

IPSECME Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

L. Xia
W. Jiang
Huawei Technologies
7 July 2025

Stateless Encryption Scheme of Enhanced Encapsulating Security Payload
(EESP)
draft-xia-ipsecme-eesp-stateless-encryption-01

Abstract

This draft first introduces several use cases for stateless encryption, analyzes and compares some existing stateless encryption schemes in the industry, and then attempts to propose a general and flexible stateless encryption scheme based on the summarized requirements.

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 8 January 2026.

Copyright Notice

Copyright (c) 2025 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
2. Use Cases	2
2.1. General Computing of Cloud Service	3
2.2. Cluster Communication in HPC Network	3
2.3. NIC/DPU Pool for General Computing	5
2.4. AI Computing	6
3. Requirement Summary	7
4. EESP Stateless Encryption Scheme	8
5. Security Considerations	9
6. IANA Considerations	9
7. Informative References	9
Authors' Addresses	9

1. Introduction

Recently, with the emergence of more new scenarios such as high-performance cloud services, AI large model computing, and 5G mobile backhaul networks, higher requirements have been put forward for the hardware friendliness, performance, and flexibility of the IPsec ESP protocol. A new protocol design, EESP [I-D.ietf-ipsecme-eesp] [I-D.ietf-ipsecme-eesp-ikev2], is being discussed and formulated. EESP focuses on solving issues such as introducing more fine-grained sub-child-SAs, adapting the ESP header and trailer format, and allowing parts of the transport layer header to be unencrypted, and implementing flexible expansion of EESP new features through options.

In addition to the issues listed above that are being addressed, stateless encryption is also a very important point. Its basic idea is to dynamically calculate data keys based on a small number of master keys (for AES-GCM, the encryption key and authentication key are combined), which helps optimize hardware resource limitations, performance optimization, and key negotiation complexity in large-scale IPsec session scenarios. This draft first introduces several use cases for stateless encryption, analyzes and compares some existing stateless encryption schemes in the industry, and then attempts to propose a general and flexible stateless encryption scheme based on the summarized requirements.

2. Use Cases

2.1. General Computing of Cloud Service

Public cloud services provide IPsec VPN access for massive users, and the servers in their infrastructure need to support massive IPsec session access. If hardware supports IPsec, the hardware should support session-based encryption and decryption, and the data keys of different sessions are isolated. The server needs to maintain the security connection context between the server and a large number of clients, and the hardware with limited memory cannot store the huge context. Note that the client and server do not belong to the same trusted domain in this case.

The stateless encryption scheme in the [PSP] solution proposed by Google is used to address the above hardware memory overhead problem. Its main principle is to derive a data key based on the master key on the server side, and the client side obtains the data key through an out-of-band method. It has:

- * Pros: Save half of total session contexts. Furthermore, since the master key is owned by server and not shared, key leakage affects only one server;
- * Cons: When a large number of new sessions are created, the data key negotiation is along the out of band slow path in real time, the first packet transmit will be delayed, and which results in performance degrade.

2.2. Cluster Communication in HPC Network

As shown in the below figure, encrypted communication is required between different instances of large-scale HPC jobs, the security session number is at the scale of $O(M * N * N)$. So, an efficient security context management mechanism is required to solve the problem of large-scale security sessions. Note that all communication instances of a HPC job belong to the same trusted domain.

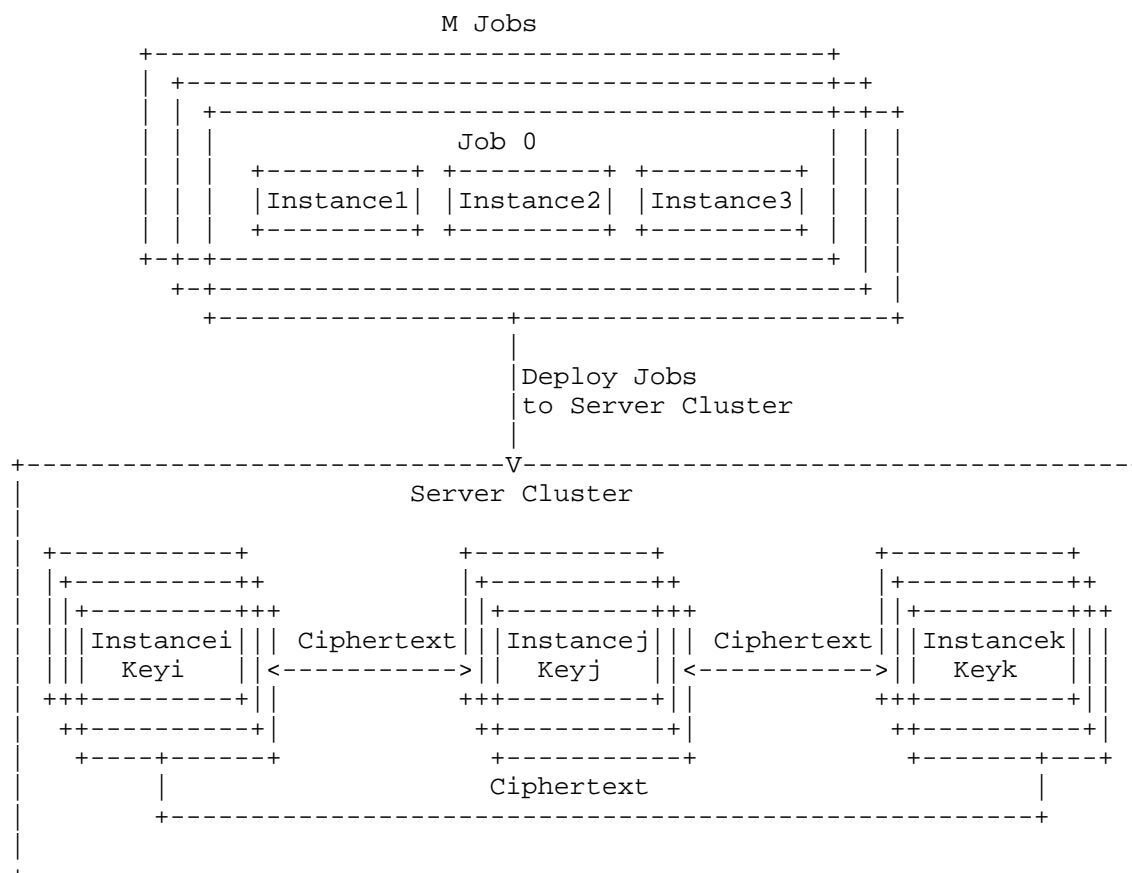


Figure 1: Encrypted Communication for Large Scale HPC Networks

The stateless encryption scheme defined by [UEC_TSS] can be used to solve the above problem. The main principle is that all communication instances of a HPC job belong to the same trust domain and share the same master key for both receiving and sending directions. It has:

* Pros:

- Better than Google PSP, it saves all security session contexts;
- The communication parties do not need to store data keys, and the increase of the number of instances and connections of the HPC job does not affect the number of security contexts;

- Without out of band slow path data key negotiation, the first packet delay is small;
- Data keys can be updated through the TSC.epoch.

* Cons:

- Master key leakage affects the entire trusted domain;
- The context content can be generated based on the SSI / Source IP / Destination IP field. Although the context content is flexible, the calculation overhead increases.

2.3. NIC/DPU Pool for General Computing

To cope with large-scale traffic access (e.g., computing server access to storage networks) and efficiently utilize network card resources, NIC resource pooling is an effective solution. For north-south traffic from client access to servers, the NIC resource pool must be transparent to the application, allowing a client to access resources behind any NIC in the pool. When using encrypted connections, all NICs must share the same key for a client's access. At this point, the NICs in a resource pool belong to the same trust domain, so stateless encryption sharing the master key is applicable. This saves data key synchronization between NICs and reduces the storage of security sessions and data keys on them in scenarios with a large number of secure client connections. The client obtains the data key for this encrypted connection through an out-of-band method, which can be derived from the master key and context. Encrypted connections and contexts can be isolated based on flows or VM instances. As shown in the figure below:

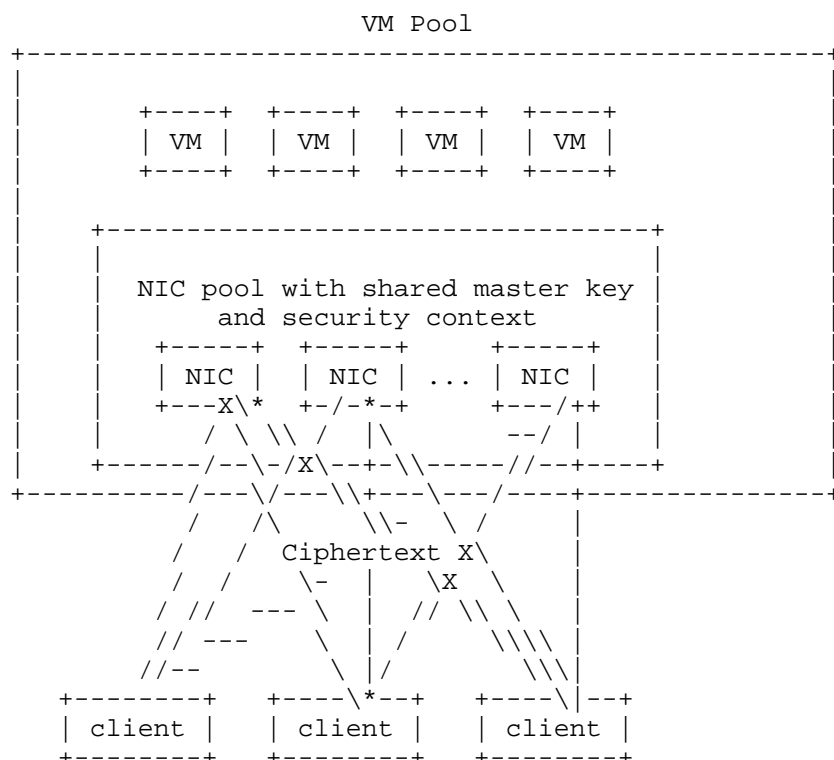


Figure 2: Encrypted Communication for NIC Pool

Similarly, the NIC resource pool can also be used for east-west traffic access between VMs. In this case, all NICs are in the same security domain and can share a master key, and different data keys can be dynamically generated based on different encryption connection contexts.

2.4. AI Computing

As shown in the figure below, in a AI computing network, a computing task is collaboratively executed by a group of CPUs & XPU's located in the same trust domain or across trust domains (in the case of cross-trust domains, they are interconnected as proxies through DPU). For CPUs & XPU's within the same trust domain, stateless encryption sharing the same master key can eliminate the complexity and latency of key negotiation between chips. For interconnection across trust domains, the DPU needs to perform encryption connection proxy functions between two trust domains (local trusted domain and global trusted domain). At this time, the DPU simultaneously possesses the master keys of the two trust domains, calculates the data key for

intra-domain communication in each domain based on its context, and then uses the calculated two data keys to complete the secure connection proxy across trust domains.

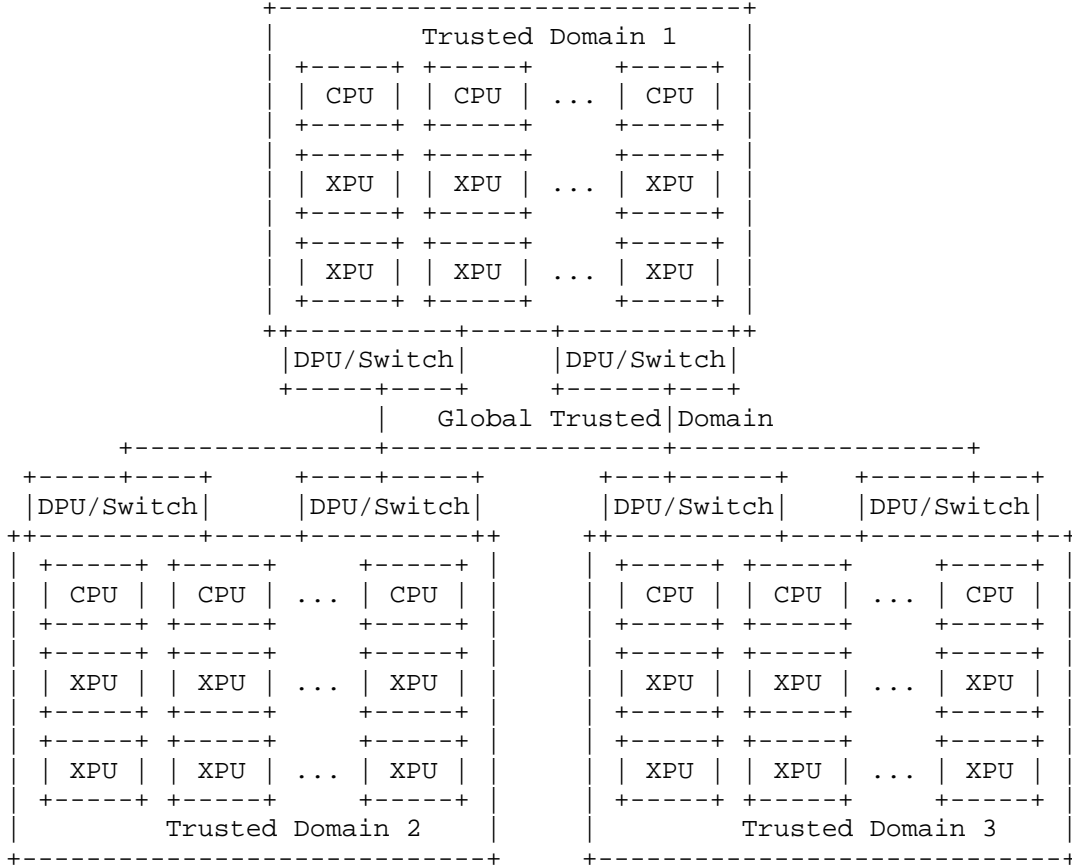


Figure 3: Encrypted Communication for AI Computing Network

3. Requirement Summary

Based on the above use cases, the requirements for a general and flexible stateless encryption scheme are as follows:

- * Support nodes within a trusted trust domain to share the same master key;
- * Master key supports multi-level combination design. In a trust domain, the master key is composed of multiple root keys of different types and levels, such as trust domain root key, tenant

root key, task group root key, etc. This enhances the overall security of the master key and supports fine-grained encryption traffic isolation (e.g., all nodes in a trust domain, nodes of the same tenant in a trust domain, nodes of the same computing task in a trust domain, etc.);

- * Different types of root keys have different security levels and lifecycles, and corresponding key rotation mechanisms need to be defined. The master key update will trigger the data key update;
- * The key rotation of each type of root key should support multiple key rotations, such as `pre_key`, `current_key`, and `next_key`, to support rapid rotation while ensuring that real-time encryption and decryption are not affected;
- * The key derivation of the data key is based on the master key, context, and KDF. KDF must support packet-by-packet data key calculation in most cases (except when the data key is cached in memory), which requires extremely high performance and must support cryptographically secure, hardware-concurrent high-performance algorithms;
- * To support real-time derivation of the Data Key, context information and IV information need to be carried with the message. To support different scenarios and different granularities of data key calculation and encryption traffic isolation (based on stream, based on source IP, based on source ID, etc.), multiple combinations of context and IV need to be supported, and different combination algorithms need to be distinguished through specific fields in the message;
- * Context information enables dynamic updates of the data key, such as carrying an epoch in the context. When the epoch changes, the data key is also refreshed accordingly;
- * It is necessary to support encryption proxy capabilities across trust domains. At the edge nodes across trust domains (such as DPU, Switch, etc.), support for master keys and stateless encryption of two trust domains (local trust domain and global trust domain) is required, and proxy conversion of message encryption and decryption between the two trust domains must be completed.

4. EESP Stateless Encryption Scheme

TBD.

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

7. Informative References

- [PSP] "PSP Architecture Specification", n.d.,
<https://github.com/google/psp/blob/main/doc/PSP_Arch_Spec.pdf>.
- [UEC_TSS] "Ultra Ethernet Specification v1.0", n.d.,
<<https://ultraethernet.org/wp-content/uploads/sites/20/2025/06/UE-Specification-6.11.25.pdf>>.
- [I-D.ietf-ipsecme-eesp]
Klassert, S., Antony, A., and C. Hopps, "Enhanced Encapsulating Security Payload (EESP)", Work in Progress, Internet-Draft, draft-ietf-ipsecme-eesp-01, 3 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-eesp-01>>.
- [I-D.ietf-ipsecme-eesp-ikev2]
Klassert, S., Antony, A., Brunner, T., and V. Smyslov, "IKEv2 negotiation for Enhanced Encapsulating Security Payload (EESP)", Work in Progress, Internet-Draft, draft-ietf-ipsecme-eesp-ikev2-00, 13 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-eesp-ikev2-00>>.

Authors' Addresses

Liang Xia
Huawei Technologies
Email: frank.xialiang@huawei.com

Weiyu Jiang
Huawei Technologies
Email: jiangweiyul@huawei.com