

IP Encapsulating Security Payload (ESP)

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.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Table of Contents

1. Introduction.....	2
2. Encapsulating Security Payload Packet Format.....	3
2.1 Security Parameters Index.....	4
2.2 Sequence Number	4
2.3 Payload Data.....	5
2.4 Padding (for Encryption).....	5
2.5 Pad Length.....	7
2.6 Next Header.....	7
2.7 Authentication Data.....	7
3. Encapsulating Security Protocol Processing.....	7
3.1 ESP Header Location.....	7
3.2 Algorithms.....	10
3.2.1 Encryption Algorithms.....	10
3.2.2 Authentication Algorithms.....	10
3.3 Outbound Packet Processing.....	10
3.3.1 Security Association Lookup.....	11
3.3.2 Packet Encryption.....	11
3.3.3 Sequence Number Generation.....	12
3.3.4 Integrity Check Value Calculation.....	12
3.3.5 Fragmentation.....	13
3.4 Inbound Packet Processing.....	13
3.4.1 Reassembly.....	13
3.4.2 Security Association Lookup.....	13
3.4.3 Sequence Number Verification.....	14
3.4.4 Integrity Check Value Verification.....	15

3.4.5 Packet Decryption.....	16
4. Auditing.....	17
5. Conformance Requirements.....	18
6. Security Considerations.....	18
7. Differences from RFC 1827.....	18
Acknowledgements.....	19
References.....	19
Disclaimer.....	20
Author Information.....	21
Full Copyright Statement.....	22

1. Introduction

The Encapsulating Security Payload (ESP) header is designed to provide a mix of security services in IPv4 and IPv6. ESP may be applied alone, in combination with the IP Authentication Header (AH) [KA97b], or in a nested fashion, e.g., through the use of tunnel mode (see "Security Architecture for the Internet Protocol" [KA97a], hereafter referred to as the Security Architecture document). Security services can be provided between a pair of communicating hosts, between a pair of communicating security gateways, or between a security gateway and a host. For more details on how to use ESP and AH in various network environments, see the Security Architecture document [KA97a].

The ESP header is inserted after the IP header and before the upper layer protocol header (transport mode) or before an encapsulated IP header (tunnel mode). These modes are described in more detail below.

ESP is used to provide confidentiality, data origin authentication, connectionless integrity, an anti-replay service (a form of partial sequence integrity), and limited traffic flow confidentiality. The set of services provided depends on options selected at the time of Security Association establishment and on the placement of the implementation. Confidentiality may be selected independent of all other services. However, use of confidentiality without integrity/authentication (either in ESP or separately in AH) may subject traffic to certain forms of active attacks that could undermine the confidentiality service (see [Bel96]). Data origin authentication and connectionless integrity are joint services (hereafter referred to jointly as "authentication") and are offered as an option in conjunction with (optional) confidentiality. The anti-replay service may be selected only if data origin authentication is selected, and its election is solely at the discretion of the receiver. (Although the default calls for the sender to increment the Sequence Number used for anti-replay, the service is effective only if the receiver checks the Sequence Number.) Traffic flow

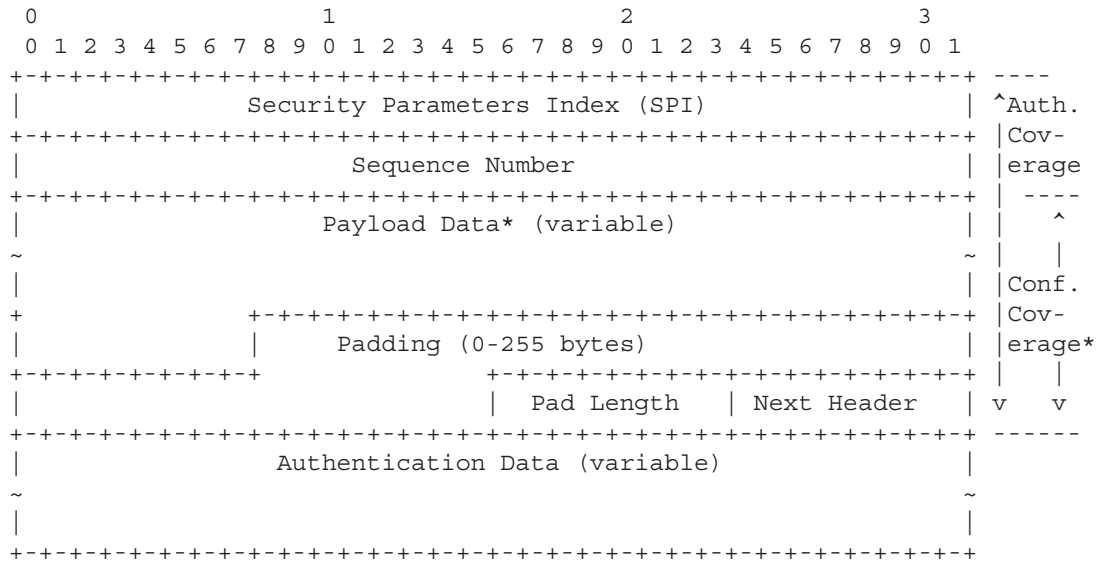
confidentiality requires selection of tunnel mode, and is most effective if implemented at a security gateway, where traffic aggregation may be able to mask true source-destination patterns. Note that although both confidentiality and authentication are optional, at least one of them MUST be selected.

It is assumed that the reader is familiar with the terms and concepts described in the Security Architecture document. In particular, the reader should be familiar with the definitions of security services offered by ESP and AH, the concept of Security Associations, the ways in which ESP can be used in conjunction with the Authentication Header (AH), and the different key management options available for ESP and AH. (With regard to the last topic, the current key management options required for both AH and ESP are manual keying and automated keying via IKE [HC98].)

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in RFC 2119 [Bra97].

2. Encapsulating Security Payload Packet Format

The protocol header (IPv4, IPv6, or Extension) immediately preceding the ESP header will contain the value 50 in its Protocol (IPv4) or Next Header (IPv6, Extension) field [STD-2].



* If included in the Payload field, cryptographic synchronization data, e.g., an Initialization Vector (IV, see



Section 2.3), usually is not encrypted per se, although it often is referred to as being part of the ciphertext.

The following subsections define the fields in the header format. "Optional" means that the field is omitted if the option is not selected, i.e., it is present in neither the packet as transmitted nor as formatted for computation of an Integrity Check Value (ICV, see Section 2.7). Whether or not an option is selected is defined as part of Security Association (SA) establishment. Thus the format of ESP packets for a given SA is fixed, for the duration of the SA. In contrast, "mandatory" fields are always present in the ESP packet format, for all SAs.

2.1 Security Parameters Index

The SPI is an arbitrary 32-bit value that, in combination with the destination IP address and security protocol (ESP), uniquely identifies the Security Association for this datagram. The set of SPI values in the range 1 through 255 are reserved by the Internet Assigned Numbers Authority (IANA) for future use; a reserved SPI value will not normally be assigned by IANA unless the use of the assigned SPI value is specified in an RFC. It is ordinarily selected by the destination system upon establishment of an SA (see the Security Architecture document for more details). The SPI field is mandatory.

The SPI value of zero (0) is reserved for local, implementation-specific use and MUST NOT be sent on the wire. For example, a key management implementation MAY use the zero SPI value to mean "No Security Association Exists" during the period when the IPsec implementation has requested that its key management entity establish a new SA, but the SA has not yet been established.

2.2 Sequence Number

This unsigned 32-bit field contains a monotonically increasing counter value (sequence number). It is mandatory and is always present even if the receiver does not elect to enable the anti-replay service for a specific SA. Processing of the Sequence Number field is at the discretion of the receiver, i.e., the sender MUST always transmit this field, but the receiver need not act upon it (see the discussion of Sequence Number Verification in the "Inbound Packet Processing" section below).

The sender's counter and the receiver's counter are initialized to 0 when an SA is established. (The first packet sent using a given SA will have a Sequence Number of 1; see Section 3.3.3 for more details on how the Sequence Number is generated.) If anti-replay is enabled

(the default), the transmitted Sequence Number must never be allowed to cycle. Thus, the sender's counter and the receiver's counter MUST be reset (by establishing a new SA and thus a new key) prior to the transmission of the 2³²nd packet on an SA.

2.3 Payload Data

Payload Data is a variable-length field containing data described by the Next Header field. The Payload Data field is mandatory and is an integral number of bytes in length. If the algorithm used to encrypt the payload requires cryptographic synchronization data, e.g., an Initialization Vector (IV), then this data MAY be carried explicitly in the Payload field. Any encryption algorithm that requires such explicit, per-packet synchronization data MUST indicate the length, any structure for such data, and the location of this data as part of an RFC specifying how the algorithm is used with ESP. If such synchronization data is implicit, the algorithm for deriving the data MUST be part of the RFC.

Note that with regard to ensuring the alignment of the (real) ciphertext in the presence of an IV:

- o For some IV-based modes of operation, the receiver treats the IV as the start of the ciphertext, feeding it into the algorithm directly. In these modes, alignment of the start of the (real) ciphertext is not an issue at the receiver.
- o In some cases, the receiver reads the IV in separately from the ciphertext. In these cases, the algorithm specification MUST address how alignment of the (real) ciphertext is to be achieved.

2.4 Padding (for Encryption)

Several factors require or motivate use of the Padding field.

- o If an encryption algorithm is employed that requires the plaintext to be a multiple of some number of bytes, e.g., the block size of a block cipher, the Padding field is used to fill the plaintext (consisting of the Payload Data, Pad Length and Next Header fields, as well as the Padding) to the size required by the algorithm.
- o Padding also may be required, irrespective of encryption algorithm requirements, to ensure that the resulting ciphertext terminates on a 4-byte boundary. Specifically,

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.