

Security Architecture for the Internet Protocol

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

1. INTRODUCTION

This memo describes the security mechanisms for IP version 4 (IPv4) and IP version 6 (IPv6) and the services that they provide. Each security mechanism is specified in a separate document. This document also describes key management requirements for systems implementing those security mechanisms. This document is not an overall Security Architecture for the Internet and is instead focused on IP-layer security.

1.1 Technical Definitions

This section provides a few basic definitions that are applicable to this document. Other documents provide more definitions and background information [VK83, HA94].

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 that data.

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.

Security Association

The set of security information relating to a given network connection or set of connections. This is described in detail below.

Traffic Analysis

The analysis of network traffic flow for the purpose of deducing 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. [Sch94].

1.2 Requirements Terminology

In this document, the words that are used to define the significance of each particular requirement are usually capitalised. These words are:

- MUST

This word or the adjective "REQUIRED" means that the item is an absolute requirement of the specification.

- SHOULD

This word or the adjective "RECOMMENDED" means that there might exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before taking a different course.

- MAY

This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor might choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

1.3 Typical Use

There are two specific headers that are used to provide security services in IPv4 and IPv6. These headers are the "IP Authentication Header (AH)" [Atk95a] and the "IP Encapsulating Security Payload (ESP)" [Atk95b] header. There are a number of ways in which these IP security mechanisms might be used. This section describes some of the more likely uses. These descriptions are not complete or exhaustive. Other uses can also be envisioned.

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. 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., an Ethernet) isn't being attacked.

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.

A security gateway which receives a datagram containing a recognised sensitivity label, for example IPSO [Ken91], from a trusted host should take that label's value into consideration when creating/selecting an Security Association for use with AH between the gateway and the external destination. In such an environment, a gateway which receives a IP packet containing the IP Encapsulating Security Payload (ESP) should add appropriate authentication, including implicit (i.e., contained in the Security Association used) or explicit label information (e.g., IPSO), for the decrypted packet that it forwards to the trusted host that is the ultimate destination. The IP Authentication Header should always be used on packets containing explicit sensitivity labels to ensure end-to-end

label integrity. In environments using security gateways, those gateways MUST perform address-based IP packet filtering on unauthenticated packets purporting to be from a system known to be using IP security.

The IP Encapsulating Security Payload (ESP) is designed to provide integrity, authentication, and confidentiality to IP datagrams [Atk95b]. 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. A security gateway is a system which acts as the communications gateway between external untrusted systems and trusted hosts on their own subnetwork and provides security services for the trusted hosts when they communicate with 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., an Ethernet) isn't being attacked. Trusted systems always should be trustworthy, but in practice they often are not trustworthy.

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, and indeed the two 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.

A gateway which receives a datagram containing a recognised sensitivity label from a trusted host should take that label's value into consideration when creating/selecting a Security Association for use with ESP between the gateway and the external destination. In such an environment, a gateway which receives a IP packet containing the ESP should appropriately label the decrypted packet that it forwards to the trusted host that is the ultimate destination. The IP Authentication Header should always be used on packets containing explicit sensitivity labels to ensure end-to-end label integrity.

If there are no security gateways present in the connection, then two end systems that implement ESP may also use it to encrypt only the user data (e.g., TCP or UDP) being carried between the two systems. ESP is designed to provide maximum flexibility so that users may select and use only the security that they desire and need.

Routing headers for which integrity has not been cryptographically protected SHOULD be ignored by the receiver. If this rule is not strictly adhered to, then the system will be vulnerable to various kinds of attacks, including source routing attacks [Bel89] [CB94] [CERT95].

While these documents do not specifically discuss IPv4 broadcast, these IP security mechanisms MAY be used with such packets. Key distribution and Security Association management are not trivial for broadcast applications. Also, if symmetric key algorithms are used the value of using cryptography with a broadcast packet is limited because the receiver can only know that the received packet came from one of many systems knowing the correct key to use.

1.4 Security Associations

The concept of a "Security Association" is fundamental to both the IP Encapsulating Security Payload and the IP Authentication Header. The combination of a given Security Parameter Index (SPI) and Destination Address uniquely identifies a particular "Security Association". An implementation of the Authentication Header or the Encapsulating Security Payload MUST support this concept of a Security Association. An implementation MAY also support other parameters as part of a Security Association. A Security Association normally includes the parameters listed below, 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 synchronisation or initialisation 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 the key or time when key change should occur [RECOMMENDED for all 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].
- Sensitivity level (for example, Secret or Unclassified) of the protected data [REQUIRED for all systems claiming to provide multi-level security, RECOMMENDED for all other systems].

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 and Destination Address to distinguish the correct association. Hence, 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. When a formerly valid Security Association becomes invalid, the destination system(s) SHOULD NOT immediately reuse that SPI value and instead SHOULD let that SPI value become stale before reusing it for some other Security Association.

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 Destination Address may be a unicast address or a multicast group address.

The receiver-orientation of the Security Association implies that, in the case of unicast traffic, 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. For multicast traffic, there are multiple destination systems but a single destination multicast group, so some system or person will need to select SPIs on behalf of that multicast group and then communicate the information to all of the legitimate members of that multicast group via mechanisms not defined here.

Multiple senders to a multicast group MAY use a single Security Association (and hence Security Parameter Index) for all traffic to that group. In that case, the receiver only knows that the message came from a system knowing the security association data for that multicast group. A receiver cannot generally authenticate which system sent the multicast traffic when symmetric algorithms (e.g., DES, IDEA) are in use. Multicast traffic MAY also use a separate Security Association (and hence SPI) for each sender to the multicast group. If each sender has its own Security Association and asymmetric algorithms are used, then data origin authentication is also a provided service.

2. DESIGN OBJECTIVES

This section describes some of the design objectives of this security architecture and its component mechanisms. The primary objective of this work 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 Internet users who do not employ these security mechanisms for their traffic. These mechanisms are intended to be algorithm-independent so that the cryptographic algorithms can be altered without affecting the other parts of the implementation. These security mechanisms should be useful in enforcing a variety of security policies.

Standard default algorithms (keyed MD5, DES CBC) are specified to ensure interoperability in the global Internet. The selected default algorithms are the same as the standard default algorithms used in SNMPv2 [GM93].

3. IP-LAYER SECURITY MECHANISMS

There are two cryptographic security mechanisms for IP. The first is the Authentication Header which provides integrity and authentication without confidentiality [Atk95a]. The second is the Encapsulating Security Payload which always provides confidentiality, and (depending on algorithm and mode) might also provide integrity and authentication [Atk95b]. The two IP security mechanisms may be used together or separately.

These IP-layer mechanisms do not provide security against a number of traffic analysis attacks. However, there are several techniques outside the scope of this specification (e.g., bulk link encryption) that might be used to provide protection against traffic analysis [VK83].

3.1 AUTHENTICATION HEADER

The IP Authentication Header holds authentication information for its IP datagram [Atk95a]. 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 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, such as the "TTL" (IPv4) or "Hop Limit" (IPv6) field, which is decremented on each hop, are omitted from the authentication calculation. However the omission of the Hop Limit field does not adversely impact the security provided. Non-repudiation might be provided by some authentication algorithms (e.g., asymmetric algorithms when both sender and receiver keys are used in the authentication calculation) used with the Authentication Header, but it is not necessarily provided by all authentication algorithms that might be used with the Authentication Header. The default authentication algorithm is keyed MD5, which, like all symmetric algorithms, cannot provide non-repudiation by itself. 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).

The Authentication Header provides much stronger security than exists in most of the current Internet and should not affect exportability or significantly increase implementation cost. While the Authentication Header might be implemented by a security gateway on behalf of hosts on a trusted network behind that security gateway, this mode of operation is not encouraged. Instead, the Authentication Header should be used from origin to final destination.

All IPv6-capable hosts MUST implement the IP Authentication Header with at least the MD5 algorithm using a 128-bit key. IPv4-systems claiming to implement the Authentication Header MUST implement the IP Authentication Header with at least the MD5 algorithm using a 128-bit key [MS95]. An implementation MAY support other authentication algorithms in addition to keyed MD5.

3.2 ENCAPSULATING SECURITY PAYLOAD

The IP Encapsulating Security Payload (ESP) is designed to provide integrity, authentication, and confidentiality to IP datagrams [Atk95b]. It does this by encapsulating either an entire IP datagram or only the upper-layer protocol (e.g., TCP, UDP, ICMP) data inside the ESP, encrypting most of the ESP contents, and then appending a new cleartext IP header to the now encrypted Encapsulating Security Payload. This cleartext IP header is used to carry the protected data through the internetwork.

3.2.1 Description of the ESP Modes

There are two modes within ESP. The first mode, which is known as Tunnel-mode, encapsulates an entire IP datagram within the ESP header. The second mode, which is known as Transport-mode, encapsulates an upper-layer protocol (for example UDP or TCP) inside ESP and then prepends a cleartext IP header.

3.2.2 Usage of ESP

ESP works between hosts, between a host and a security gateway, or between security gateways. This support for security gateways permits trustworthy networks behind a security gateway to omit encryption and thereby avoid the performance and monetary costs of encryption, while still providing confidentiality for traffic transiting untrustworthy network segments. When both hosts directly implement ESP and there is no intervening security gateway, then they may use the Transport-mode (where only the upper layer protocol data (e.g., TCP or UDP) is encrypted and there is no encrypted IP header). This mode reduces both the bandwidth consumed and the protocol processing costs for users that don't need to keep the entire IP datagram confidential.

ESP works with both unicast and multicast traffic.

3.2.3 Performance Impacts of ESP

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.

For interoperability throughout the worldwide Internet, all conforming implementations of the IP Encapsulating Security Payload MUST support the use of the Data Encryption Standard (DES) in Cipher-Block Chaining (CBC) Mode as detailed in the ESP specification. Other confidentiality algorithms and modes may also be implemented in addition to this mandatory algorithm and mode. Export and use of encryption are regulated in some countries [OTA94].

3.3 COMBINING SECURITY MECHANISMS

In some cases the IP Authentication Header might be combined with the IP Encapsulating Security Protocol to obtain the desired security properties. The Authentication Header always provides integrity and authentication and can provide non-repudiation if used with certain authentication algorithms (e.g., RSA). The Encapsulating Security Payload always provides integrity and confidentiality and can also provide authentication if used with certain authenticating encryption algorithms. Adding the Authentication Header to a IP datagram prior to encapsulating that datagram using the Encapsulating Security Protocol might be desirable for users wishing to have strong integrity, authentication, confidentiality, and perhaps also for users who require strong non-repudiation. When the two mechanisms are combined, the placement of the IP Authentication Header makes clear which part of the data is being authenticated. Details on combining the two mechanisms are provided in the IP Encapsulating Security Payload specification [At94b].

3.4 OTHER SECURITY MECHANISMS

Protection from traffic analysis is not provided by any of the security mechanisms described above. It is unclear whether meaningful protection from traffic analysis can be provided economically at the Internet Layer and it appears that few Internet users are concerned about traffic analysis. One traditional method for protection against traffic analysis is the use of bulk link encryption. Another technique is to send false traffic in order to increase the noise in the data provided by traffic analysis. Reference [VK83] discusses traffic analysis issues in more detail.

4. KEY MANAGEMENT

The Key Management protocol that will be used with IP layer security is not specified in this document. However, because the key management protocol is coupled to AH and ESP only via the Security Parameters Index (SPI), we can meaningfully define AH and ESP without having to fully specify how key management is performed. We envision that several key management systems will be usable with these specifications, including manual key configuration. Work is ongoing within the IETF to specify an Internet Standard key management protocol.

Support for key management methods where the key management data is carried within the IP layer is not a design objective for these IP-layer security mechanisms. Instead these IP-layer security mechanisms will primarily use key management methods where the key management data will be carried by an upper layer protocol, such as UDP or TCP, on some specific port number or where the key management data will be distributed manually.

This design permits clear decoupling of the key management mechanism from the other security mechanisms, and thereby permits one to substitute new and improved key management methods without having to modify the implementations of the other security mechanisms. This separation of mechanism is clearly wise given the long history of subtle flaws in published key management protocols [NS78, NS81]. What follows in this section is a brief discussion of a few alternative approaches to key management. Mutually consenting systems may additionally use other key management approaches by private prior agreement.

4.1 Manual Key Distribution

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 quite practical in

small, static environments but does not scale. It is not a viable medium-term or long-term approach, but might be appropriate and useful in many environments in the near-term. For example, within a small LAN it is entirely practical to manually configure keys for each system. Within a single administrative domain it is practical to configure keys for each router so that the routing data can be protected and to reduce the risk of an intruder breaking into a router. Another case is where an organisation 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 organisation using a manually configured key, while not encrypting traffic for other destinations. It also might be appropriate when only selected communications need to be secured.

4.2 Some Existing Key Management Techniques

There are a number of key management algorithms that have been described in the public literature. Needham & Schroeder have proposed a key management algorithm which relies on a centralised key distribution system [NS78, NS81]. This algorithm is used in the Kerberos Authentication System developed at MIT under Project Athena [KB93]. Diffie and Hellman have devised an algorithm which does not require a centralised key distribution system [DH76]. Unfortunately, the original Diffie-Hellman technique is vulnerable to an active "man in the middle" attack [Sch93, p. 43-44]. However, this vulnerability can be mitigated by using signed keys to authentically bootstrap into the Diffie-Hellman exchange [Sch93, p. 45].

4.3 Automated Key Distribution

Widespread deployment and use of IP security will require an Internet-standard scalable key management protocol. Ideally such a protocol would support a number of protocols in the Internet protocol suite, not just IP security. There is work underway within the IETF to add signed host keys to the Domain Name System [EK94]. The DNS keys enable the originating party to authenticate key management messages with the other key management party using an asymmetric algorithm. The two parties would then have an authenticatable communications channel that could be used to create a shared session key using Diffie-Hellman or other means [DH76] [Sch93].

4.4 Keying Approaches for IP

There are two keying approaches for IP. The first approach, called host-oriented keying, has all users on host 1 share the same key for use on traffic destined for all users on host 2. The second approach, called user-oriented keying, lets user A on host 1 have one

or more unique session keys for its traffic destined for host 2; such session keys are not shared with other users on host1. For example, user A's ftp session might use a different key than user A's telnet session. In systems claiming to provide multi-level security, user A will typically have at least one key per sensitivity level in use (e.g., one key for UNCLASSIFIED traffic, a second key for SECRET traffic, and a third key for TOP SECRET traffic). Similarly, with user-oriented keying one might use separate keys for information sent to a multicast group and control messages sent to the same multicast group.

In many cases, a single computer system will have at least two mutually suspicious users A and B that do not trust each other. When host-oriented keying is used and mutually suspicious users exist, it is sometimes possible for user A to determine the host-oriented key via well known methods, such as a Chosen Plaintext attack. Once user A has improperly obtained the key in use, user A can then either read user B's encrypted traffic or forge traffic from user B. When user-oriented keying is used, certain kinds of attack from one user onto another user's traffic are not possible.

IP Security is intended to be able to provide Authentication, Integrity, and Confidentiality for applications operating on connected end systems when appropriate algorithms are in use. Integrity and Confidentiality can be provided by host-oriented keying when appropriate dynamic key management techniques and appropriate algorithms are in use. However, authentication of principals using applications on end-systems requires that processes running applications be able to request and use their own Security Associations. In this manner, applications can make use of key distribution facilities that provide authentication.

Hence, support for user-oriented keying SHOULD be present in all IP implementations, as is described in the "IP Key Management Requirements" section below.

4.5 Multicast Key Distribution

Multicast key distribution is an active research area in the published literature as of this writing. For multicast groups having relatively few members, manual key distribution or multiple use of existing unicast key distribution algorithms such as modified Diffie-Hellman appears feasible. For very large groups, new scalable techniques will be needed. The use of Core-Based Trees (CBT) to provide session key management as well as multicast routing might be an approach used in the future [BFC93].

4.6 IP Key Management Requirements

This section defines key management requirements for all IPv6 implementations and for those IPv4 implementations that implement the IP Authentication Header, the IP Encapsulating Security Payload, or both. It applies equally to the IP Authentication Header and the IP Encapsulating Security Payload.

All such implementations MUST support manual configuration of Security Associations.

All such implementations SHOULD support an Internet standard Security Association establishment protocol (e.g., IKMP, Photuris) once such a protocol is published as an Internet standards-track RFC.

Implementations MAY also support other methods of configuring Security Associations.

Given two endpoints, it MUST be possible to have more than one concurrent Security Association for communications between them. Implementations on multi-user hosts SHOULD support user granularity (i.e., "user-oriented") Security Associations.

All such implementations MUST permit the configuration of host-oriented keying.

A device that encrypts or authenticates IP packets originated other systems, for example a dedicated IP encryptor or an encrypting gateway, cannot generally provide user-oriented keying for traffic originating on other systems. Such systems MAY additionally implement support for user-oriented keying for traffic originating on other systems.

The method by which keys are configured on a particular system is implementation-defined. A flat file containing security association identifiers and the security parameters, including the key(s), is an example of one possible method for manual key distribution. An IP system MUST take reasonable steps to protect the keys and other security association information from unauthorised examination or modification because all of the security lies in the keys.

5. USAGE

This section describes the possible use of the security mechanisms provided by IP in several different environments and applications in order to give the implementer and user a better idea of how these mechanisms can be used to reduce security risks.

5.1 USE WITH FIREWALLS

Firewalls are not uncommon in the current Internet [CB94]. While many dislike their presence because they restrict connectivity, they are unlikely to disappear in the near future. Both of these IP mechanisms can be used to increase the security provided by firewalls.

Firewalls used with IP often need to be able to parse the headers and options to determine the transport protocol (e.g., UDP or TCP) in use and the port number for that protocol. Firewalls can be used with the Authentication Header regardless of whether that firewall is party to the appropriate Security Association, but a firewall that is not party to the applicable Security Association will not normally be able to decrypt an encrypted upper-layer protocol to view the protocol or port number needed to perform per-packet filtering OR to verify that the data (e.g., source, destination, transport protocol, port number) being used for access control decisions is correct and authentic. Hence, authentication might be performed not only within an organisation or campus but also end to end with remote systems across the Internet. This use of the Authentication Header with IP provides much more assurance that the data being used for access control decisions is authentic.

Organisations with two or more sites that are interconnected using commercial IP service might wish to use a selectively encrypting firewall. If an encrypting firewall were placed between each site of a company and the commercial IP service provider, the firewall could provide an encrypted IP tunnel among all the company's sites. It could also encrypt traffic between the company and its suppliers, customers, and other affiliates. Traffic with the Network Information Center, with public Internet archives, or some other organisations might not be encrypted because of the unavailability of a standard key management protocol or as a deliberate choice to facilitate better communications, improved network performance, and increased connectivity. Such a practice could easily protect the company's sensitive traffic from eavesdropping and modification.

Some organisations (e.g., governments) might wish to use a fully encrypting firewall to provide a protected virtual network over commercial IP service. The difference between that and a bulk IP encryption device is that a fully encrypting firewall would provide filtering of the decrypted traffic as well as providing encryption of IP packets.

5.2 USE WITH IP MULTICAST

In the past several years, the Multicast Backbone (MBONE) has grown rapidly. IETF meetings and other conferences are now regularly multicast with real-time audio, video, and whiteboards. Many people are now using teleconferencing applications based on IP Multicast in the Internet or in private internal networks. Others are using IP multicasting to support distributed simulation or other applications. Hence it is important that the security mechanisms in IP be suitable for use in an environment where multicast is the general case.

The Security Parameters Indexes (SPIs) used in the IP security mechanisms are receiver-oriented, making them well suited for use in IP multicast [Atk95a, Atk95b]. Unfortunately, most currently published multicast key distribution protocols do not scale well. However, there is active research in this area. As an interim step, a multicast group could repeatedly use a secure unicast key distribution protocol to distribute the key to all members or the group could pre-arrange keys using manual key distribution.

5.3 USE TO PROVIDE QOS PROTECTION

The recent IAB Security Workshop identified Quality of Service protection as an area of significant interest [BCCH]. The two IP security mechanisms are intended to provide good support for real-time services as well as multicasting. This section describes one possible approach to providing such protection.

The Authentication Header might be used, with appropriate key management, to provide authentication of packets. This authentication is potentially important in packet classification within routers. The IPv6 Flow Identifier might act as a Low-Level Identifier (LLID). Used together, packet classification within routers becomes straightforward if the router is provided with the appropriate keying material. For performance reasons the routers might authenticate only every Nth packet rather than every packet, but this is still a significant improvement over capabilities in the current Internet. Quality of service provisioning is likely to also use the Flow ID in conjunction with a resource reservation protocol, such as RSVP [ZDESZ93]. Thus, the authenticated packet classification can be used to help ensure that each packet receives appropriate handling inside routers.

5.4 USE IN COMPARTMENTED OR MULTI-LEVEL NETWORKS

A multi-level secure (MLS) network is one where a single network is used to communicate data at different sensitivity levels (e.g., Unclassified and Secret) [DoD85] [DoD87]. Many governments have

significant interest in MLS networking [DIA]. The IP security mechanisms have been designed to support MLS networking. MLS networking requires the use of strong Mandatory Access Controls (MAC), which ordinary users are incapable of controlling or violating [BL73]. This section pertains only to the use of these IP security mechanisms in MLS environments.

The Authentication Header can be used to provide strong authentication among hosts in a single-level network. The Authentication Header can also be used to provide strong assurance for both mandatory access control decisions in multi-level networks and discretionary access control decisions in all kinds of networks. If explicit IP sensitivity labels (e.g., IPSO) [Ken91] are used and confidentiality is not considered necessary within the particular operational environment, the Authentication Header is used to provide authentication for the entire packet, including cryptographic binding of the sensitivity level to the IP header and user data. This is a significant improvement over labeled IPv4 networks where the label is trusted even though it is not trustworthy because there is no authentication or cryptographic binding of the label to the IP header and user data. IPv6 will normally use implicit sensitivity labels that are part of the Security Association but not transmitted with each packet instead of using explicit sensitivity labels. All explicit IP sensitivity labels MUST be authenticated using either ESP, AH, or both.

The Encapsulating Security Payload can be combined with appropriate key policies to provide full multi-level secure networking. In this case each key must be used only at a single sensitivity level and compartment. For example, Key "A" might be used only for sensitive Unclassified packets, while Key "B" is used only for Secret/No-compartments traffic, and Key "C" is used only for Secret/No-Foreign traffic. The sensitivity level of the protected traffic MUST NOT dominate the sensitivity level of the Security Association used with that traffic. The sensitivity level of the Security Association MUST NOT dominate the sensitivity level of the key that belongs to that Security Association. The sensitivity level of the key SHOULD be the same as the sensitivity level of the Security Association. The Authentication Header can also have different keys for the same reasons, with the choice of key depending in part on the sensitivity level of the packet.

Encryption is very useful and desirable even when all of the hosts are within a protected environment. The Internet-standard encryption algorithm could be used, in conjunction with appropriate key management, to provide strong Discretionary Access Controls (DAC) in conjunction with either implicit sensitivity labels or explicit sensitivity labels (such as IPSO provides for IPv4 [Ken91]). Some

environments might consider the Internet-standard encryption algorithm sufficiently strong to provide Mandatory Access Controls (MAC). Full encryption SHOULD be used for all communications between multi-level computers or compartmented mode workstations even when the computing environment is considered to be protected.

6. SECURITY CONSIDERATIONS

This entire memo discusses the Security Architecture for the Internet Protocol. It is not a general security architecture for the Internet, but is instead focused on the IP layer.

Cryptographic transforms for ESP which use a block-chaining algorithm and lack a strong integrity mechanism are vulnerable to a cut-and-paste attack described by Bellovin and should not be used unless the Authentication Header is always present with packets using that ESP transform [Bel95].

If more than one sender uses shares a Security Association with a destination, then the receiving system can only authenticate that the packet was sent from one of those systems and cannot authenticate which of those systems sent it. Similarly, if the receiving system does not check that the Security Association used for a packet is valid for the claimed Source Address of the packet, then the receiving system cannot authenticate whether the packet's claimed Source Address is valid. For example, if senders "A" and "B" each have their own unique Security Association with destination "D" and "B" uses its valid Security Association with D but forges a Source Address of "A", then "D" will be fooled into believing the packet came from "A" unless "D" verifies that the claimed Source Address is party to the Security Association that was used.

Users need to understand that the quality of the security provided by the mechanisms provided by these two IP security mechanisms depends completely on the strength of the implemented cryptographic algorithms, the strength of the key being used, the correct implementation of the cryptographic algorithms, the security of the key management protocol, and the correct implementation of IP and the several security mechanisms in all of the participating systems. The security of the implementation is in part related to the security of the operating system which embodies the security implementations. For example, if the operating system does not keep the private cryptologic keys (that is, all symmetric keys and the private asymmetric keys) confidential, then traffic using those keys will not be secure. If any of these is incorrect or insufficiently secure, little or no real security will be provided to the user. Because different users on the same system might not trust each other, each user or each session should usually be keyed separately. This will

also tend to increase the work required to cryptanalyse the traffic since not all traffic will use the same key.

Certain security properties (e.g., traffic analysis protection) are not provided by any of these mechanisms. One possible approach to traffic analysis protection is appropriate use of link encryption [VK83]. Users must carefully consider which security properties they require and take active steps to ensure that their needs are met by these or other mechanisms.

Certain applications (e.g., electronic mail) probably need to have application-specific security mechanisms. Application-specific security mechanisms are out of the scope of this document. Users interested in electronic mail security should consult the RFCs describing the Internet's Privacy-Enhanced Mail system. Users concerned about other application-specific mechanisms should consult the online RFCs to see if suitable Internet Standard mechanisms exist.

ACKNOWLEDGEMENTS

Many of the concepts here are derived from or were influenced by the US Government's SDNS security protocol specifications, the ISO/IEC's NLSP specification, or from the proposed swIPe security protocol [SDNS, ISO, IB93, IBK93]. The work done for SNMP Security and SNMPv2 Security influenced the choice of default cryptological algorithms and modes [GM93]. Steve Bellovin, Steve Deering, Richard Hale, George Kamis, Phil Karn, Frank Kastenholz, Perry Metzger, Dave Mihelcic, Hilarie Orman and Bill Simpson provided careful critiques of early versions of this document.

REFERENCES

- [Atk95a] Atkinson, R., "IP Authentication Header", RFC 1826, NRL, August 1995.
- [Atk95b] Atkinson, R., "IP Encapsulating Security Payload", RFC 1827, NRL, August 1995.
- [BCCH94] Braden, R., Clark, D., Crocker, S., and C. Huitema, "Report of IAB Workshop on Security in the Internet Architecture", RFC 1636, USC/Information Sciences Institute, MIT, Trusted Information Systems, INRIA, June 1994.
- [Bel89] Steven M. Bellovin, "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, Vol. 19, No. 2, March 1989.

- [Bel95] Steven M. Bellovin, Presentation at IP Security Working Group Meeting, Proceedings of the 32nd Internet Engineering Task Force, March 1995, Internet Engineering Task Force, Danvers, MA.
- [BFC93] A. Ballardie, P. Francis, & J. Crocroft, "Core Based Trees: An Architecture for Scalable Inter-Domain Multicast Routing", Proceedings of ACM SIGCOMM 93, ACM Computer Communications Review, Volume. 23, Number 4, October 1993, pp. 85-95.
- [BL73] Bell, D.E. & LaPadula, L.J., "Secure Computer Systems: Mathematical Foundations and Model", Technical Report M74-244, The MITRE Corporation, Bedford, MA, May 1973.
- [CB94] William R. Cheswick & Steven M. Bellovin, Firewalls & Internet Security, Addison-Wesley, Reading, MA, 1994.
- [DIA] US Defense Intelligence Agency, "Compartmented Mode Workstation Specification", Technical Report DDS-2600-6243-87.
- [DoD85] US National Computer Security Center, "Department of Defense Trusted Computer System Evaluation Criteria", DoD 5200.28-STD, US Department of Defense, Ft. Meade, MD., December 1985.
- [DoD87] US National Computer Security Center, "Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria", NCSC-TG-005, Version 1, US Department of Defense, Ft. Meade, MD., 31 July 1987.
- [DH76] W. Diffie & M. Hellman, "New Directions in Cryptography", IEEE Transactions on Information Theory, Vol. IT-22, No. 6, November 1976, pp. 644-654.
- [EK94] D. Eastlake III & C. Kaufman, "Domain Name System Protocol Security Extensions", Work in Progress.
- [GM93] Galvin J., and K. McCloghrie, "Security Protocols for version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1446, Trusted Information Systems, Hughes LAN Systems, April 1993.
- [HA94] Haller, N., and R. Atkinson, "On Internet Authentication", RFC 1704, Bell Communications Research, NRL, October 1994.
- [Hin94] Bob Hinden (Editor), Internet Protocol version 6 (IPv6) Specification, Work in Progress, October 1994.

- [ISO] ISO/IEC JTC1/SC6, Network Layer Security Protocol, ISO-IEC DIS 11577, International Standards Organisation, Geneva, Switzerland, 29 November 1992.
- [IB93] John Ioannidis and Matt Blaze, "Architecture and Implementation of Network-layer Security Under Unix", Proceedings of USENIX Security Symposium, Santa Clara, CA, October 1993.
- [IBK93] John Ioannidis, Matt Blaze, & Phil Karn, "swIpe: Network-Layer Security for IP", presentation at the Spring 1993 IETF Meeting, Columbus, Ohio.
- [Ken91] Kent, S., "US DoD Security Options for the Internet Protocol", RFC 1108, BBN Communications, November 1991.
- [Ken93] Kent, S., "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", RFC 1422, BBN Communications, February 1993.
- [KB93] Kohl, J., and B. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, Digital Equipment Corporation, USC/Information Sciences Institute, September 1993.
- [MS95] Metzger, P., and W. Simpson, "IP Authentication with Keyed MD5", RFC 1828, Piermont, Daydreamer, August 1995.
- [KMS95] Karn, P., Metzger, P., and W. Simpson, "The ESP DES-CBC Transform", RFC 1829, Qualcomm, Inc., Piermont, Daydreamer, August 1995.
- [NS78] R.M. Needham & M.D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers", Communications of the ACM, Vol. 21, No. 12, December 1978, pp. 993-999.
- [NS81] R.M. Needham & M.D. Schroeder, "Authentication Revisited", ACM Operating Systems Review, Vol. 21, No. 1., 1981.
- [OTA94] US Congress, Office of Technology Assessment, "Information Security & Privacy in Network Environments", OTA-TCT-606, Government Printing Office, Washington, DC, September 1994.
- [Sch94] Bruce Schneier, Applied Cryptography, Section 8.6, John Wiley & Sons, New York, NY, 1994.

- [SDNS] SDNS Secure Data Network System, Security Protocol 3, SP3, Document SDN.301, Revision 1.5, 15 May 1989, published in NIST Publication NIST-IR-90-4250, February 1990.
- [VK83] V.L. Voydock & S.T. Kent, "Security Mechanisms in High-level Networks", ACM Computing Surveys, Vol. 15, No. 2, June 1983.
- [ZDESZ93] Zhang, L., Deering, S., Estrin, D., Shenker, S., and D. Zappala, "RSVP: A New Resource ReSerVation Protocol", IEEE Network magazine, September 1993.

DISCLAIMER

The views expressed in this note are those of the author and are not necessarily those of his employer. The Naval Research Laboratory has not passed judgement on the merits, if any, of this work. The author and his employer specifically disclaim responsibility for any problems arising from correct or incorrect implementation or use of this design.

AUTHOR'S ADDRESS

Randall Atkinson
Information Technology Division
Naval Research Laboratory
Washington, DC 20375-5320
USA

Phone: (202) 767-2389
Fax: (202) 404-8590
EMail: atkinson@itd.nrl.navy.mil

