VPNs) with IP Security*, March 1997. Internet Draft: Expires in six months.
- [16] Angelos D. Keromytis, John Ioannidis, and Jonathan M. Smith. *Implementing IPsec*. angelos@dsl.cis.upenn.edu, ji@reseach.att.com, jms@central.cis.upenn.edu, August 1997.