

IP Security Maintenance and Extensions
Internet-Draft
Updates: 9347 (if approved)
Intended status: Standards Track
Expires: 5 November 2026

A. Antony
S. Klassert
secunet
4 May 2026

Encrypted ESP Echo Protocol
draft-ietf-ipsecme-encrypted-esp-ping-03

Abstract

This document defines the Encrypted ESP Echo Function, a mechanism to assess the reachability of IP Security (IPsec) network paths using Encapsulating Security Payload (ESP) packets. It detects end-to-end path status by exchanging only encrypted ESP packets between IPsec peers. The Encrypted Echo message can either use existing congestion control payloads from RFC9347 or a new message format defined here, with an option to specify a preferred return path when there is more than one pair of IPsec SAs between the same set of IPsec peers.

A peer can announce support using a new IKEv2 Status Notification ENCRYPTED_PING_SUPPORTED.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 November 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	3
2. Requirements Language	3
3. Use cases	3
3.1. ESP Blocked or Filtered	4
3.2. Probing Multiple Paths	4
3.3. Probe Return Path	4
3.4. Manually Probing a Constant Rate on the AGGFRAG Tunnel	5
3.5. Why Not Use Existing IP Tools	5
3.6. Also Track Incoming Traffic for liveness check	5
4. Protocol Specification	6
4.1. Using Congestion Control Payload	6
4.2. Encrypted ESP Ping Payload Format	7
4.3. Return Path Validation	8
5. IKEv2 Notification	8
6. IANA Considerations	8
7. Operational Considerations	9
8. Acknowledgments	9
9. Security Considerations	9
10. Normative References	9
11. Informative References	10
Appendix A. Additional Stuff	10
Authors' Addresses	10

1. Introduction

In response to the operational need for a data-plane failure-detection mechanism for IP Security (IPsec) Encapsulating Security Payload (ESP) from [RFC4303], this document introduces Encrypted ESP Ping, including the Echo Request and Response. This protocol assesses network path reachability dynamically and can optionally specify a return path for echo Reply messages.

Encrypted ESP Ping supports two modes of operation. In the first mode, peers that have negotiated USE_AGGFRAG [RFC9347] can use the existing Congestion Control payload for echo requests and responses. In the second mode, peers negotiate ENCRYPTED_PING_SUPPORTED via IKEv2, which uses a new payload format defined in this document and commits the peer to responding.

The approach for specifying a preferred return path is inspired by Return Path Specified LSP Ping [RFC7110].

This document covers only Encrypted ESP Ping, typically used after an IKE negotiation, while [I-D.ietf-ipsecme-esp-ping] specifies an unauthenticated ESP Ping to be used before IKE negotiation.

This document updates [RFC9347] by defining new AGGFRAG_PAYLOAD subtypes for ESP Echo Request and Response, which may be used under ENCRYPTED_PING_SUPPORTED negotiation without requiring USE_AGGFRAG.

1.1. Terminology

This document uses the following terms defined in [RFC4301]: Encapsulating Security Payload (ESP), Security Association (SA), Security Policy Database (SPD).

This document uses the following terms defined in [RFC9347]: AGGFRAG tunnel.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Use cases

Diagnosing operational problems in IPsec can be challenging. Encrypted ESP Echo addresses these challenges; Encrypted ESP Ping is one such tool.

3.1. ESP Blocked or Filtered

An IPsec session typically employs ESP, using IP or IPv6 packets. ESP parameters are negotiated using IKEv2, with default IKEv2 messages exchanged over UDP port 500. In scenarios where ESP packets are not encapsulated in UDP (i.e., using the ESP protocol), successful IKE negotiation may occur, but ESP packets might fail to reach the peer due to differences in the packet path or filtering policies compared to IKE packets (e.g., UDP is allowed while ESP is filtered, typically due to misconfiguration). When using UDP encapsulation, ESP packets may encounter different filtering policies. This is typically due to broken packet filtering. Although this is less likely, it is still possible and can be difficult to diagnose. Operational experience suggests that networks and some home routers that drop ESP packets are common enough to cause problems for general-purpose VPN applications that require reliable performance on the Internet. Encrypted ESP Ping would greatly assist in diagnosing these scenarios. When IKEv2 operates over TCP to accommodate large post-quantum key exchanges while ESP may use a separate network path, Encrypted ESP Ping can verify ESP reachability independently of the IKEv2 control channel [I-D.ietf-ipsecme-ikev2-reliable-transport].

3.2. Probing Multiple Paths

When there are multiple paths created using multiple Child SAs with identical Traffic Selectors as specified in [RFC7296], there is a need to probe each Child SA, including the network path, independently from an IPsec peer. Each SA may traverse different network paths and may have different policies. Encrypted ESP Ping determines the reachability of each path independently.

3.3. Probe Return Path

IPsec Security Associations (SAs) are negotiated as a pair, consisting of two unidirectional SAs in one exchange. IKEv2 [RFC7296] Section 2.9 allows installing multiple Child SAs with identical Traffic Selectors. When there are multiple paths, the Encrypted ESP Ping SHOULD support requesting an echo response via a specific return path IPsec SA. To request a return path, additional attributes are necessary. The initiator would propose a specific SPI as the preferred return path. A specific return path SPI is necessary when to probe a specific path among multiple possible SAs between same peer. Multiple paths can exist for various reasons, such as a primary and secondary path scenario. For example over a satellite link and over fiber, the receiving peer may have a policy to respond via the fiber path even when the request arrives via the satellite link. If the initiator requests a return path, the

responder SHOULD try to respond via that IPsec SA. However, the final decision is up to the responder. If the responder decides to send the response via a different path than the requested return path, the initiator SHOULD detect this mismatch and report it, e.g., via a status code to the requesting, local, process.

3.4. Manually Probing a Constant Rate on the AGGFRAG Tunnel

AGGFRAG probes, whether sent at a constant or variable rate, can derive reachability information for the IPsec SA, similar to Encrypted ESP Ping. Since the probes are sent over an IPsec tunnel, the reachability information is reliable. IP Traffic Flow Security (IP-TFS) [RFC9347] uses AGGFRAG at a constant rate; an Encrypted ESP Ping application may also send AGGFRAG payloads for manual diagnostic purposes.

3.5. Why Not Use Existing IP Tools

Existing tools such as ICMP ping or traceroute assume IP connectivity. However, in IPsec gateway setups, the gateway itself may not have an IP address that matches the IPsec Security Policy Database (SPD). A peer MUST accept Encrypted ESP Ping messages even when it does not match a local SPD.

In the case of multiple SAs as mentioned above, IP tools would find it hard, if not impossible, to generate IP traffic to explore multiple paths specifically

3.6. Also Track Incoming Traffic for liveness check

Outgoing probes alone do not cover IPsec path health; incoming traffic must also be monitored. Incoming per-SA byte and packet counters maintained by IPsec are cryptographically verified: a counter can only increment if a valid authenticated ESP packet was received. These counters reliably indicate a viable path. This should be taken into account when probing IPsec paths. For example, when the crypto subsystem is overloaded, the responder may miss out on Encrypted ESP Ping responses. Tracking the incoming traffic after the ping probe is sent allows applications to confirm the IPsec path is still viable.

4. Protocol Specification

After an IPsec SA negotiation, [RFC7296], an IPsec peer wishing to verify the viability of the current network path for ESP packets MAY initiate an ESP Echo Request. The ESP Echo Request MUST be sent as an ESP packet using the established SA, receiving the same level of protection as any other traffic on that SA. It SHOULD use an SPI value previously established, whether negotiated through IKEv2 or configured manually.

The initiator sets the ESP Next Header field to AGGFRAG_PAYLOAD (144), as specified in [RFC9347]. The ESP payload contains one of the echo request sub-type payloads defined in this document, optionally followed by padding.

The receiving IPsec peer, having established ESP through IKE, MAY respond with an ESP Echo Response. The ESP Echo Response MUST be sent as an ESP packet using the corresponding SA and SPI. The responder also sets the ESP Next Header field to AGGFRAG_PAYLOAD (144), followed by the response sub-type payload.

Two payload formats are defined: the existing Congestion Control payload (Section 4.1) and the new Encrypted ESP Ping payload (Section 4.2).

4.1. Using Congestion Control Payload

IP-TFS Congestion Control AGGFRAG_PAYLOAD Payload Format as specified in [RFC9347] Section 6.1.2 can be used for Echo Request and response. When using this payload for Echo Request and response, IPv4 or IPv6 Data Block MUST NOT be concatenated, especially when USE_AGGFRAG is not successfully negotiated. This request does not support requesting a specific return path.

[AA when using USE_AGGFRAG tunnel is negotiated, responder may concatenate AGGFRAG_PAYLOAD Congestion control probe]


The Echo request and response payloads are not subject to IPsec Security Policy (SP), typically negotiated using IKEv2 and manually configured. End padding would be necessary if the tunnel is always sending fixed size ESP payload or to possibly detect path anomalies.

When probing do not take the lack of a response alone as an indication of the unreachability of the return path using ESP echo; also consider the received bytes on the return path. IPsec has a unique advantage over other tunneling protocols when the return path shows incoming bytes, indicating that the path is partially functional. This is especially useful when used as a liveness check

on busy paths. When there is no response, instead of concluding that the path is not viable and taking action, such as tearing down the IPsec connection, read the incoming bytes. This would help avoid tearing down busy paths due to the missing ESP echo response.

4.2. Encrypted ESP Ping Payload Format

Control Payload Format



- * Sub-Type: ESP-ECHO-REQUEST or ESP-ECHO-RESPONSE
- * Return path (R): 1 bit flag, set when requesting a specific return path SPI
- * Reserved: 7 bits
- * Data Length: number of data octets following, length 16 bits
- * Identifier : A 16-bit request identifier. The identifier SHOULD be set to a unique value to distinguish between different ESP Request sessions. The responder copies it from the request.
- * Sequence number: A 16-bit field that increments with each echo request sent.
- * Return path: 32 bits, optional requested return path SPI, when R is set to one.
- * Data : Optional data that follows the Echo request.

The responder SHOULD copy the request message and MUST change the Sub-type to ESP-ECHO-RESPONSE. The SHOULD allows a responder to send a minimal response with a shorter or empty data section when resources are constrained.

4.3. Return Path Validation

On the initiator, the return path SPI in the request MUST be in the local SADB with the same peer as the destination. The responder should also validate the requested return path SPI. When the SPI does not match the initiator in the SPD, the responder MUST NOT respond via the requested SPI. This is specifically to avoid amplification or DDoS. However, the responder MAY respond using the SPI of the SA on which the request was received.

5. IKEv2 Notification

A peer negotiates support for Encrypted ESP Ping Mode 2 using the Notification Status Type 'ENCRYPTED_PING_SUPPORTED' during the IKEv2 negotiation, in the IKE_AUTH exchange. When this has been negotiated, the peer has committed to responding to ESP Echo Requests on all ESP SAs between the two peers, and the initiator can reliably interpret a lack of response as path failure.

When ENCRYPTED_PING_SUPPORTED has been negotiated, the peer SHOULD respond to ESP Echo Requests. Responding remains subject to local resource constraints, but the negotiation represents a commitment to support this functionality. A lack of response after a configurable timeout MAY be treated as an indication of path failure.

For Mode 1 (USE_AGGFRAG), the commitment to respond is implicit in the USE_AGGFRAG negotiation itself.

6. IANA Considerations

This document defines two new registrations for the IANA ESP [AGGFRAG] PAYLOAD Sub-Types.

Value	ESP AGGFRAG_PAYLOAD Sub-Type	Reference
-----	-----	-----
[TBD2]	ESP-ECHO-REQUEST	[this document]
[TBD3]	ESP-ECHO-RESPONSE	[this document]

This document defines one new registration for the IANA "IKEv2 Notify Message Status Types" [STATUSNOTIFY].

Value	Notify Message Status Type	Reference
-----	-----	-----
[TBD1]	ENCRYPTED_PING_SUPPORTED.	[this document]

7. Operational Considerations

When an explicit return path is requested and the ESP Echo responder SHOULD make best effort to respond via this path, however, if local policies do not allow this respond via another SA.

A typical implementation involves creating an ESP Echo socket, which allows setting an outgoing SPI during initialization, and matching source and destination address. Once socket is setup before sending any data, only write payload with optionally specifying return path.

8. Acknowledgments

The authors thank Valery Smyslov for his thorough review and comments.

9. Security Considerations

The security considerations are similar to other unconnected request-reply protocols such as ICMP or ICMPv6 echo. The proposed ESP echo and response does not constitute an amplification attack because the ESP Echo Reply is almost same size as the ESP Echo Request. It can also be rate limited or filtered using ingress filtering per BCP 38 [RFC2827]

10. Normative References

- [AGGFRAG] IANA, "ESP AGGFRAG_PAYLOAD Registry",
<<https://www.iana.org/assignments/esp-aggfrag-payload/esp-aggfrag-payload.xhtml>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.

- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC9347] Hopps, C., "Aggregation and Fragmentation Mode for Encapsulating Security Payload (ESP) and Its Use for IP Traffic Flow Security (IP-TFS)", RFC 9347, DOI 10.17487/RFC9347, January 2023, <<https://www.rfc-editor.org/info/rfc9347>>.
- [STATUSNOTIFY] IANA, "IKEv2 Notify Message Status Types", <<https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#ikev2-parameters-16>>.

11. Informative References

- [I-D.ietf-ipsecme-esp-ping] Colitti, L., Linkova, J., and M. Richardson, "ESP Echo Protocol", Work in Progress, Internet-Draft, draft-ietf-ipsecme-esp-ping-01, 1 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-esp-ping-01>>.
- [I-D.ietf-ipsecme-ikev2-reliable-transport] Smyslov, V. and T. Reddy, "Separate Transports for IKE and ESP", Work in Progress, Internet-Draft, draft-ietf-ipsecme-ikev2-reliable-transport-02, 18 April 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-ikev2-reliable-transport-02>>.
- [RFC7110] Chen, M., Cao, W., Ning, S., Jounay, F., and S. Delord, "Return Path Specified Label Switched Path (LSP) Ping", RFC 7110, DOI 10.17487/RFC7110, January 2014, <<https://www.rfc-editor.org/info/rfc7110>>.

Appendix A. Additional Stuff

TBD

Authors' Addresses

Antony Antony
secunet Security Networks AG
Email: antony.antony@secunet.com

Steffen Klassert
secunet Security Networks AG
Email: steffen.klassert@secunet.com