Designation: E 2086 – 00

An American National Standard

Standard Guide for Internet and Intranet Healthcare Security[1]

This standard is issued under the fixed designation E 2086; the number immediately following the designation indicates the year of original adoption or, in the case of revision, the year of last revision. A number in parentheses indicates the year of last reapproval. A superscript epsilon (ε) indicates an editorial change since the last revision or reapproval.

1. Scope

1.1 This guide covers mechanisms that can be used to protect healthcare information which is being transmitted over networks using the Internet Protocol Suite (IPS). This includes the actual Internet itself, as well as corporate intranets constructed from off–the–shelf components implementing these protocols. An organization's security policy will determine when these mechanisms are used, based on risk analysis.

1.2 The Internet Engineering Task Force (IETF) is defining security standards for use with the IPS. This guide covers the relevant standards and recommends, where needed, particular options (such as cryptographic transformations) to be used with the standards. Most standards referenced here are proposed standards issued as Requests for Comments (RFCs). Some are in the draft stage, but are stable enough (and widely enough implemented) to be recommended for use at this time.

2. Referenced Documents

2.1 *IETF Standards:*[2]

RFC 1510 Kerberos Authentication Service
RFC 1777 Lightweight Directory Access Protocol (v2)
RFC 2251 Lightweight Directory Access Protocol (v3)
RFCs 1901–1910 Simple Network Management Protocol
RFC 1945 Hypertext Transfer Protocol
RFC 1964 Kerberos v5 GSS-API Mechanism
RFC 2246 The TLS Protocol Version 1.0
RFC 2401 Security Architecture for the Internet Protocol
RFC 2402 IP Authentication Header
RFC 2403 The Use of HMAC-MD5–96 within ESP and AH
RFC 2404 The Use of HMAC-SHA-196 within ESP and AH
RFC 2406 IP Encapsulating Security Payload (ESP)
RFC 2407 The Internet IP Security Domain of Interpretation for ISAKMP
RFC 2408 Internet Security Association and Key Management Protocol (ISAKMP)
RFC 2409 The Internet Key Exchange (IKE)
RFC 2411 IP Security Document Roadmap
RFC 2440 OpenPGP Message Format
RFC 2451 The ESP CBC-Mode Cipher Algorithms
RFC 2560 Internet X.509 Public Key Infrastructure Online Certificate Status Protocol
RFC 2630 Cryptographic Message Syntax
RFC 2631 Diffie-Hellman Key Agreement Method
RFC 2632 S/MIME Version 3 Certificate Handling
RFC 2633 S/MIME Version 3 Message Specification
RFC 2634 Enhanced Security Services for S/MIME

2.2 *Other Standards:*

FIPS PUB 180–1 Secure Hash Algorithm

3. Terminology

3.1 *Definitions:*

3.1.1 *algorithm*—a clearly specified mathematical process for computation; a set of rules which, if followed, will give a prescribed result.

3.1.2 *asymmetric cryptography*—cryptographic algorithm that uses two related keys, a public key and a private key; the two algorithm keys have the property that, given the public key, it is computationally infeasible to derive the private key.

3.1.3 *authentication*—the corroboration that the source of data received is as claimed.

3.1.4 *authorization*—the granting of rights.

3.1.5 *cipher text*—data in its enciphered form.

3.1.6 *clear text*—data in its original, unencrypted form.

3.1.7 *confidentiality*—the property that information is not made available to or disclosed to unauthorized individuals, entities, and processes.

3.1.8 *cryptographic checkvalue*—a value computed using a shared secret key and a data unit, which can be used to provide data integrity and authentication services.

3.1.9 *cryptography*—the discipline which embodies principles, means, and methods for the transformation of data in order to hide its information content, prevent its undetected modification, prevent its unauthorized use or a combination thereof.

3.1.10 *datagram*—a data unit that is delivered independently of other data units transmitted over a network.

3.1.11 *data integrity*—a property whereby data has not been altered or destroyed.

[1] This guide is under the jurisdiction of ASTM Committee E31 on Healthcare Informatics, and is the direct responsibility of Subcommittee E31.20 on Data and System Security for Health Information.

Current edition approved April 10, 2000. Published June 2000.

[2] Available on line at ftp://ds.internic.net.

3.1.12 *decryption*—a process of transforming ciphertext into plaintext.

3.1.13 *digital signature*—a cryptographic transformation of data which, when associated with a data unit, provides the services of origin authentication, data integrity, and signer non–repudiation.

3.1.14 *encryption*—a process of transforming plain text (readable) into cipher text (unreadable) for the purpose of security or privacy.

3.1.15 *encryption key*—a binary number used to transform plain text into ciphertext.

3.1.16 *gateway*—a computer system or other device that acts as a translator between two systems that do not use the same communications protocols, data formatting, structures, languages, or architecture, or a combination thereof.

3.1.17 *intranet*—an internal corporate network which uses the Internet protocol suite (TCP, IP, etc.)

3.1.18 *non–repudiation*—this service provides proof of the integrity and origin of data, both in an unforgeable relationship, which can be verified by any party.

3.1.19 *plain text*—data in its original, unencrypted form.

3.1.20 *repudiation*—the denial by a user of having participated in part or all of a communication. (See *non–repudiation*, which has the opposite meaning.)

3.1.21 *replay*—the process of sending a previously sent message as a method of perpetrating a fraud.

3.1.22 *security association*—the relationship between two entities which allows the protection of information communicated between the entities.

3.1.22.1 *Discussion*—This relationship includes a shared symmetric key, and security attributes describing the relationship. The security association is used to negotiate the characteristics of these protection mechanisms, but does not include the protection mechanisms themselves.

3.1.23 *session*—logical relationship between two network endpoints that supports a user or network application.

3.1.24 *subnetwork*—a network segment, usually with its own address.

3.1.25 *symmetric encryption*—encryption using a single key to encrypt and decrypt which both the sender and receiver hold privately.

3.1.26 *virtual private network*—a network which uses the Internet as a carrier, but is operated as a dedicated point-to-point network.

3.1.26.1 *Discussion*—Encryption is used to segregate and protect the VPN's data when it is conveyed over the Internet.

3.2 *Acronyms:Acronyms:*

3.2.1 *AH*—Authentication Header

3.2.2 *API*—Application Programming Interface

3.2.3 *ASTM*—American Society for Testing and Materials

3.2.4 *ATM*—Asynchronous Transfer Mode

3.2.5 *DEC*—Digital Equipment Corporation

3.2.6 *DES*—Data Encryption Standard

3.2.7 *DSA*—Digital Signature Algorithm

3.2.8 *EDI*—Electronic Data Interchange

3.2.9 *ESP*—Encapsulating Security Payload

3.2.10 *FTP*—File Transfer Protocol

3.2.11 *GSS*—Generic Security Services

3.2.12 *HMAC*—Hashed Message Authentication Code

3.2.13 *HTTP*—HyperText Transfer Protocol

3.2.14 *IDUP*—Independent Data Unit Protection

3.2.15 *IETF*—Internet Engineering Task Force

3.2.16 *IP*—Internet Protocol

3.2.17 *IPS*—Internet Protocol Suite

3.2.18 *IPSEC*—Internet Protocol Security

3.2.19 *ISAKMP*—Internet Security Association and Key Management Protocol

3.2.20 *LAN*—Local Area Network

3.2.21 *MD*—Message Digest

3.2.22 *MIME*—Multipurpose Internet Mail Extension

3.2.23 *PCT*—Private Communications Technology

3.2.24 *PIN*—Personal Identification Number

3.2.25 *PKCS*—Public–Key Cryptography Standards

3.2.26 *RFC*—Requests for Comment

3.2.27 *RSA*—Rivest, Shamir, and Adelman

3.2.28 *SHA-1*—Secure Hash Algorithm

3.2.29 *S–HTTP*—Secure HyperText Transfer Protocol

3.2.30 *S/MIME*—Secure/Multipurpose Internet Mail Extensions

3.2.31 *SMTP*—Simple Mail Transfer Protocol

3.2.32 *SSL*—Secure Socket Layer

3.2.33 *TCP*—Transmission Control Protocol

3.2.34 *TLSP*—Transport Layer Security Protocol

3.2.35 *UDP*—User Datagram Protocol

3.2.36 *VPN*—Virtual Private Network

3.2.37 *WAN*—Wide Area Network

3.2.38 *WWW*—World Wide Web

4. Significance and Use

4.1 This guide recommends security mechanisms for protection of healthcare information transmitted using the IPS. The IPS consists of multiple protocol layers.

4.2 The lowest layer which can provide end–to–end security is the Internet Protocol (IP). IP may run over a variety of subnetwork technologies, such as Ethernet, X.25, ATM, and even asynchronous dial–up lines. While it is possible to provide security services directly over those technologies, such approaches only protect a single subnetwork and are not discussed further.

4.3 A variety of protocols may be run on top of IP. These include the Transmission Control Protocol (TCP), which provides reliable, sequenced data delivery (sessions), and the User Datagram Protocol (UDP), which provides unsequenced data delivery (datagrams). Other protocols at this layer include various routing and configuration protocols used by the network itself.

4.4 Application protocols typically make use of either TCP or UDP. A variety of standard application protocols have been defined for such applications as file transfer (FTP), electronic mail (SMTP), and the World Wide Web (HTTP). Some applications have their own security requirements, dictated by the structure of the application or its protocols.

4.5 The remainder of this guide is organized as follows: Section 5 discusses security threats and the countermeasures which can be used to protect against these threats. Section 6 presents a brief overview of cryptography, as most network security mechanisms rely on its use. Section 7 distinguishes

between network and application security and discusses when each level of security might be useful. The remaining sections recommend specific security protocols and mechanisms for both network and application security needs.

5. Threats and Countermeasures

5.1 This section covers the principal threats to a system. In some cases, security services can prevent an attack; in other cases, they merely detect an attack.

5.1.1 *Masquerade* occurs when an entity successfully pretends to be another entity. This includes impersonation of users or system components, as well as falsely claiming origination or acknowledging receipt of a message or transaction.

5.1.2 *Modification of information* can include modification of message or data content, as well as destruction of messages, data, or management information. This includes message sequencing threats, which occur when the order of messages is altered.

5.1.3 *Unauthorized disclosure* threats include revealing message contents or other data, as well as information derived from observing traffic flow, as well as revealing information held in storage on an open system.

5.1.4 *Repudiation* occurs when a user or the system denies having performed some action, such as origination or reception of a message.

5.1.5 *Denial of service* threats prevent the system from performing its functions. This may be accomplished by attacks on the underlying communications infrastructure, attacks on the underlying applications, or by flooding the system with extra traffic.

5.2 The following services protect against the threats described in 5.1.1-5.1.5.

5.2.1 *Peer entity authentication* provides proof of the identity of communicating parties. Various types of authentication exchanges have been discussed in the literature; most are based on digital signatures or other cryptographic mechanisms.

5.2.2 *Data origin authentication* counters the threat of masquerade and is provided using digital signatures or other cryptographic integrity mechanisms.

5.2.3 *Access control* counters the threat of unauthorized disclosure or modification of data. This is particularly appropriate on an end system. A variety of access control strategies can be found in the standards, including access control lists and security labels. Since access control is typically provided on an end system, it is not discussed further in this guide.

5.2.4 *Confidentiality* counters the threat of unauthorized disclosure, particularly during the transfer of information. Confidentiality can be applied to entire messages or other data units or to selected fields. Encryption is used to provide this service.

5.2.5 *Integrity* counters the threat of unauthorized modification of data. This can be provided with various types of integrity check values. To protect against deliberate modification, a cryptographic check value or digital signature should be used. This also provides the service of data origin authentication. As with confidentiality, this service may be applied to entire messages or selected fields. One particularly useful application of selective field integrity is message sequence integrity, in which the integrity service is applied to a sequence number or other sequencing information.

5.2.6 *Non–repudiation* of origin and delivery protect against an originator or recipient falsely denying originating or receiving a message. This service provides proof (to a third party) of origin or receipt, and is provided using digital signatures.

6. Cryptography Overview

6.1 Cryptography is the art or science of keeping data secure from disclosure, modification, and forgery. It is particularly appropriate in today's computing environment, given the increasing use of networks to connect systems (implying more, possibly unknown users may access data), the increasing amount of sensitive data being conveyed on these networks, legal requirements for protection of data, and the ease and low cost of network attack.

6.2 *Encryption* can be used to provide confidentiality and integrity services. Following are two types of encryption systems:

6.2.1 In *symmetric* (conventional) cryptography, the sender and recipient share a secret key. This key is used by the originator to encrypt a message and by the recipient to decrypt a message. The Data Encryption Standard (DES) is an example of a symmetric cryptosystem. Confidentiality is provided by encrypting the message under a shared key. Integrity and authentication are supported by computing a cryptographic checkvalue, or *authenticator*, over the message, using a key shared by the originator and recipient.

6.2.2 In *asymmetric* (public key) cryptography, different keys are used to encrypt and decrypt a message. Each user is associated with a pair of keys. One key (the *public key*) is publicly known and is used to encrypt messages destined for that user. The *private key* is known only to the user and is used to decrypt incoming messages. RSA (named after the inventors' initials) is the most well-known asymmetric algorithm.

6.3 Some asymmetric algorithms, such as RSA, can also provide authentication, integrity, and non-repudiation when used as follows:

6.3.1 To *sign* a data unit, the user encrypts it under his private key.

6.3.2 To *verify* the data unit, the recipient decrypts it with the originator's public key.

6.3.3 If the message is successfully decrypted, it must have been encrypted by the originator, who is the only entity that knows the corresponding private key.

6.3.4 A *digital signature* is, then, a piece of data appended to a message, generated from the message and the signer's private key, which allows the recipient to prove the origin of the message and to protect against modification and forgery.

6.4 Note the digital signature can be used to provide non–repudiation services. Unlike the authenticator discussed in 6.2, the private key used to sign a message is known only by the signer. This prevents the signer from claiming that another party (for example, the recipient) generated a given digital signature.

7. Network and Application Security

7.1 *Network Security*:

7.1.1 Network security services protect data in transit between systems. This would be appropriate in the following situations[3]:

7.1.1.1 The end system is trusted, but the underlying network is not trusted, or

7.1.1.2 Protection is required for all (or most) traffic between systems.

7.1.2 The security services are transparent to applications, which require no modification. Furthermore, performance of bulk data protection services is improved, since they can operate on larger data units and handle all applications the same way.

7.1.3 In some cases, it might be more cost effective to protect data crossing a given link or subnetwork. For example, a source system on a LAN can send sensitive data through a router onto the Internet; the data is sent to a destination router and onward to the destination system on another LAN. The LANs are typically as secure as the end–system (frequently the end–systems and LAN might share a security administrator. Conversely, the Internet is less trusted, so encryption between routers (on the Internet subnetwork) is appropriate. This solution is generally cheapest in terms of equipment, since there are many more end–systems than there are subnetwork gateways (for example, routers).

7.1.4 Some level of access control can be provided by firewalls, which can filter packets based on network addresses and target application. Additional protection is provided using security protocols. These protocols provide confidentiality using symmetric encryption, origin authentication and integrity using authenticators, and peer entity authentication using (typically) digital signatures. Management of encryption keys may be done using either secret–key or public–key techniques; public–key approaches scale better and are being adopted, for example, by the IETF. Cryptographic services are described in the IPSEC series of standards discussed in Section 8. Interoperability testing among about a dozen vendors has been ongoing.

7.2 *Application Security*:

7.2.1 Application security measures are built into a particular application, such as record storage and retrieval, imaging, or claims processing, and are independent of network layer security measures. Security shall be placed at this level if the following conditions exist:

7.2.1.1 Security services are application–specific, or

7.2.1.2 Services traverse multiple application programs when data is moved from source to destination.

7.2.2 An example of 7.2.1.1 would be secure insurance claims using Electronic Data Interchange (EDI). Standards are defined both for the claim format (a transaction set), and for securing individual transaction sets and functional groups within an EDI interchange. As another example, ATM applications which encrypt only the PIN portion of a financial transaction. The major example of 7.2.1.2 is store–and–forward electronic mail, in which sender and recipient(s) never directly communicate, and in which only the content portion of a message is protected.

7.2.3 Application layer security provides confidentiality using a combination of symmetric and asymmetric encryption, and authentication and integrity using authenticators or digital signatures. Non–repudiation may be provided in conjunction with authentication, using digital signatures.

7.3 *Placement of Security Services*:

7.3.1 Network and application security are useful in different circumstances. Following are some criteria for choosing the layer(s) in which to place security.

7.3.1.1 Application traffic is typically multiplexed onto network layer connections. Therefore, it is likely that security services at the network layer will be protecting a data stream containing traffic to and from different sources and destinations. If the security policy dictates that all (or most) traffic requires a certain degree of protection, use of lower level security is desirable for efficiency reasons. If security is at the discretion of individual users, lower level services may not be desirable due to the cost of unnecessarily protecting data which does not require protection. In such a case, application level security is a better choice.

7.3.1.2 At the network layer, there is more knowledge of the security characteristics of particular routes and links. If these characteristics vary greatly within different portions of the network, using network layer security is preferable, since appropriate security services can be selected on a per–subnetwork or per–link basis rather than being implemented in all end–systems.

7.3.1.3 As mentioned in 7.3.1.2, the minimum number of protection points is at the subnetwork layer. This level of security might be the most cost-effective, compared to direct link level security. Placing services at the direct link layer requires security devices at the end of every link. Placing services at higher layers requires their implementation in every end–system or sensitive application. Since much of this would be done in (relatively inexpensive) software rather than in hardware, a cost analysis would be needed to determine which approach is cheapest. One particularly useful option is to encrypt traffic traversing a WAN used to interconnect multiple corporate LANs. Since LANs are typically confined to a single (somewhat secure) facility, there may not be a requirement for encryption within a single LAN. However, traffic between LANs (using and untrusted WAN such as the Internet) would require encryption. Such encryption would be implemented in (or immediately outboard of) the routers which interconnect the LANs to the WAN, and would encrypt traffic destined for the other corporate LANs. This type of configuration is known as a “virtual private network.”

7.3.1.4 Those services which associate data with an originator or recipient (for example, authentication and non–repudiation) are best provided at the application layer. This provides the greatest granularity (typically to the individual user). When provided at lower levels, trusted hardware or software is needed to bind the originator to the originating end system.

7.4 *System Security*:

[3] Ford, Warwick, *Computer Communications Security: Principles, Standard Protocols and Techniques*, Prentice Hall, 1994.

7.4.1 Cryptographic protection across the network is useless without proper security measures on the end system. Such measures include the following:

7.4.1.1 *Access control* protects data from unauthorized disclosure or modification. Some operating systems already provide access control features. For other platforms, a variety of add–on products are available. Even stronger protection can be provided by encrypting data stored on the local system or even the fileserver when there is a client server environment.

7.4.1.2 Access control is predicated on proper authentication of the user. A variety of token–based authentication products are available to improve on operating system authentication mechanisms. In some environments, it is necessary to forward authentication information (or evidence of local authentication) to other systems. A number of protocols have been designed to do this, including Kerberos (RFC 1510 and RFC 1964).

7.4.1.3 Proper configuration management will ensure that all software security updates are performed, and that only required applications are run on end–systems.

7.4.1.4 Robust audit procedures will minimize the impact of security breaches, by ensuring prompt detection of successful penetrations and quick implementation of preventive measures.

8. IP Security Recommendations

8.1 IP security mechanisms are specified in RFC 2401, 2402, 2406, and 2411 , which describe the security protocol architecture, authentication header (AH), and encapsulating security payload (ESP), respectively. The AH provides authentication and integrity, while the ESP provides confidentiality and, optionally, integrity and authentication. The AH and ESP each consist of a protocol header. Additionally, the ESP includes the actual data being sent, after being subjected to some cryptographic transform (for example, encryption). Additional RFCs describe particular authentication mechanisms and cryptographic transforms.

8.2 IP security mechanisms may be applied on a subnetwork basis, using security gateways. For example, all users on a LAN might be trusted, so no security is needed. When sending data over a WAN (for example, the Internet), a security gateway could add the AH and ESP to provide the necessary security services. Similarly, a recipient security gateway might process the AH and ESP, delivering unprotected IP packets to a recipient on the destination LAN. This method of using the untrusted Internet to bridge two portions of a corporate intranet has been called a “virtual private network.”

8.3 Security headers identify a particular security association, which defines the security relationship between source and destination. Attributes of a security association include, for example, identifiers of the authentication and encryption algorithms being used, the cryptographic keys being used, lifetime of the association, and security labels.

8.4 The following AH (RFC 2402) mechanisms are recommended, where the AH is used. ESP may provide all required services, as well.

8.4.1 HMAC using SHA–1 (RFC 2404).

8.4.2 HMAC using MD5 (RFC 2403).

8.5 The ESP mechanism can encapsulate an entire IP datagram (tunnel mode) or just the upper-layer (for example, TCP) protocol data included in the datagram (transport mode). In the first case, even the actual IP addresses in the inner datagram can be replaced with other addresses in the outer datagram, providing some protection against traffic analysis. This guide does not recommend one mode over the other. The following ESP transforms are recommended:

8.5.1 *Triple–DES–CBC, with explicit IVs*—This is appropriate where information must remain secret for long periods of time (months or years) (see RFC 2451).

8.5.2 Other encryption algorithms are available, but have not received sufficient evaluation to be recommended at this time. NIST is currently in the process of choosing an Advanced Encryption Standard (AES) to replace DES.

8.6 A protocol for key management is also needed. This guide recommends the ISAKMP key management protocol (RFC 2407 and RFC 2408), with the Internet Key Exchange mechanism (RFC 2409). IKE is based on the Diffie–Hellman key agreement algorithm. The IKE aggressive mode and quick mode exchanges are recommended, using the standard Diffie–Hellman group parameters defined in the standard. The standards discuss when each mode would be used.

9. Application Security Recommendations

9.1 As described in 7.2.2, electronic mail applications require protection of the message content but not the envelope (routing information). Additionally, each message shall be protected independently. There have been several attempts to establish a standard for a secure mail structure. This guide recommends the use of S/MIME (RFC 2630, 2631, 2632, 2633, 2634), since it has been adopted by most of the large E–mail vendors. S/MIME is an encapsulation of a PKCS #7 message as a MIME content–type. PKCS #7 is being used in a variety of other standards, such as the SET proposal for credit card transactions, and S–HTTP (see 9.3). S/MIME provides authentication, integrity, and non–repudiation using digital signatures, and confidentiality using encryption. A fresh encryption key is used for each message, and public key encryption is used to send this message key to each recipient. Recommended algorithms include triple–DES–CBC for encryption, RSA or DSA for signature, and RSA for key management. (Work is also under way to standardize the use of Diffie-Hellman for key management in this scenario.) For closed systems with a relatively small number of users, PGP (RFC 2440) might be used as an alternative to S/MIME.

9.2 Secure client–server interaction can be done using the IETF’s Transport Layer Security (TLS) protocol (RFC 2246). This protocol is based on the Secure Socket Layer (SSL) protocol developed by Netscape, along with features from a Microsoft variant called Private Communications Technology (PCT). TLS runs at a layer between the application and TCP, and can be used with any session–oriented application. Currently, it is widely implemented for WWW (protecting HTTP interactions), and work is underway to integrate it into other application protocols as well. TLS provides the following security services:

9.2.1 Server authentication at session startup,

9.2.2 Client authentication at session startup,

9.2.3 Confidentiality of all data on the session, and

9.2.4 Integrity and origin authentication of all data on the session.

9.2.5 Recommended algorithms include DES–CBC for encryption, RSA or DSA for signature (client authentication), and RSA or Diffie–Hellman for key management.

9.3 In some cases, it may be desirable to provide non–repudiation of particular documents being accessed via HTTP. TLS does not provide the capability to sign individual messages. Therefore, the use of S–HTTP is recommended if this is a requirement. S–HTTP is basically an encapsulation of HTTP messages in a secure E–mail format, and it therefore provides all of the services described in 9.1. The use of CMS (RFC 2560) as the encapsulation format is recommended; for encryption, it is also possible to use pre–arranged symmetric keys, so it can be used without client certificates. The S-HTTP protocol is currently a work in process in the IETF. The current Internet draft documents have expired.

10. Keywords

10.1 internet; internet security; intranet