SDNS ARCHITECTURE AND END-TO-END ENCRYPTION

Ruth Nelson and John Heimann
Electronic Defense Communications Division
GTE Government Systems Corporation
100 First Avenue
Waltham, MA 02254

INTRODUCTION

The Secure Data Network System (SDNS) is intended to provide secure data communications to a variety of DoD and commercial users. SDNS services include key management and system management as well as data encryption, authentication and access control. The program is a U. S. Government/Industry effort, with participation by the National Security Agency, National Institute for Standards and Technology, other government agencies and about a dozen government contractors. During the concept definition and prototyping phases, a joint working group defined the set of security services to be provided and developed protocols for key management and for secure communications [1]. The protocols and architecture are compatible with the International Standards Organization (ISO) Reference Model for Open Systems Interconnection (OSI), and the end-to-end encryption (E3) protocols are being proposed as U.S. and international standards. The E3 protocols are publicly released and appropriate for the OSI environment.

SDNS includes security protocols at the Application, Transport, Network and Link layers of the OSI protocol hierarchy [2], as shown in Figure 1. The security services provided at each layer are a subset of those in the Security Addendum to the OSI Reference Model [3]. MSP, an application layer protocol, provides electronic mail security compatible with the CCITT X.400 series recommendations. KMP, a single key management protocol, negotiates keys and security parameters for all of the lower layer encryption services. The transport and network layer E3 protocols, which are the main subject of this paper, provide protection for the user data but allow the forwarding and handling of the data units by the network packet switches. Link encryption protocols are being defined by the working group for a variety of communication links and local networks.

G. Brassard (Ed.): Advances in Cryptology - CRYPTO '89, LNCS 435, pp. 356-366, 1990.

END-TO-END ENCRYPTION

In packet networks and internets, the term *end-to-end encryption* has been used to refer to an encryption scheme that encrypts the user data but provides unencrypted network headers. This allows the data to be routed and delivered by the network communications switches. A number of schemes have been proposed at various layers of the protocol architecture. If the E3 is too high in the protocol hierarchy, then the scheme is sensitive to the higher level protocols and formats used by the communicating parties, and different security protocols may be required for different applications. If it is too low in the hierarchy, then the encryption must be undone at each switch so that the control information can be read. These constraints have meant that most of the E3 schemes are either at the transport or the network protocol layer. Both network layer and transport layer encryption are permitted by the OSI Security Addendum. Each location has its avid proponents, and each has certain advantages and limitations.

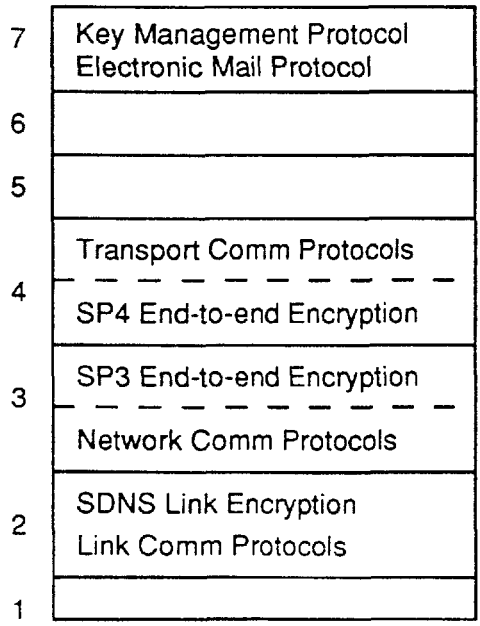

Figure 1. Placement of SDNS Protocols in the OSI Architecture

SECURITY PROTOCOL (SP)

SDNS has defined a security protocol (SP), which logically fits between the network and transport layer of the protocol hierarchy. Since OSI does not permit the addition of extra layers, we have defined this single protocol in two ways: as SP3 at the network layer and as SP4 at the transport layer [4,5]. The two representations share a common format and common basic functionality. Extensions to each protocol provide the extra services that are layer-specific. The common protocol subset is extremely simple: it encapsulates the user data (encryption and/or integrity protection) and identifies the encryption key to the receiver. SP3 and SP4 implementations can interoperate within their common domain. The protocols are independent of key management; they assume that the two parties have a common key and that each knows how to designate that key to the other. For SDNS, a key management protocol has been defined at the application layer, to provide a pairwise key for each secure association. If the key is shared by only two parties, then the protocol supports data origin authentication.

The SP service is negotiated by a separate mechanism before use, and the service is fixed for each key. The basic services are confidentiality (encryption) and connectionless integrity (each encapsulated unit protected from undetectable change). Either or both can be negotiated. The service is "connectionless;" i.e., it is provided on a per-protocol-data-unit basis and does not include sequencing or reliability. This is a major simplification in the security protocol. The needed reliability can be provided by the transport communications protocol functions operating above the security protocol, because the transport control information is encapsulated and protected by SP. SP relies on lower layer protocols for the communications functions (delivery, routing, bit-error detection on the links), and this makes the protocol simpler than some earlier designs.

SP does allow the use of multiple keys between two parties by providing a key identifier with the encrypted data. There is also an SP option (negotiated ahead of time) to carry a security label with each protected data unit. Another option allows the sender of the data to add octets of padding to the data unit.

SP3 and SP4 each have OSI-compliant definitions. SP3 is defined as a subnet-independent-convergence-protocol [6]. It is a separate protocol sublayer which sits on top of the network communications sublayer(s). The resulting network service is a secure version of the connectionless network service delivered by the OSI Connectionless Network Protocol (CLNP), which is functionally equivalent to the DoD Internet Protocol (IP). SP4 is defined as an addendum to OSI Transport Protocol (TP), which, together with TP4, can be used to provide a secure, connection-oriented service. SP3 and SP4 have identical formats. The common format is compatible with TP, so that SP4 can legitimately

be an addendum to TP. Figure 2 shows the SP format. The SP data unit includes a clear header, a protected header and protected data. The header and data are protected by an integrity check value (ICV) if integrity is provided, and they are encrypted if confidentiality is provided. The clear header identifies the data as secure (Type SE) and specifies the key for decrypting and/or integrity checking.

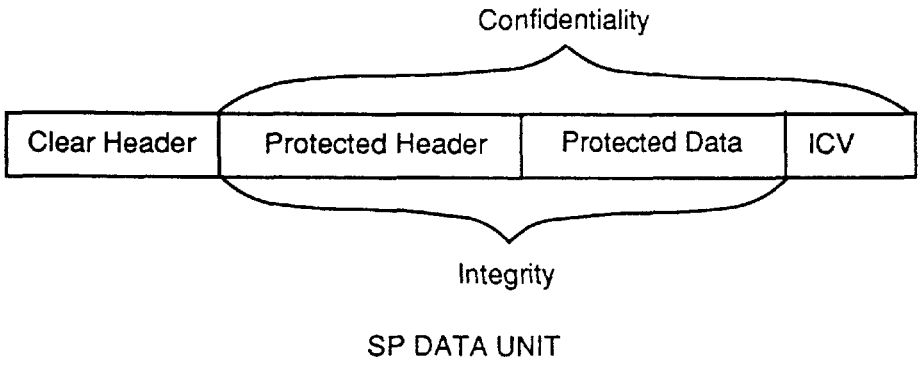

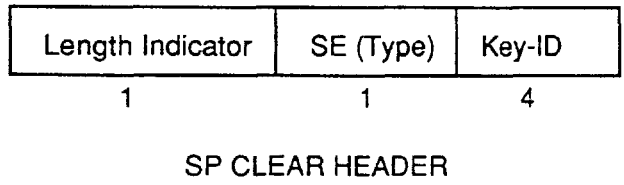

Figure 2. SP Header Formats

NETWORK LAYER ENCRYPTION (SP3)

OSI restricts transport layer protocols to end-systems, which are the computer systems that originate or terminate user data. At the network layer, however, the protocols can terminate at end-systems or intermediate systems, which are communications switching elements like packet switches or network gateways. It is often convenient to terminate "end-to-end" encryption at an intermediate system, and this requires that the encryption protocol be at the network layer or below. Terminating the encryption makes sense at the entrance to a local area network (LAN), for example, where all of the users are allowed access to all of the traffic or where another protection mechanism is used on the LAN. Terminating the encryption is also necessary if the communicating parties do not share a key and an

intermediate system must translate the protection to another security system. Figure 3 shows the interconnection of a protected LAN and two different E3 networks, one using SP3.

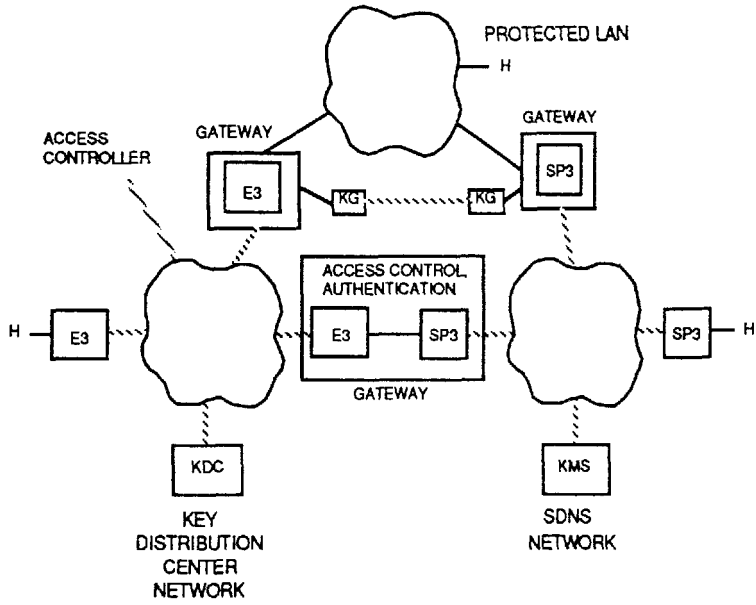

H - Host Computer
KMS - Key Management System
KDC - Key Distribution Center
KG - Key Generation, for link encryption
E^3 - Non-SDNS end-to-end encryption device
SP3 - SDNS SP3 devices

Figure 3. Interconnection of SDNS and Other Security Systems

Figure 4 shows the relationship between SP3, the transport protocol above it and the network protocol below it. The OSI model allows the network service to be provided by protocols that form sublayers of the network layer. An external key management service provides keys and security parameters, which are accessible to SP3 through a management information base. The service request to SP3 specifies a unit of data (Transport Protocol Data Unit) and some control information. The requestor can specify the source and destination addresses and a set of Quality-of-Service parameters. The Quality-of-Service parameters specify confidentiality, integrity, or both and can also request the inclusion of an explicit security label with the data.

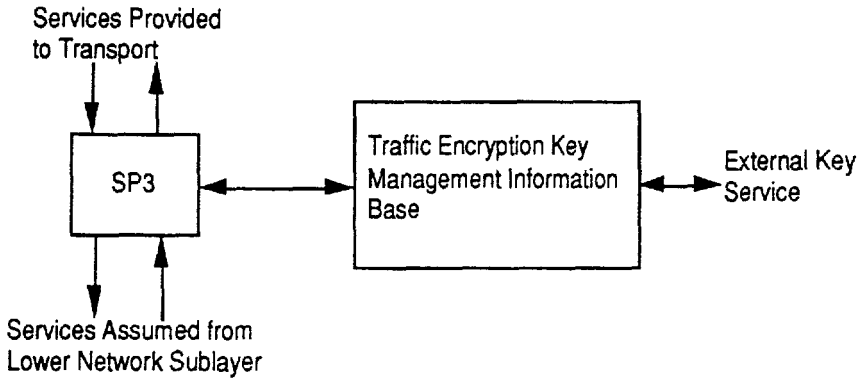

Figure 4. SP3 Interfaces

We have defined extensions to the basic SP protocol to handle the termination of SP3 at an intermediate point. The most basic of these extensions, called SP3A, simply encapsulates the end-system source and destination network addresses in the protected data unit. To completely protect these addresses, SP3A provides integrity protection, even if it was not specifically requested by the service user. The communications protocol below SP3 uses only the network address of the communicating SP3 entity and does not need to have access to the actual end-system address. At the encryption gateway, the encapsulated data is unsealed and the end-system destination address is used for routing and relaying the data unit to the end-system.

Figure 5 shows the protocol relationships in a gateway encryption scheme. In this figure, the E3 terminates at the gateway and is replaced by CLNP. The network protocols below SP3 and CLNP also are terminated at the gateway. Because the protocol on the unencrypted network is CLNP, the translation is very straightforward: the addressing schemes and data units are the same. CLNP can also be the communications protocol on NET B, operating below SP3, but this does not imply the existence of a single network protocol end-to-end. In a situation like this, end-to-end security is not provided by the E3 protocol, and other mechanisms must be used within the gateway and on the non-SP3 network.

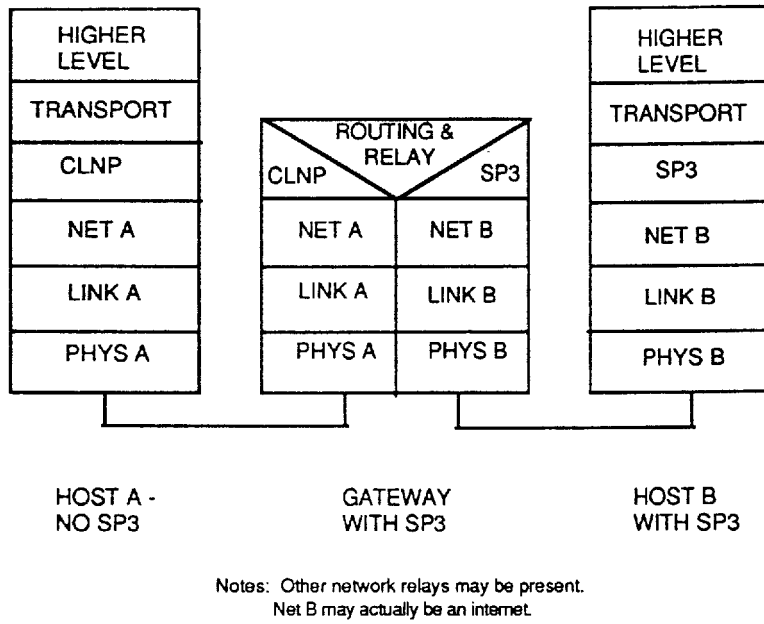

Figure 5. SP3 and CLNP at a Gateway

Additional extensions to SP3, SP3I for CLNP and SP3D for IP, permit the independent encryption and decryption of datagram "fragments," which are allowed by both CLNP and IP. Fragments are portions of a data unit which contain enough control information to be reassembled in order. They can be formed at any intermediate point in the communications path, normally when data unit size constraints require breaking up the data for passage through a particular network. They are normally reassembled only at the destination end-system, although reassembly can take place at an intermediate point if all the fragments arrive at that point. SP3I and SP3D encapsulate entire network data units or fragments, including both the address and the rest of the CLNP or IP header information. When the SP3 data unit is received at a gateway, the encapsulation is removed and normal CLNP or IP processing is done on the communications header. This mode of operation has some additional flexibility, because fragments do not have to be reassembled before encryption, but it is also more complex, because the communications processing takes place on the unencrypted data unit and so must be done securely. Two separate versions of the encryption protocol are needed because the formats of CLNP and IP headers are different. Both SP3I and SP3D provide integrity protection, even if not specifically requested, to protect the addressing and control information in the encapsulated CLNP or IP headers.

Figure 6 illustrates the SP3 addressing options. SP3N is the basic SP protocol and operates only between end-systems.

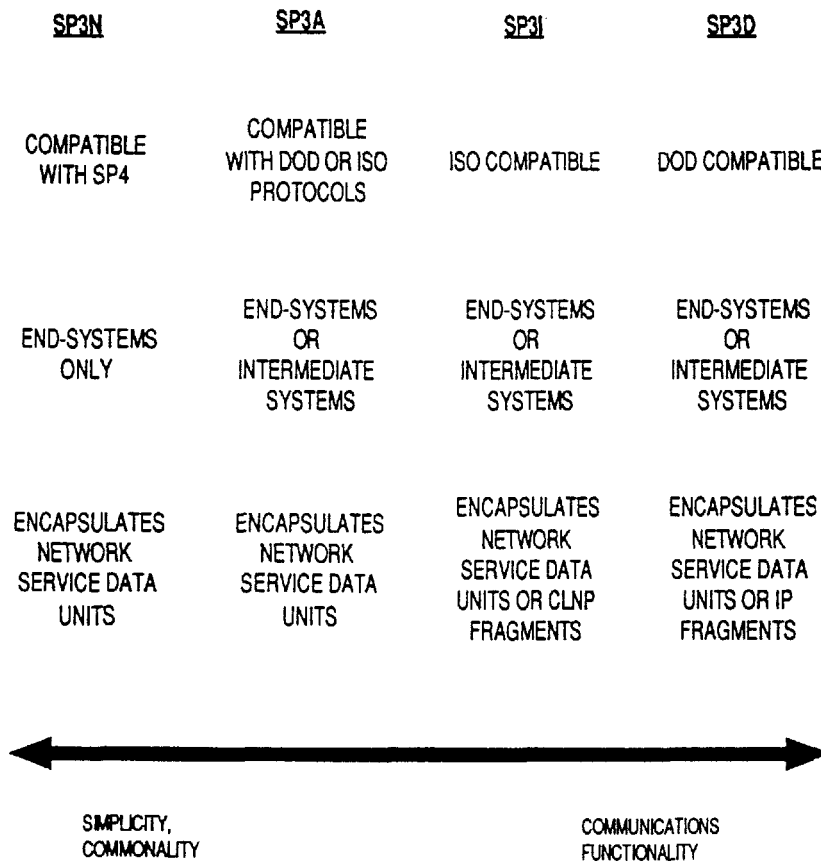

Figure 6. SP3 Addressing Options

TRANSPORT LAYER ENCRYPTION (SP4)

SP4 is defined as an addendum to the OSI transport protocols, TP. Because it is integrated with the transport protocol, it has access to all of the TP control information. The simplest form of SP4 acts identically to the simplest SP3. On output to the network, after all of the transport layer communications processing is finished, it encapsulates the Transport Protocol Data Unit (TPDU) by encrypting it, integrity-protecting it or both. This secure PDU is designated as a TPDU-type SE, reserved for SP4, and is then handed to the

network layer for processing. On receipt, the security processing is completed to expose a TPDU which is then processed normally. The extensions to SP4 allow provision of a connection-oriented secure service with truncation protection. This service includes use of a separate key for each transport connection, integrity protection on the data and the transport sequence numbers, and transmission of a Final Sequence Number at the close of a connection. SP3 with integrity protection can protect transport sequence numbers for normal TP processing, but it cannot provide the key-per-transport-connection granularity or final sequence numbers. Figure 7 shows the protected header formats used with each of the SP3 and SP4 options.

CONCLUSIONS

The SDNS end-to-end encryption protocols offer a simple, secure solution to the problem of sending sensitive data through a packet network or internet. The common subset, SP, provides a minimal solution suitable for end-system to end-system protection. The SP3 extensions provide added communications flexibility and the SP4 extensions provide added connection-oriented integrity.

ACKNOWLEDGEMENTS

The architecture and protocols described in this paper were developed by the SDNS Protocols and Signalling Working Group during the concept definition and prototyping phases of the program. The members of this group represented ten SDNS terminal contractors, GTE (the key management contractor), Analytics, the National Computer Security Center (NCSC) and various other government groups. Ruth Nelson had the privilege of chairing the group and John Heimann was an active member. Some of the concepts for the SDNS security architecture, and particularly for the end-to-end encryption protocols, were developed by GTE under an NCSC research program on Internet Security Architecture and Protocols, whose participants included GTE, Unisys, NCSC and the Defense Communications Agency.

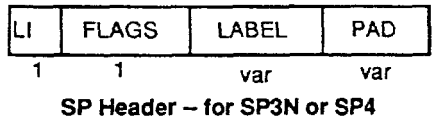

SP Header – for SP3N or SP4

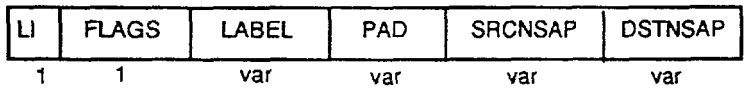

SP3A Header

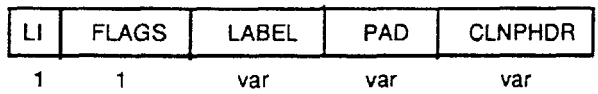

SP3I Header

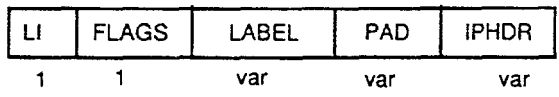

SP3D Header

SP4 Header with FSN

LI indicates length of protected header in octets
Label and PAD are optional
SRCNSAP and DSTNSAP are network addresses
CLNPHDR and IPHDR are complete network headers
FSN is present only for connection release TPDUs
All variable length fields are given as Type, Length, Value

Figure 7. SP Protected Header Options

REFERENCES

1. Ruth Nelson, SDNS Services and Architecture, National Computer Security Conference, Baltimore, Maryland, October, 1988

2. ISO 7498, Information Processing Systems -- Open Systems Interconnection -- Basic Reference Model

3. ISO 7498/2, Information Processing Systems -- Open Systems Interconnection -- Security Architecture

4. SDNS Program Office, SDN.301, Revision 1.3, Security Protocol 3 (SP3), July 1988

5. SDNS Program Office, SDN.401, Revision 1.2, Security Protocol 4 (SP4), July 1988

6. ISO 8648, Information Processing Systems -- Data Communications -- Internal Organization of the Network Layer