Network Working Group
Internet-Draft
Expires: September 3, 2006

P. Jokela
Ericsson Research NomadicLab
R. Moskowitz
ICSAlabs, a Division of TruSecure Corporation
P. Nikander
Ericsson Research NomadicLab

March 2, 2006

Using ESP transport format with HIP
draft-ietf-hip-esp-02

Status of this Memo

Copyright Notice

Abstract

This memo specifies an Encapsulated Security Payload (ESP) based mechanism for transmission of user data packets, to be used with the Host Identity Protocol (HIP).

Table of Contents

1. Introduction 4
2. Conventions used in this document 5
3. Using ESP with HIP 6
3.1. ESP Packet Format 6
3.2. Conceptual ESP Packet Processing 6
3.2.1. Semantics of the Security Parameter Index (SPI) . . . 7
3.3. Security Association Establishment and Maintenance 7
3.3.1. ESP Security Associations 8
3.3.2. Rekeying 8
3.3.3. Security Association Management 9
3.3.4. Security Parameter Index (SPI) 9
3.3.5. Supported Transforms 9
3.3.6. Sequence Number 10
3.3.7. Lifetimes and Timers 10
4. The Protocol 11
4.1. ESP in HIP 11
4.1.1. Setting up an ESP Security Association 11
4.1.2. Updating an Existing ESP SA 12
5. Parameter and Packet Formats 13
5.1. New Parameters 13
5.1.1. ESP_INFO 13
5.1.2. ESP_TRANSFORM 14
5.1.3. NOTIFY Parameter 15
5.2. HIP ESP Security Association Setup 16
5.2.1. Setup During Base Exchange 16
5.3. HIP ESP Rekeying 17
5.3.1. Initializing Rekeying 18
5.3.2. Responding to the Rekeying Initialization 18
5.4. ICMP Messages 19
5.4.1. Unknown SPI 19
6. Packet Processing 20
6.1. Processing Outgoing Application Data 20
6.2. Processing Incoming Application Data 20
6.3. HMAC and SIGNATURE Calculation and Verification 21
6.4. Processing Incoming ESP SA Initialization (R1) 21
6.5. Processing Incoming Initialization Reply (I2) 22
6.6. Processing Incoming ESP SA Setup Finalization (R2) 22
6.7. Dropping HIP Associations 22
6.8. Initiating ESP SA Rekeying 22
6.9. Processing Incoming UPDATE Packets 24
6.9.1. Processing UPDATE Packet: No Outstanding Rekeying Request 24
6.10. Finalizing Rekeying 25
6.11. Processing NOTIFY Packets 26
7. Keying Material 27
8. Security Considerations 28

9. IANA Considerations 29
10. References 30
10.1. Normative references 30
10.2. Informative references 30
Appendix A. A Note on Implementation Options 31
Authors' Addresses 32
Intellectual Property and Copyright Statements 33

1. Introduction

In the Host Identity Protocol Architecture [7], hosts are identified with public keys. The Host Identity Protocol [5] base exchange allows any two HIP-supporting hosts to authenticate each other and to create a HIP association between themselves. During the base exchange, the hosts generate a piece of shared keying material using an authenticated Diffie-Hellman exchange.

The HIP base exchange specification [5] does not describe any transport formats, or methods for user data, to be used during the actual communication; it only defines that it is mandatory to implement the Encapsulated Security Payload (ESP) [4] based transport format and method. This document specifies how ESP is used with HIP to carry actual user data.

To be more specific, this document specifies a set of HIP protocol extensions and their handling. Using these extensions, a pair of ESP Security Associations (SAs) is created between the hosts during the base exchange. The resulting ESP Security Associations use keys drawn from the keying material (KEYMAT) generated during the base exchange. After the HIP association and required ESP SAs have been established between the hosts, the user data communication is protected using ESP. In addition, this document specifies methods to update an existing ESP Security Association.

It should be noted that representations of host identity are not carried explicitly in the headers of user data packets. Instead, the ESP Security Parameter Index (SPI) is used to indicate the right host context. The SPIs are selected during the HIP ESP setup exchange. For user data packets, ESP SPIs (in possible combination with IP addresses) are used indirectly to identify the host context, thereby avoiding any additional explicit protocol headers.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [1].

3. Using ESP with HIP

The HIP base exchange is used to set up a HIP association between two hosts. The base exchange provides two-way host authentication and key material generation, but it does not provide any means for protecting data communication between the hosts. In this document we specify the use of ESP for protecting user data traffic after the HIP base exchange. Note that this use of ESP is intended only for host-to-host traffic; security gateways are not supported.

To support ESP use, the HIP base exchange messages require some minor additions to the parameters transported. In the R1 packet, the responder adds the possible ESP transforms in a new ESP_TRANSFORM parameter before sending it to the Initiator. The Initiator gets the proposed transforms, selects one of those proposed transforms, and adds it to the I2 packet in an ESP_TRANSFORM parameter. In this I2 packet, the Initiator also sends the SPI value that it wants to be used for ESP traffic flowing from the Responder to the Initiator. This information is carried using the new ESP_INFO parameter. When finalizing the ESP SA setup, the Responder sends its SPI value to the Initiator in the R2 packet, again using ESP_INFO.

3.1. ESP Packet Format

The ESP specification [4] defines the ESP packet format for IPsec. The HIP ESP packet looks exactly the same as the IPsec ESP transport format packet. The semantics, however, are a bit different and are described in more detail in the next subsection.

3.2. Conceptual ESP Packet Processing

ESP packet processing can be implemented in different ways in HIP. It is possible to implement it in a way that a standards compliant, unmodified IPsec implementation [4] can be used.

When a standards compliant IPsec implementation that uses IP addresses in the SPD and SAD is used, the packet processing may take the following steps. For outgoing packets, assuming that the upper layer pseudoheader has been built using IP addresses, the implementation recalculates upper layer checksums using HITs and, after that, changes the packet source and destination addresses back to corresponding IP addresses. The packet is sent to the IPsec ESP for transport mode handling and from there the encrypted packet is sent to the network. When an ESP packet is received, the packet is first put to the IPsec ESP transport mode handling, and after decryption, the source and destination IP addresses are replaced with HITs and finally, upper layer checksums are verified before passing the packet to the upper layer.

An alternative way to implement the packet processing is the BEET (Bound End-to-End Tunnel) [11] mode. In BEET mode, the ESP packet is formatted as a transport mode packet, but the semantics of the connection are the same as for tunnel mode. The "outer" addresses of the packet are the IP addresses and the "inner" addresses are the HITs. For outgoing traffic, after the packet has been encrypted, the packet's IP header is changed to a new one, containing IP addresses instead of HITs and the packet is sent to the network. When ESP packet is received, the SPI value, together with the integrity protection, allow the packet to be securely associated with the right HIT pair. The packet header is replaces with a new header, containing HITs and the packet is decrypted.

3.2.1. Semantics of the Security Parameter Index (SPI)

SPIs are used in ESP to find the right Security Association for received packets. The ESP SPIs have added significance when used with HIP; they are a compressed representation of a pair of HITs. Thus, SPIs MAY be used by intermediary systems in providing services like address mapping. Note that since the SPI has significance at the receiver, only the < DST, SPI >, where DST is a destination IP address, uniquely identifies the receiver HIT at any given point of time. The same SPI value may be used by several hosts. A single < DST, SPI > value may denote different hosts and contexts at different points of time, depending on the host that is currently reachable at the DST.

Each host selects for itself the SPI it wants to see in packets received from its peer. This allows it to select different SPIs for different peers. The SPI selection SHOULD be random; the rules of Section 2.1 of the ESP specification [4] must be followed. A different SPI SHOULD be used for each HIP exchange with a particular host; this is to avoid a replay attack. Additionally, when a host rekeys, the SPI MUST be changed. Furthermore, if a host changes over to use a different IP address, it MAY change the SPI.

One method for SPI creation that meets the above criteria would be to concatenate the HIT with a 32-bit random or sequential number, hash this (using SHA1), and then use the high order 32 bits as the SPI.

The selected SPI is communicated to the peer in the third (I2) and fourth (R2) packets of the base HIP exchange. Changes in SPI are signaled with ESP_INFO parameters.

3.3. Security Association Establishment and Maintenance

3.3.1. ESP Security Associations

In HIP, ESP Security Associations are setup between the HIP nodes during the base exchange [5]. Existing ESP SAs can be updated later using UPDATE messages. The reason for updating the ESP SA later can be e.g. need for rekeying the SA because of sequence number rollover.

Upon setting up a HIP association, each association is linked to two ESP SAs, one for incoming packets and one for outgoing packets. The Initiator's incoming SA corresponds with the Responder's outgoing one, and vice versa. The Initiator defines the SPI for its incoming association, as defined in Section 3.2.1. This SA is herein called SA-RI, and the corresponding SPI is called SPI-RI. Respectively, the Responder's incoming SA corresponds with the Initiator's outgoing SA and is called SA-IR, with the SPI being called SPI-IR.

The Initiator creates SA-RI as a part of R1 processing, before sending out the I2, as explained in Section 6.4. The keys are derived from KEYMAT, as defined in Section 7. The Responder creates SA-RI as a part of I2 processing, see Section 6.5.

The Responder creates SA-IR as a part of I2 processing, before sending out R2; see Section 6.5. The Initiator creates SA-IR when processing R2; see Section 6.6.

The initial session keys are drawn from the generated keying material, KEYMAT, after the HIP keys have been drawn as specified in [5].

When the HIP association is removed, the related ESP SAs MUST also be removed.

3.3.2. Rekeying

After the initial HIP base exchange and SA establishment, both hosts are in the ESTABLISHED state. There are no longer Initiator and Responder roles and the association is symmetric. In this subsection, the party that initiates the rekey procedure is denoted with I' and the peer with R'.

An existing HIP-created ESP SA may need updating during the lifetime of the HIP association. This document specifies the rekeying of an existing HIP-created ESP SA, using the UPDATE message. The ESP_INFO parameter introduced above is used for this purpose.

I' initiates the ESP SA updating process when needed (see Section 6.8). It creates an UPDATE packet with required information and sends it to the peer node. The old SAs are still in use, local

policy permitting.

R', after receiving and processing the UPDATE (see Section 6.9), generates new SAs: SA-I'R' and SA-R'I'. It does not take the new outgoing SA into use, but still uses the old one, so there temporarily exists two SA pairs towards the same peer host. The SPI for the new outgoing SA, SPI-R'I', is specified in the received ESP_INFO parameter in the UPDATE packet. For the new incoming SA, R' generates the new SPI value, SPI-I'R', and includes it in the response UPDATE packet.

When I' receives a response UPDATE from R', it generates new SAs, as described in Section 6.9: SA-I'R' and SA-R'I'. It starts using the new outgoing SA immediately.

R' starts using the new outgoing SA when it receives traffic on the new incoming SA or when it receives the UPDATE ACK confirming completion of rekeying. After this, R' can remove the old SAs. Similarly, when the I' receives traffic from the new incoming SA, it can safely remove the old SAs.

3.3.3. Security Association Management

An SA pair is indexed by the 2 SPIs and 2 HITs (both local and remote HITs since a system can have more than one HIT). An inactivity timer is RECOMMENDED for all SAs. If the state dictates the deletion of an SA, a timer is set to allow for any late arriving packets.

3.3.4. Security Parameter Index (SPI)

The SPIs in ESP provide a simple compression of the HIP data from all packets after the HIP exchange. This does require a per HIT-pair Security Association (and SPI), and a decrease of policy granularity over other Key Management Protocols like IKE.

When a host updates the ESP SA, it provides a new inbound SPI to and gets a new outbound SPI from its partner.

3.3.5. Supported Transforms

All HIP implementations MUST support AES [3] and HMAC-SHA-1-96 [2]. If the Initiator does not support any of the transforms offered by the Responder, it should abandon the negotiation and inform the peer with a NOTIFY message about a non-supported transform.

In addition to AES, all implementations MUST implement the ESP NULL encryption algorithm. When the ESP NULL encryption is used, it MUST be used together with SHA1 or MD5 authentication as specified in

Section 5.1.2

3.3.6. Sequence Number

The Sequence Number field is MANDATORY when ESP is used with HIP. Anti-replay protection MUST be used in an ESP SA established with HIP. This means that each host MUST rekey before its sequence number reaches 2^32, or if extended sequence numbers are used, 2^64.

In some instances, a 32-bit sequence number is inadequate. In the ESP_TRANSFORM parameter, a peer MAY require that a 64-bit sequence numbers be used. In this case the higher 32 bits are NOT included in the ESP header, but are simply kept local to both peers. The 64-bit sequence number is required in fast networks when there is a risk that the sequence number will rollover too often. See [9].

3.3.7. Lifetimes and Timers

HIP does not negotiate any lifetimes. All ESP lifetimes are local policy. The only lifetimes a HIP implementation MUST support are sequence number rollover (for replay protection), and SHOULD support timing out inactive ESP SAs. An SA times out if no packets are received using that SA. The default timeout value is 15 minutes. Implementations MAY support lifetimes for the various ESP transforms. Each implementation SHOULD implement per-HIT configuration of the inactivity timeout, allowing statically configured HIP associations to stay alive for days, even when inactive.

4. The Protocol

In this section, the protocol for setting up an ESP association to be used with HIP association is described.

4.1. ESP in HIP

4.1.1. Setting up an ESP Security Association

Setting up an ESP Security Association between hosts using HIP consists of three messages passed between the hosts. The parameters are included in R1, I2, and R2 messages during base exchange.

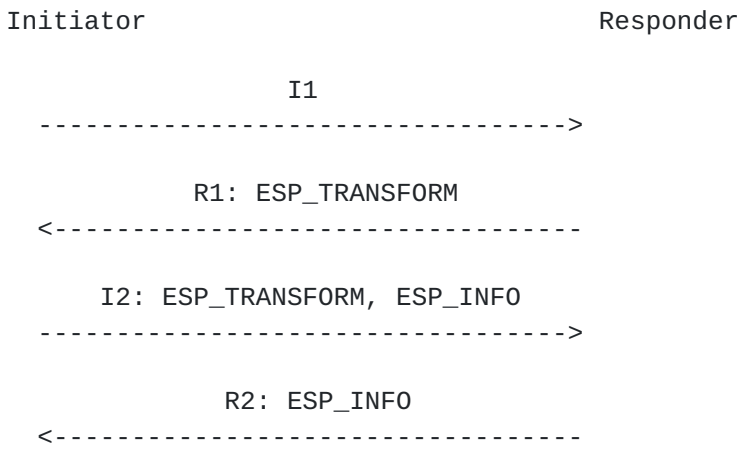

Setting up an ESP Security Association between HIP hosts requires three messages to exchange the information that is required during an ESP communication.

The R1 message contains the ESP_TRANSFORM parameter, in which the sending host defines the possible ESP transforms it is willing to use for the ESP SA.

The I2 message contains the response to an ESP_TRANSFORM received in the R1 message. The sender must select one of the proposed ESP transforms from the ESP_TRANSFORM parameter in the R1 message and include the selected one in the ESP_TRANSFORM parameter in the I2 packet. In addition to the transform, the host includes the ESP_INFO parameter, containing the SPI value to be used by the peer host.

In the R2 message, the ESP SA setup is finalized. The packet contains the SPI information required by the Initiator for the ESP SA.

4.1.2. Updating an Existing ESP SA

The update process is accomplished using two messages. The HIP UPDATE message is used to update the parameters of an existing ESP SA. The UPDATE mechanism and message is defined in [5] and the additional parameters for updating an existing ESP SA are described here.

The following picture shows a typical exchange when an existing ESP SA is updated. Messages include SEQ and ACK parameters required by the UPDATE mechanism.

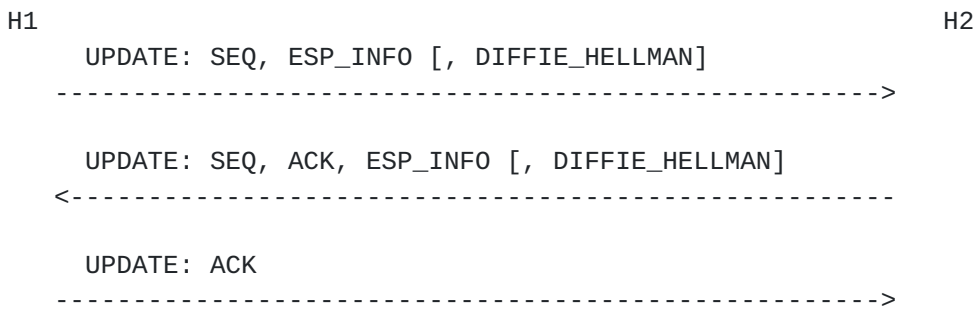

The host willing to update the ESP SA creates and sends an UPDATE message. The message contains the ESP_INFO parameter, containing the old SPI value that was used, the new SPI value to be used, and the index value for the keying material, giving the point from where the next keys will be drawn. If new keying material must be generated, the UPDATE message will also contain the DIFFIE_HELLMAN parameter, defined in [5].

The host receiving the UPDATE message requesting update of an existing ESP SA, MUST reply with an UPDATE message. In the reply message, the host sends the ESP_INFO parameter containing the corresponding values: old SPI, new SPI, and the keying material index. If the incoming UPDATE contained a DIFFIE_HELLMAN parameter, the reply packet MUST also contain a DIFFIE_HELLMAN parameter.

5. Parameter and Packet Formats

In this section, new and modified HIP parameters are presented, as well as modified HIP packets.

5.1. New Parameters

Two new HIP parameters are defined for setting up ESP transport format associations in HIP communication and for rekeying existing ones. Also, the NOTIFY parameter, described in [5], has two new error parameters.

Parameter	Type	Length	Data
ESP_INFO	65	12	Remote's old SPI, new SPI and other info
ESP_TRANSFORM	4095	variable	ESP Encryption and Authentication Transform(s)

5.1.1. ESP_INFO

During the establishment and update of an ESP SA, the SPI value of both hosts must be transmitted between the hosts. Additional information that is required when the hosts are drawing keys from the generated keying material is the index value into the KEYMAT from where the keys are drawn. The ESP_INFO parameter is used to transmit this information between the hosts.

During the initial ESP SA setup, the hosts send the SPI value that they want the peer to use when sending ESP data to them. The value is set in the New SPI field of the ESP_INFO parameter. In the initial setup, an old value for the SPI does not exist, thus the Old SPI value field is set to zero. The Old SPI field value may also be zero when additional SAs are set up between HIP hosts, e.g. in case of multihomed HIP hosts [12]. However, such use is beyond the scope of this specification.

The Keymat index value points to the place in the KEYMAT from where the keying material for the ESP SAs is drawn. The Keymat index value is zero only when the ESP_INFO is sent during a rekeying process and new keying material is generated.

During the life of an SA established by HIP, one of the hosts may need to reset the Sequence Number to one and rekey. The reason for rekeying might be an approaching sequence number wrap in ESP, or a local policy on use of a key. Rekeying ends the current SAs and starts new ones on both peers.

During the rekeying process, the ESP_INFO parameter is used to transmit the changed SPI values and the keying material index.

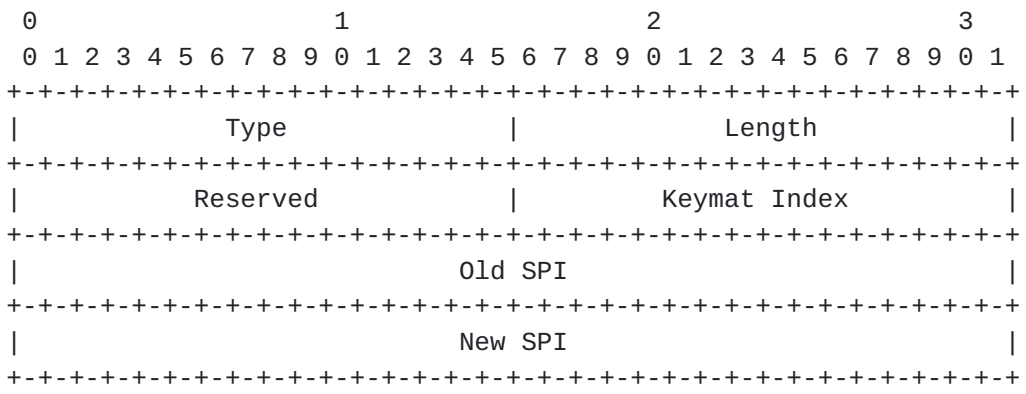

Type	65
Length	12
Keymat Index	Index, in bytes, where to continue to draw ESP keys from KEYMAT. If the packet includes a new Diffie-Hellman key and the ESP_INFO is sent in an UPDATE packet, the field MUST be zero. If the ESP_INFO is included in base exchange messages, the Keymat Index must have the index value of the point from where the ESP SA keys are drawn. Note that the length of this field limits the amount of keying material that can be drawn from KEYMAT. If that amount is exceeded, the packet MUST contain a new Diffie-Hellman key.
Old SPI	Old SPI for data sent to address(es) associated with this SA. If this is an initial SA setup, the Old SPI value is zero.
New SPI	New SPI for data sent to address(es) associated with this SA.

5.1.2. ESP_TRANSFORM

The ESP_TRANSFORM parameter is used during ESP SA establishment. The first party sends a selection of transfrom families in the ESP_TRANSFORM parameter and the peer must select one of the proposed values and include it in the response ESP_TRANSFORM parameter.

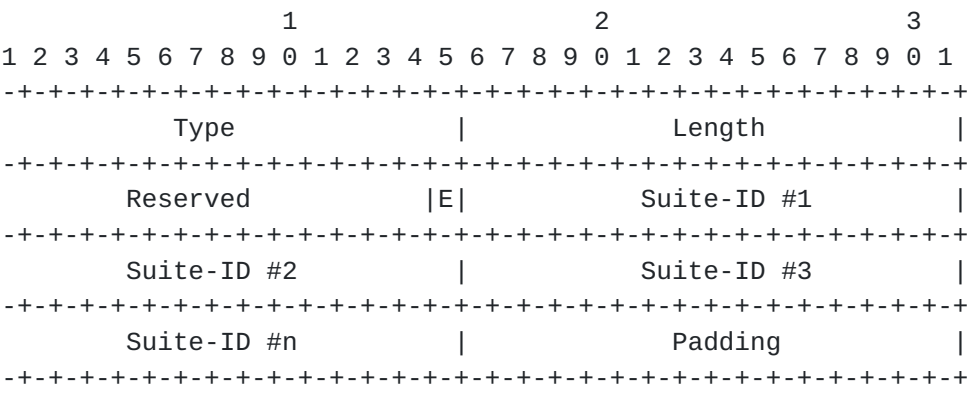

Type	4095
Length	length in octets, excluding Type, Length, and padding
E	One if the ESP transform requires 64-bit sequence numbers (see Section 3.3.6
Reserved	zero when sent, ignored when received
Suite-ID	defines the ESP Suite to be used

The following Suite-IDs are defined ([6],[8]):

Suite-ID	Value
RESERVED	0
ESP-AES-CBC with HMAC-SHA1	1
ESP-3DES-CBC with HMAC-SHA1	2
ESP-3DES-CBC with HMAC-MD5	3
ESP-BLOWFISH-CBC with HMAC-SHA1	4
ESP-NULL with HMAC-SHA1	5
ESP-NULL with HMAC-MD5	6

There MUST NOT be more than six (6) ESP Suite-IDs in one ESP_TRANSFORM parameter. The limited number of Suite-IDs sets the maximum size of ESP_TRANSFORM parameter. The ESP_TRANSFORM MUST contain at least one of the mandatory Suite-IDs.

Mandatory implementations: ESP-AES-CBC with HMAC-SHA1 and ESP-NULL with HMAC-SHA1.

5.1.3. NOTIFY Parameter

The HIP base specification defines a set of NOTIFY error types. The following error types are required for describing errors in ESP Transform crypto suites during negotiation.

NOTIFY PARAMETER - ERROR TYPES	Value
NO_ESP_PROPOSAL_CHOSEN	18

None of the proposed ESP Transform crypto suites was acceptable.

INVALID_ESP_TRANSFORM_CHOSEN	19

The ESP Transform crypto suite does not correspond to one offered by the responder.

5.2. HIP ESP Security Association Setup

The ESP Security Association is set up during the base exchange. The following subsections define the ESP SA setup procedure both using base exchange messages (R1, I2, R2) and using UPDATE messages.

5.2.1. Setup During Base Exchange

5.2.1.1. Modifications in R1

The ESP_TRANSFORM contains the ESP modes supported by the sender, in the order of preference. All implementations MUST support AES [3] with HMAC-SHA-1-96 [2].

The following figure shows the resulting R1 packet layout.

The HIP parameters for the R1 packet:

```
IP ( HIP ( [ R1_COUNTER, ]
           PUZZLE,
           DIFFIE_HELLMAN,
           HIP_TRANSFORM,
           ESP_TRANSFORM,
           HOST_ID,
           [ ECHO_REQUEST, ]
           HIP_SIGNATURE_2 )
           [, ECHO_REQUEST ])
```

5.2.1.2. Modifications in I2

The ESP_INFO contains the sender's SPI for this association as well as the keymat index from where the ESP SA keys will be drawn. The Old SPI value is set to zero.

The ESP_TRANSFORM contains the ESP mode selected by the sender of R1. All implementations MUST support AES [3] with HMAC-SHA-1-96 [2].

The following figure shows the resulting I2 packet layout.

The HIP parameters for the I2 packet:

```
IP ( HIP ( ESP_INFO,
           [R1_COUNTER,]
           SOLUTION,
           DIFFIE_HELLMAN,
           HIP_TRANSFORM,
           ESP_TRANSFORM,
           ENCRYPTED { HOST_ID },
           [ ECHO_RESPONSE ,]
           HMAC,
           HIP_SIGNATURE
           [, ECHO_RESPONSE] ) )
```

5.2.1.3. Modifications in R2

The R2 contains an ESP_INFO parameter, which has the SPI value of the sender of the R2 for this association. The ESP_INFO also has the keymat index value specifying where the ESP SA keys are drawn.

The following figure shows the resulting R2 packet layout.

The HIP parameters for the R2 packet:

```
IP ( HIP ( ESP_INFO, HMAC_2, HIP_SIGNATURE ) )
```

5.3. HIP ESP Rekeying

In this section, the procedure for rekeying an existing ESP SA is presented.

Conceptually, the process can be represented by the following message sequence using the host names I' and R' defined in Section 3.3.2. For simplicity, HMAC and HIP_SIGNATURE are not depicted, and DIFFIE_HELLMAN keys are optional. The UPDATE with ACK_I need not be piggybacked with the UPDATE with SEQ_R; it may be acked separately (in which case the sequence would include four packets).

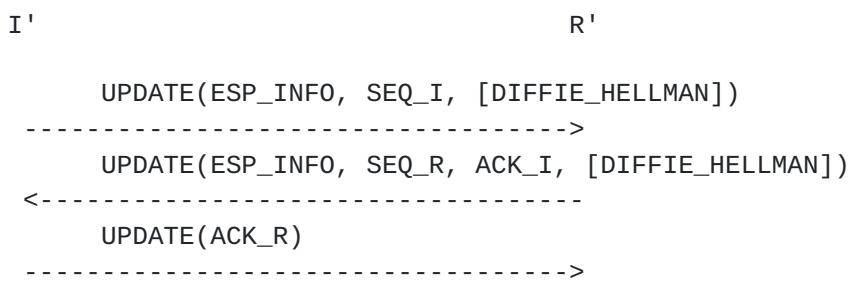

Below, the first two packets in this figure are explained.

5.3.1. Initializing Rekeying

When HIP is used with ESP, the UPDATE packet is used to initiate rekeying. The UPDATE packet MUST carry an ESP_INFO and MAY carry a DIFFIE_HELLMAN parameter.

Intermediate systems that use the SPI will have to inspect HIP packets for those that carry rekeying information. The packet is signed for the benefit of the intermediate systems. Since intermediate systems may need the new SPI values, the contents cannot be encrypted.

The following figure shows the contents of a rekeying initialization UPDATE packet.

The HIP parameters for the UPDATE packet initiating rekeying:

```
IP ( HIP ( ESP_INFO,
           SEQ,
           [DIFFIE_HELLMAN, ]
           HMAC,
           HIP_SIGNATURE ) )
```

5.3.2. Responding to the Rekeying Initialization

The UPDATE ACK is used to acknowledge the received UPDATE rekeying initialization. The acknowledgement UPDATE packet MUST carry an ESP_INFO and MAY carry a DIFFIE_HELLMAN parameter.

Intermediate systems that use the SPI will have to inspect HIP packets for packets carrying rekeying information. The packet is signed for the benefit of the intermediate systems. Since intermediate systems may need the new SPI values, the contents cannot be encrypted.

The following figure shows the contents of a rekeying acknowledgement UPDATE packet.

The HIP parameters for the UPDATE packet:

```
IP ( HIP ( ESP_INFO,
           SEQ,
           ACK,
           [ DIFFIE_HELLMAN, ]
           HMAC,
           HIP_SIGNATURE ) )
```

5.4. ICMP Messages

The ICMP message handling is mainly described in the HIP base specification [5]. In this section, we describe the actions related to ESP security associations.

5.4.1. Unknown SPI

If a HIP implementation receives an ESP packet that has an unrecognized SPI number, it MAY respond (subject to rate limiting the responses) with an ICMP packet with type "Parameter Problem", with the Pointer pointing to the the beginning of SPI field in the ESP header.

6. Packet Processing

Packet processing is mainly defined in the HIP base specification [5]. This section describes the changes and new requirements for packet handling when the ESP transport format is used. Note that all HIP packets (currently protocol 99) MUST bypass ESP processing.

6.1. Processing Outgoing Application Data

Outgoing application data handling is specified in the HIP base specification [5]. When ESP transport format is used, and there is an active HIP session for the given < source, destination > HIT pair, the outgoing datagram is protected using the ESP security association. In a typical implementation, this will result in a BEET-mode ESP packet being sent. BEET-mode [11] was introduced above in Section 3.2.

1. Detect the proper ESP SA using the HITs in the packet header or other information associated with the packet

2. Process the packet normally, as if the SA was a transport mode SA.

3. Ensure that the outgoing ESP protected packet has proper IP header format depending on the used IP address family, and proper IP addresses in its IP header, e.g., by replacing HITs left by the ESP processing. Note that this placement of proper IP addresses MAY also be performed at some other point in the stack, e.g., before ESP processing.

6.2. Processing Incoming Application Data

Incoming HIP user data packets arrive as ESP protected packets. In the usual case the receiving host has a corresponding ESP security association, identified by the SPI and destination IP address in the packet. However, if the host has crashed or otherwise lost its HIP state, it may not have such an SA.

The basic incoming data handling is specified in the HIP base specification. Additional steps are required when ESP is used for protecting the data traffic. The following steps define the conceptual processing rules for incoming ESP protected datagrams targeted to an ESP security association created with HIP.

1. Detect the proper ESP SA using the SPI. If the resulting SA is a non-HIP ESP SA, process the packet according to standard IPsec rules. If there are no SAs identified with the SPI, the host MAY send an ICMP packet as defined in Section 5.4. How to handle

lost state is an implementation issue.

2. If the SPI matches with an active HIP-based ESP SA, the IP addresses in the datagram are replaced with the HITs associated with the SPI. Note that this IP-address-to-HIT conversion step MAY also be performed at some other point in the stack, e.g., after ESP processing. Note also that if the incoming packet has IPv4 addresses, the packet must be converted to IPv6 format before replacing the addresses with HITs (such that the transport checksum will pass if there are no errors).

3. The transformed packet is next processed normally by ESP, as if the packet were a transport mode packet. The packet may be dropped by ESP, as usual. In a typical implementation, the result of successful ESP decryption and verification is a datagram with the associated HITs as source and destination.

4. The datagram is delivered to the upper layer. Demultiplexing the datagram to the right upper layer socket is performed as usual, except that the HITs are used in place of IP addresses during the demultiplexing.

6.3. HMAC and SIGNATURE Calculation and Verification

The new HIP parameters described in this document, ESP_INFO and ESP_TRANSFORM, must be protected using HMAC and signature calculations. In a typical implementation, they are included in R1, I2, R2, and UPDATE packet HMAC and SIGNATURE calculations as described in [5].

6.4. Processing Incoming ESP SA Initialization (R1)

The ESP SA setup is initialized in the R1 message. The receiving host (Initiator) select one of the ESP transforms from the presented values. If no suitable value is found, the negotiation is terminated. The selected values are subsequently used when generating and using encryption keys, and when sending the reply packet. If the proposed alternatives are not acceptable to the system, it may abandon the ESP SA establishment negotiation, or it may resend the I1 message within the retry bounds.

After selecting the ESP transform, and performing other R1 processing, the system prepares and creates an incoming ESP security association. It may also prepare a security association for outgoing traffic, but since it does not have the correct SPI value yet, it cannot activate it.

6.5. Processing Incoming Initialization Reply (I2)

The following steps are required to process the incoming ESP SA initialization replies in I2. The steps below assume that the I2 has been accepted for processing (e.g., has not been dropped due to HIT comparisons as described in [5]).

- The ESP_TRANSFORM parameter is verified and it MUST contain a single value in the parameter and it MUST match one of the values offered in the initialization packet.

- The ESP_INFO New SPI field is parsed to obtain the SPI that will be used for the Security Association outbound from the Responder and inbound to the Initiator. For this initial ESP SA establishment, the Old SPI value MUST be zero. The Keymat Index field MUST contain the index value to the KEYMAT from where the ESP SA keys are drawn.

- The system prepares and creates both incoming and outgoing ESP security associations.

- Upon successful processing of the initialization reply message, the possible old Security Associations (as left over from an earlier incarnation of the HIP association) are dropped and the new ones are installed, and a finalizing packet, R2, is sent. Possible ongoing rekeying attempts are dropped.

6.6. Processing Incoming ESP SA Setup Finalization (R2)

Before the ESP SA can be finalized, the ESP_INFO New SPI field is parsed to obtain the SPI that will be used for the ESP Security Association inbound to the sender of the finalization message R2. The system uses this SPI to create or activate the outgoing ESP security association used for sending packets to the peer.

6.7. Dropping HIP Associations

When the system drops a HIP association, as described in the HIP base specification, the associated ESP SAs MUST also be dropped.

6.8. Initiating ESP SA Rekeying

During ESP SA rekeying, the hosts draw new keys from the existing keying material, or a new keying material is generated from where the new keys are drawn.

A system may initiate the SA rekeying procedure at any time. It MUST initiate a rekey if its incoming ESP sequence counter is about to

overflow. The system MUST NOT replace its keying material until the rekeying packet exchange successfully completes.

Optionally, a system may include a new Diffie-Hellman key for use in new KEYMAT generation. New KEYMAT generation occurs prior to drawing the new keys.

The rekeying procedure uses the UPDATE mechanism defined in [5]. Because each peer must update its half of the security association pair (including new SPI creation), the rekeying process requires that each side both send and receive an UPDATE. A system will then rekey the ESP SA when it has sent parameters to the peer and has received both an ACK of the relevant UPDATE message and corresponding peer's parameters. It may be that the ACK and the required HIP parameters arrive in different UPDATE messages. This is always true if a system does not initiate ESP SA update but responds to an update request from the peer, but may also occur if two systems initiate update nearly simultaneously. In such a case, if the system has an outstanding update request, it saves the one parameter and waits for the other before completing rekeying.

The following steps define the processing rules for initiating an ESP SA update:

1. The system decides whether to continue to use the existing KEYMAT or to generate new KEYMAT. In the latter case, the system MUST generate a new Diffie-Hellman public key.

2. The system creates an UPDATE packet, which contains the ESP_INFO parameter. In addition, the host may include the optional DIFFIE_HELLMAN parameter. If the UDPATE contains the DIFFIE_HELLMAN parameter, the Keymat Index in the ESP_INFO parameter MUST be zero, and the Diffie-Hellman group ID must be unchanged from that used in the initial handshake. If the UPDATE does not contain DIFFIE_HELLMAN, the ESP_INFO Keymat Index MUST be greater or equal to the index of the next byte to be drawn from the current KEYMAT.

3. The system sends the UPDATE packet. For reliability, the underlying UPDATE retransmission mechanism MUST be used.

4. The system MUST NOT delete its existing SAs, but continue using them if its policy still allows. The rekeying procedure SHOULD be initiated early enough to make sure that the SA replay counters do not overflow.

5. In case a protocol error occurs and the peer system acknowledges the UPDATE but does not itself send an ESP_INFO, the system may

not finalize the outstanding ESP SA update request. To guard against this, a system MAY re-initiate the ESP SA update procedure after some time waiting for the peer to respond, or it MAY decide to abort the ESP SA after waiting for an implementation-dependent time. The system MUST NOT keep an oustanding ESP SA update request for an indefinite time.

To simplify the state machine, a host MUST NOT generate new UPDATEs while it has an outstanding ESP SA update request, unless it is restarting the update process.

6.9. Processing Incoming UPDATE Packets

When a system receives an UPDATE packet, it must be processed if the following conditions hold (in addition to the generic conditions specified for UPDATE processing in Section 6.12 of [5]):

1. A corresponding HIP association must exist. This is usually ensured by the underlying UPDATE mechanism.

2. The state of the HIP association is ESTABLISHED or R2-SENT.

If the above conditions hold, the following steps define the conceptual processing rules for handling the received UPDATE packet:

1. If the received UPDATE contains a DIFFIE_HELLMAN parameter, the received Keymat Index MUST be zero and the Group ID must match the Group ID in use on the association. If this test fails, the packet SHOULD be dropped and the system SHOULD log an error message.

2. If there is no outstanding rekeying request, the packet processing continues as specified in Section 6.9.1.

3. If there is an outstanding rekeying request, the UPDATE MUST be acknowledged, the received ESP_INFO (and possibly DIFFIE_HELLMAN) parameters must be saved, and the packet processing continues as specified in Section 6.10.

6.9.1. Processing UPDATE Packet: No Outstanding Rekeying Request

The following steps define the conceptual processing rules for handling a received UPDATE packet with ESP_INFO parameter:

1. The system consults its policy to see if it needs to generate a new Diffie-Hellman key, and generates a new key (with same Group ID) if needed. The system records any newly generated or received Diffie-Hellman keys, for use in KEYMAT generation upon

finalizing the ESP SA update.

2. If the system generated a new Diffie-Hellman key in the previous step, or if it received a DIFFIE_HELLMAN parameter, it sets ESP_INFO Keymat Index to zero. Otherwise, the ESP_INFO Keymat Index MUST be greater or equal to the index of the next byte to be drawn from the current KEYMAT. In this case, it is RECOMMENDED that the host use the Keymat Index requested by the peer in the received ESP_INFO.

3. The system creates an UPDATE packet, which contains an ESP_INFO parameter, and the optional DIFFIE_HELLMAN parameter. This UPDATE would also typically acknowledge the peer's UPDATE with an ACK parameter, although a separate UPDATE ACK may be sent.

4. The system sends the UPDATE packet and stores any received ESP_INFO, and DIFFIE_HELLMAN parameters. At this point, it only needs to receive an acknowledgement for the newly sent UPDATE to finish ESP SA update. In the usual case, the acknowledgement is handled by the underlying UPDATE mechanism.

6.10. Finalizing Rekeying

A system finalizes rekeying when it has both received the corresponding UPDATE acknowledgement packet from the peer and it has successfully received the peer's UPDATE. The following steps are taken:

1. If the received UPDATE messages contains a new Diffie-Hellman key, the system has a new Diffie-Hellman key due to initiating ESP SA update, or both, the system generates new KEYMAT. If there is only one new Diffie-Hellman key, the old existing key is used as the other key.

2. If the system generated new KEYMAT in the previous step, it sets Keymat Index to zero, independent of whether the received UPDATE included a Diffie-Hellman key or not. If the system did not generate new KEYMAT, it uses the greater Keymat Index of the two (sent and received) ESP_INFO parameters.

3. The system draws keys for new incoming and outgoing ESP SAs, starting from the Keymat Index, and prepares new incoming and outgoing ESP SAs. The SPI for the outgoing SA is the new SPI value received in an ESP_INFO parameter. The SPI for the incoming SA was generated when the ESP_INFO was sent to the peer. The order of the keys retrieved from the KEYMAT during rekeying process is similar to that described in Section 7. Note, that only IPsec ESP keys are retrieved during rekeying process, not

the HIP keys.

4. The system starts to send to the new outgoing SA and prepares to start receiving data on the new incoming SA. Once the system receives data on the new incoming SA it may safely delete the old SAs.

6.11. Processing NOTIFY Packets

The processing of NOTIFY packets is described in the HIP base specification.

7. Keying Material

The keying material is generated as described in the HIP base specification. During the base exchange, the initial keys are drawn from the generated material. After the HIP association keys have been drawn, the ESP keys are drawn in the following order:

SA-gl ESP encryption key for HOST_g's outgoing traffic

SA-gl ESP authentication key for HOST_g's outgoing traffic

SA-lg ESP encryption key for HOST_l's outgoing traffic

SA-lg ESP authentication key for HOST_l's outgoing traffic

HOST_g denotes the host with the greater HIT value, and HOST_l the host with the lower HIT value. When HIT values are compared, they are interpreted as positive (unsigned) 128-bit integers in network byte order.

The four HIP keys are only drawn from KEYMAT during a HIP I1->R2 exchange. Subsequent rekeys using UPDATE will only draw the four ESP keys from KEYMAT. Section 6.9 describes the rules for reusing or regenerating KEYMAT based on the rekeying.

The number of bits drawn for a given algorithm is the "natural" size of the keys. For the mandatory algorithms, the following sizes apply:

AES 128 bits

SHA-1 160 bits

NULL 0 bits

8. Security Considerations

In this document the usage of ESP [4] between HIP hosts to protect data traffic is introduced. The Security Considerations for ESP are discussed in the ESP specification.

There are different ways to establish an ESP Security Association between two nodes. This can be done, e.g. using IKE [10]. This document specifies how Host Identity Protocol is used to establish ESP Security Associations.

The following issues are new, or changed from the standard ESP usage:

- Initial keying material generation
- Updating the keying material

The initial keying material is generated using the Host Identity Protocol [5] using Diffie-Hellman procedure. This document extends the usage of UDPATE packet, defined in the base specification, to modify existing ESP SAs. The hosts may rekey, i.e. force the generation of new keying material using Diffie-Hellman procedure. The initial setup of ESP SA between the hosts is done during the base ecxhange and the message exchange is protected with using methods provided by base exchange. Changing of connection parameters means basically that the old ESP SA is removed and a new one is generated once the UPDATE message exchange has been completed. The message exchange is protected using the HIP association keys. Both HMAC and signing of packets is used.

9. IANA Considerations

This document defines additional parameters for the Host Identity Protocol [5]. These parameters are defined in Section 5.1.1 and Section 5.1.2 with the following numbers:

- ESP_INFO is 65.
- ESP_TRANSFORM is 4095.

10. References

10.1. Normative references

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", RFC 2404, November 1998.

[3] Frankel, S., Glenn, R., and S. Kelly, "The AES-CBC Cipher Algorithm and Its Use with IPsec", RFC 3602, September 2003.

[4] Kent, S., "IP Encapsulating Security Payload (ESP)", draft-ietf-ipsec-esp-v3-10 (work in progress), March 2005.

[5] Moskowitz, R., "Host Identity Protocol", draft-ietf-hip-base-04 (work in progress), October 2005.

[6] Schiller, J., "Cryptographic Algorithms for use in the Internet Key Exchange Version 2", draft-ietf-ipsec-ikev2-algorithms-05 (work in progress), April 2004.

[7] Moskowitz, R. and P. Nikander, "Host Identity Protocol Architecture", draft-ietf-hip-arch-03 (work in progress), August 2005.

[8] Schneier, B., "Applied Cryptography Second Edition: protocols algorithms and source in code in C", 1996.

10.2. Informative references

[9] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", draft-ietf-ipsec-rfc2401bis-06 (work in progress), April 2005.

[10] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[11] Melen, J. and P. Nikander, "A Bound End-to-End Tunnel (BEET) mode for ESP", draft-nikander-esp-beet-mode-04 (work in progress), November 2005.

[12] Nikander, P., "End-Host Mobility and Multihoming with the Host Identity Protocol", draft-ietf-hip-mm-02 (work in progress), July 2005.

Appendix A. A Note on Implementation Options

It is possible to implement this specification in multiple different ways. As noted above, one possible way of implementing is to rewrite IP headers below IPsec. In such an implementation, IPsec is used as if it was processing IPv6 transport mode packets, with the IPv6 header containing HITs instead of IP addresses in the source and destionation address fields. In outgoing packets, after IPsec processing, the HITs are replaced with actual IP addresses, based on the HITs and the SPI. In incoming packets, before IPsec processing, the IP addresses are replaced with HITs, based on the SPI in the incoming packet. In such an implementation, all IPsec policies are based on HITs and the upper layers only see packets with HITs in the place of IP addresses. Consequently, support of HIP does not conflict with other use of IPsec as long as the SPI spaces are kept separate.

Another way for implementing is to use the proposed BEET mode (A Bound End-to-End mode for ESP) [11]. The BEET mode provides some features from both IPsec tunnel and transport modes. The HIP uses HITs as the "inner" addresses and IP addresses as "outer" addresses like IP addresses are used in the tunnel mode. Instead of tunneling packets between hosts, a conversion between inner and outer addresses is made at end-hosts and the inner address is never sent in the wire after the initial HIP negotiation. BEET provides IPsec transport mode syntax (no inner headers) with limited tunnel mode semantics (fixed logical inner addresses - the HITs - and changeable outer IP addresses).

Compared to the option of implementing the required address rewrites outside of IPsec, BEET has one implementation level benefit. The BEET-way of implementing the address rewriting keeps all the configuration information in one place, at the SADB. On the other hand, when address rewriting is implemented separately, the implementation must make sure that the information in the SADB and the separate address rewriting DB are kept in synchrony. As a result, the BEET mode based way of implementing is RECOMMENDED over the separate implementation.

Authors' Addresses

Petri Jokela
Ericsson Research NomadicLab
JORVAS FIN-02420
FINLAND

Phone: +358 9 299 1
Email: petri.jokela@nomadiclab.com

Robert Moskowitz
ICSAlabs, a Division of TruSecure Corporation
1000 Bent Creek Blvd, Suite 200
Mechanicsburg, PA
USA

Email: rgm@icsalabs.com

Pekka Nikander
Ericsson Research NomadicLab
JORVAS FIN-02420
FINLAND

Phone: +358 9 299 1
Email: pekka.nikander@nomadiclab.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.