WG Working Group
Internet-Draft
Intended status: Informational
Expires: 17 September 2026

S. Salsano
Univ. of Rome Tor Vergata / CNIT
H. ElBakoury
Consultant
D. Lopez
Telefonica, I+D

16 March 2026

Extensible In-band Processing (EIP) Architecture and Framework
draft-eip-arch-09

Abstract

Extensible In-band Processing (EIP) extends the functionality of the IPv6 protocol considering the needs of future Internet services and advanced in-band metadata processing. This document discusses the architecture and framework of EIP. Separate documents respectively analyze a number of use cases for EIP, provide protocol specifications for EIP, and describe the integration of EIP within the IOAM framework through the Global Opaque Block (GOB) of the IOAM Pre-allocated Trace Option.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://eip-home.github.io/eip-arch/draft-eip-arch.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-eip-arch/.

Discussion of this document takes place on the EIP SIG mailing list (mailto:eip@cnit.it), which is archived at http://postino.cnit.it/cgi-bin/mailman/private/eip/.

Source for this draft and an issue tracker can be found at https://github.com/eip-home/eip-arch.

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/.

Copyright Notice

Table of Contents

1. Introduction

Networking architectures need to evolve to support the needs of future Internet services and 6G networks. The networking research and standardization communities have considered different approaches for this evolution, that can be broadly classified in 3 different categories:

1. Clean slate and "revolutionary" solutions. Throw away the legacy IP networking layer.

2. Solutions above Layer 3. Do not touch the legacy networking layer (IP).

3. Evolutionary solutions. Improve the IP layer (and try to preserve backward compatibility).

The proposed EIP (Extensible In-band Processing) solution belongs to the third category; it extends the current IPv6 architecture without requiring a clean-slate revolution.

The use cases for EIP are discussed in [id-eip-use-cases]. The specification of the EIP header format is provided in [id-eip-headers].

2. Basic principles for EIP

An ongoing trend is extending the functionality of the IPv6 networking layer, going beyond the plain packet forwarding. An example of this trend is the rise of the SRv6 "network programming" model. With the SRv6 network programming model, the routers can implement "complex" functionalities and they can be controlled by a "network program" that is embedded in IPv6 packet headers. Another example is the INT (IN band Telemetry) solution for monitoring. These (and other) examples are further discussed in Section 4.

The EIP solution is aligned with this trend, which will ensure a future proof evolution of networking architectures. EIP supports a feature-rich and extensible IPv6 networking layer, in which complex dataplane functions can be executed by end-hosts, routers, virtual functions, servers in datacenters so that services can be implemented in the smartest and most efficient way.

The EIP solution introduces an EIP header in the IPv6 packet header. The proposed EIP header is extensible and it is meant to support a number of different use cases. In general, both end-hosts and transit routers can read and write the content of this header. Depending on the specific use case, only specific nodes will be capable and interested in reading or writing the EIP header. The use of the EIP header can be confined to a single domain or to a set of cooperating domains, so there is no need of a global, Internet-wide support of the new header for its introduction. Moreover, there can be usage scenarios in which legacy nodes can simply ignore the EIP header and provide transit to packets containing the EIP header.

The EIP header could be carried in different ways inside the IPv6 header: 1) as an EIP Option in the Hop-by-Hop Extension Header; 2) as an EIP TLV in the Segment Routing Header; 3) within the IOAM framework through the Global Opaque Block (GOB) of the IOAM Pre-allocated Trace Option, as specified in [id-gob-ioam] and discussed in Section 5.

An important usage scenario considers the transport of user packets over a provider network. In this scenario, we consider the network portion from the provider ingress edge node to the provider egress edge node. The ingress edge node can encapsulate the user packet coming from an access network into an outer packet. The outer packet travels in the provider network until the egress edge node, which will decapsulate the inner packet and deliver it to the destination access network or to another transit network, depending on the specific topology and service. Assuming that the IPv6/SRv6 dataplane is used in the provider network, the ingress edge node will be the source of an outer IPv6 packet in which it is possible to add the EIP header. The outer IPv6 packet (containing the EIP header) will be processed inside the "limited domain" (see [RFC8799]) of the provider network, so that the operator can make sure that all the transit routers either are EIP aware or at least they can forward packets containing the EIP header. In this usage scenario, the EIP framework operates "edge-to-edge" and the end-user packets are "tunneled" over the EIP domain.

The architectural framework for EIP is depicted in Figure 1. We refer to nodes that are not EIP capable as legacy nodes. An EIP domain is made up by EIP aware routers (EIP R) and can also include legacy routers (LEG R). At the border of the EIP domain, EIP edge nodes (EIP ER) are used to interact with legacy End Hosts / Servers (LEG H) and with other domains. It is also possible that an End Host / Server is EIP aware (EIP H), in this case the EIP framework could operate "edge-to-end" or "end-to-end".

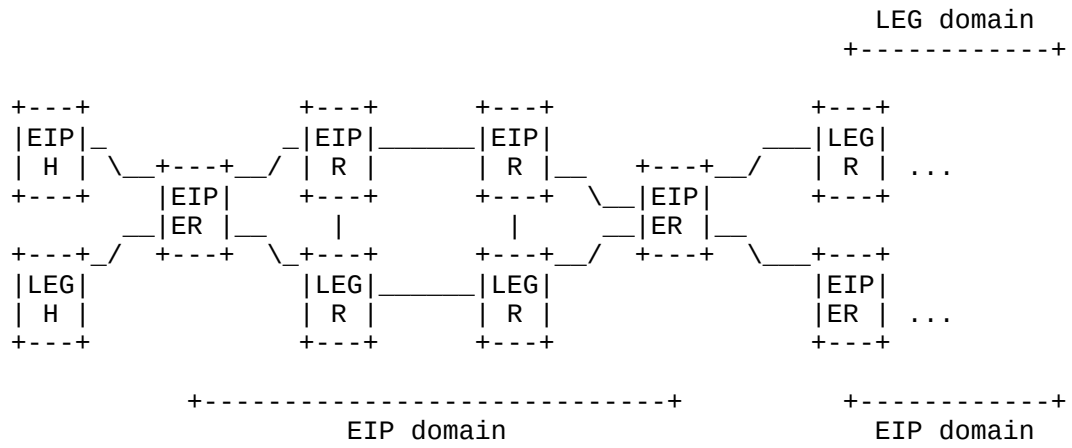

Figure 1: EIP framework

As shown in Figure 1, an EIP domain can communicate with other domains, which can be legacy domains or EIP capable domains.

3. Benefits of a common EIP header for multiple use cases

The EIP header will carry different EIP Information Elements that are defined to support the different use cases. There are reasons why it is beneficial to define a single common EIP header that supports multiple use cases using the EIP Information Elements.

1. The number of available Option Types in HBH header is limited (see Section 3.1). Likewise the number of available TLVs in the Segment Routing Header (SRH) and the number of IOAM-Data-Field-Type are limited. Defining multiple Option Types (or SRH TLVs or IOAM-Data-Field-Type) for multiple use cases is not scalable and puts pressure on the allocation of such codepoints.

2. The definition and standardization of specific EIP Information Elements for the different use cases will be simplified, compared to the need of requiring the definition of a new Option Type or SRH TLVs or IOAM-Data-Field-Type.

3. Different use cases may share a subset of common EIP Information Elements.

4. Efficient mechanisms for the processing of the EIP header (both in software and in hardware) can be defined when the different EIP Information Elements are carried inside the same EIP header.

3.1. Considerations on Hop-by-Hop Options allocation

Several proposals and already standardized solutions use the IPv6 Hop-by-Hop Options, as discussed below in Section 4. The Hop-by-Hop Options are represented with an 8-bit code. The first two bits represent the action to be taken if the option is unknown to a node that receives it, while the third bit specifies whether the content of the option can be changed in flight. In particular, Option Types that start with 000 should be ignored if unknown and cannot be changed in flight, whereas Option Types that start with 001 should be ignored if unknown and can be changed in flight. The current IANA allocation for Option Types starting with 000 and 001 is as follows (see [IANA-ipv6-parameters]):

‘‘‘ 32 possible Option Types starting with 000 7 assigned values in the current IANA registry (including IOAM and AltMark) 25 unassigned

32 possible Option Types starting with 001 6 assigned values in the current IANA registry (including IOAM) 26 unassigned ‘‘‘

We observe that there is a potential scarcity of the code points, as there are many scenarios that could require the definition of a new Hop-by-Hop option. We also observe that having only 1 code point allocated for experiments is a very restrictive limitation.

4. Review of recent activities that propose to extend the IP networking layer

4.1. Standardized and proposed evolutions of IPv6

In the last few years, we have witnessed important innovations in IPv6 networking, centered around the emergence of Segment Routing for IPv6 (SRv6) [RFC8754] and of the SRv6 "Network Programming model" [RFC8986]. With SRv6 it is possible to insert a _Network program_, i.e. a sequence of instructions (called _segments_), in a header of the IPv6 protocol, called Segment Routing Header (SRH). Recent updates (see [RFC9800]) introduced compression mechanisms for segment lists, improving scalability for long segment chains.

Another recent activity that proposed to extend the networking layer to support more complex functions concerns network monitoring. The concept of INT "In-band Network Telemetry" has been proposed since 2015 [onf-int] in the context of the definition of use cases for P4 based data plane programmability. The latest version of INT specifications dates November 2020 [int-spec]. [int-spec] specifies the format of headers that carry monitoring instructions and monitoring information along with data plane packets. The specific location for INT Headers is intentionally not specified: an INT

Header can be inserted as an option or payload of any encapsulation type. The In-band Telemetry concept has been adopted by the IPPM IETF Working Group, renaming it "In-situ Operations, Administration, and Maintenance" (IOAM). [RFC9197] discusses the data fields and associated data types for IOAM. The in-situ OAM data fields can be encapsulated in a variety of protocols, including IPv6. The specification details for carrying IOAM data inside IPv6 headers are provided in [RFC9486]. In particular, IOAM data fields can be encapsulated in IPv6 using either Hop-by-Hop Options header or Destination options header. A Direct Export variant has been defined in [RFC9326], enabling nodes to export telemetry data directly without per-hop accumulation.

Another example of extensions to IPv6 for network monitoring is specified in [RFC8250], which defines an IPv6 Destination Options header called Performance and Diagnostic Metrics (PDM). The PDM option header provides sequence numbers and timing information as a basis for measurements.

The "Alternate Marking Method" is a recently proposed performance measurement approach described in [RFC9341]. [RFC9343] defines a new Hop-by-Hop Option to support this approach.

"Path Tracing" [I-D.draft-filsfils-ippm-path-tracing] proposes an efficient solution for recording the route taken by a packet (including timestamps and load information taken at each hop along the route). This solution needs a new Hop-by-Hop Option to be defined. A new lightweight telemetry mechanism has been proposed in [I-D.draft-mzbc-ippm-transit-measurement-option], which accumulates end-to-end delay and congestion flags in a fixed-size structure.

[RFC8558] analyses the evolution of transport protocols. It recommends that explicit signals should be used when the endpoints desire that network elements along the path become aware of events related to transport protocol. Among the solutions, [RFC8558] considers the use of explicit signals at the network layer, and in particular it mentions that IPv6 hop-by-hop headers might suit this purpose.

[RFC9268] specifies a new IPv6 Hop-by-Hop option that is used to record the minimum Path MTU between a source and a destination.

[RFC9837] describes an experiment in which VPN service information for both Layer 2 and Layer 3 VPNs is encoded in an IPv6 Destination Option. This experimental IPv6 Destination Option is called the VPN Service Option.

The Internet-Draft [I-D.draft-ietf-6man-enhanced-vpn-vtn-id] proposes a new Hop-by-Hop option of IPv6 extension header to carry the Network Resource Partition (NRP) information, which could be used to identify the NRP-specific processing to be performed on the packets by each network node along a network path in the NRP.

The Internet-Draft [I-D.draft-guan-6man-ipv6-id-authentication] proposes an IPv6 based address label terminal identity authentication mechanism, which uses a new Hop-by-Hop option, called Address Label Extension (ALE).

The Internet-Draft [I-D.draft-herbert-fast] (currently expired) proposed the Firewalls and Service Tickets (FAST) protocol. This is a generic and extensible protocol for hosts to signal network nodes to request services or to gain admission into a network. Tickets are sent in IPv6 Hop-by-Hop options.

The Internet-Draft [I-D.draft-eckert-6man-qos-exthdr-discuss] (currently expired) provided considerations for common QoS IPv6 extension header, in the context of the functionality under definition in the Deterministic Networking (detnet) IETF Working Group [detnet-wg].

The Internet-Draft [I-D.draft-li-6man-topology-id] (currently expired) proposed a new Hop-by-Hop option of IPv6 extension header to carry the topology identifier, which is used to identify the forwarding table instance created by the Multi Topology Routing or Flexible Algorithm.

The Internet-Draft [I-D.draft-iurman-6man-carry-identifier] (currently expired) discussed the need of having a generic approach for carrying identifiers in IPv6 Destination Options and Hop-by-Hop Options. The EIP proposal can be seen as a superset and a further generalization of the proposal of [I-D.draft-iurman-6man-carry-identifier].

4.2. Additional relevant activities

The IETF has shown interest in carrying application or service-level metadata in IPv6. The Application-aware Networking (APN) BoF discussed embedding such metadata, leading to proposals like APN6. The recently chartered CATS (Compute-Aware Traffic Steering) WG explores approaches where traffic is steered based on in-packet compute-related information. The GREEN WG (Getting Ready for Energy-Efficient Networking), formed in 2024, investigates telemetry for carbon-aware routing. The COIN IRTF RG has discussed in-network processing requirements that also point to in-band metadata handling.

The FANTEL BoF (IETF 123, Madrid, 2025) discussed the Fast Notification for Traffic Engineering and Load Balancing framework [ietf-fantel]. FANTEL proposes in-band mechanisms to signal network conditions such as congestion or link degradation using IPv6 packets. These notifications are inserted by routers to support real-time traffic steering decisions. The goals of FANTEL align with the EIP approach, which provides an extensible container for in-band metadata through EIP Information Elements. The EIP header could encapsulate FANTEL notifications without requiring additional Hop-by-Hop Option codepoints, supporting both domain-specific and broader deployments.

Outside the IETF, the P4.org community continues its efforts on programmable dataplanes and has proposed updated INT mechanisms. Recent research includes the use of in-band headers for on-path inference and service-specific packet handling, showing increasing interest in general, extensible frameworks like EIP.

5. Integration of EIP into the IOAM Framework

The IOAM (In-situ Operations, Administration, and Maintenance) framework [RFC9197] defines a set of data fields and associated semantics for recording telemetry and operational information within packets as they traverse a network. The IOAM data can be encapsulated in IPv6 via Hop-by-Hop or Destination Options headers, as specified in [RFC9486], and can be processed by IOAM-capable nodes along the path.

An earlier integration direction for EIP within IOAM was explored in [salsano25-eipioam], where EIP was modeled as a new IOAM Data-Field-Type carried within the existing IOAM data-field structure. That approach showed that EIP Information Elements could be embedded into the IOAM processing pipeline, enabling reuse of the existing IOAM encapsulation and processing model without introducing separate Hop-by-Hop options.

The current integration direction adopts a different approach based on the Global Opaque Block (GOB) of the IOAM Pre-allocated Trace Option, as specified in [id-gob-ioam]. In this model, EIP Information Elements can be carried inside a reusable, pre-allocated global metadata region that is distinct from the per-node Trace data and can support schema-driven formats and controlled in-band updates.

Compared to the earlier Data-Field-Type approach, the GOB-based integration provides a more general and implementation-friendly solution for structured global metadata within IOAM. It preserves compatibility with the IOAM encapsulation model while avoiding the need to map EIP semantics onto the existing per-node data-field abstraction.

The IOAM-based integration of EIP is not mutually exclusive with the standalone deployment of EIP as an independent IPv6 extension header. However, for integration with IOAM, the GOB-based approach is considered the preferred solution.

6. Conventions and Definitions

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.

7. Security Considerations

EIP introduces in-band metadata that may be read or modified by on-path nodes. Unauthorized access or modification can affect telemetry, service behavior, or policy enforcement. Therefore, deployments should be limited to controlled domains and should rely on existing IPv6/IOAM security mechanisms and domain trust assumptions.

8. IANA Considerations

The definition of the EIP header as an Option for the IPv6 Hop-by-Hop Extension header requires the allocation of a codepoint from the "Destination Options and Hop-by-Hop Options" registry in the "Internet Protocol Version 6 (IPv6) Parameters" [IANA-ipv6-parameters].

The definition of the EIP header as a TLV in the Segment Routing Header requires the allocation of a codepoint from the "Segment Routing Header TLVs" registry in the "Internet Protocol Version 6 (IPv6) Parameters" [IANA-ipv6-parameters].

The definition of EIP Information Elements in the EIP header will require the creation of a new IANA registry to manage EIP Information Element type values.

An earlier integration of EIP into IOAM as a new Data-Field-Type was explored in [salsano25-eipioam], which would have required an allocation from the "IOAM Data Field Types" registry [IANA-ioam-types]. The currently preferred IOAM integration for EIP is instead based on the Global Opaque Block (GOB), whose protocol format and any related codepoint requirements are specified in [id-gob-ioam].

9. References

9.1. Normative References

[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/rfc/rfc2119>.

[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/rfc/rfc8174>.

9.2. Informative References

[detnet-wg]
IETF, "Deterministic Networking (DetNet) IETF Working Group", 2025, <https://datatracker.ietf.org/wg/detnet/about/>.

[I-D.draft-eckert-6man-qos-exthdr-discuss]
Eckert, T. T., Joung, J., Peng, S., and X. Geng, "Considerations for common QoS IPv6 extension header(s)", Work in Progress, Internet-Draft, draft-eckert-6man-qos-exthdr-discuss-00, 4 March 2024, <https://datatracker.ietf.org/doc/html/draft-eckert-6man-qos-exthdr-discuss-00>.

[I-D.draft-filsfils-ippm-path-tracing]
Filsfils, C., Abdelsalam, A., Camarillo, P., Yufit, M., Su, Y., Matsushima, S., Valentine, M., and Dhamija, "Path Tracing in SRv6 networks", Work in Progress, Internet-Draft, draft-filsfils-ippm-path-tracing-05, 4 January 2026, <https://datatracker.ietf.org/doc/html/draft-filsfils-ippm-path-tracing-05>.

[I-D.draft-guan-6man-ipv6-id-authentication]
Guan, J., Yao, S., Liu, K., Hu, X., and J. Liu, "Terminal Identity Authentication Based on Address Label", Work in Progress, Internet-Draft, draft-guan-6man-ipv6-id-authentication-04, 18 January 2026, <https://datatracker.ietf.org/doc/html/draft-guan-6man-ipv6-id-authentication-04>.

[I-D.draft-herbert-fast]
Herbert, T., "Firewall and Service Tickets", Work in Progress, Internet-Draft, draft-herbert-fast-07, 7 October 2023, <https://datatracker.ietf.org/doc/html/draft-herbert-fast-07>.

[I-D.draft-ietf-6man-enhanced-vpn-vtn-id]
Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-14, 11 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-14>.

[I-D.draft-iurman-6man-carry-identifier]
Iurman, J., "Carrying an Identifier in IPv6 packets", Work in Progress, Internet-Draft, draft-iurman-6man-carry-identifier-00, 8 February 2023, <https://datatracker.ietf.org/doc/html/draft-iurman-6man-carry-identifier-00>.

[I-D.draft-li-6man-topology-id]
Li, Z., Hu, Z., and J. Dong, "Topology Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-li-6man-topology-id-00, 20 March 2022, <https://datatracker.ietf.org/doc/html/draft-li-6man-topology-id-00>.

[I-D.draft-mzbc-ippm-transit-measurement-option]
Mizrahi, T., Zhou, T., Belkar, S., and R. Cohen, "The Transit Measurement Option", Work in Progress, Internet-Draft, draft-mzbc-ippm-transit-measurement-option-07, 5 January 2026, <https://datatracker.ietf.org/doc/html/draft-mzbc-ippm-transit-measurement-option-07>.

[IANA-ioam-types]
IANA, "IOAM Data Field Types", n.d., <https://www.iana.org/assignments/ioam/ioam.xhtml#data-field-types>.

[IANA-ipv6-parameters]
IANA, "Internet Protocol Version 6 (IPv6) Parameters", n.d., <https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml>.

[id-eip-headers]
Salsano, S. and H. ElBakoury, "Extensible In-band Processing (EIP) Headers Definitions", 2022, <https://eip-home.github.io/eip-headers/draft-eip-headers-definitions.txt>.

[id-eip-use-cases]
Salsano, S. and H. ElBakoury, "Extensible In-band Processing (EIP) Use Cases", 2022, <https://eip-home.github.io/use-cases/draft-eip-use-cases.txt>.

[id-gob-ioam]
Mayer, A. and S. Salsano, "Global Opaque Block for IOAM Pre-allocated Trace Option", Internet-Draft draft-mayer-ioam-gob, 2026, <https://datatracker.ietf.org/doc/draft-mayer-ioam-gob/>.

[ietf-fantel]
IETF, "Fast Notification for Traffic Engineering and Load Balancing (FANTEL) BoF", 2025, <https://datatracker.ietf.org/meeting/123/materials/bofdraft-fantel-00>.

[int-spec] The P4 org Applications Working Group, "In-band Network Telemetry (INT) Dataplane Specification, version 2.1", 2022, <https://p4.org/p4-spec/docs/INT v2 1.pdf>.

[onf-int] P4 org, "Improving Network Monitoring and Management with Programmable Data Planes", 2015, <https://opennetworking.org/news-and-events/blog/improving-network-monitoring-and-management-with-programmable-data-planes/>.

[RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <https://www.rfc-editor.org/rfc/rfc8250>.

[RFC8558] Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <https://www.rfc-editor.org/rfc/rfc8558>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/rfc/rfc8754>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/rfc/rfc8799>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/rfc/rfc8986>.

[RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <https://www.rfc-editor.org/rfc/rfc9197>.

[RFC9268] Hinden, R. and G. Fairhurst, "IPv6 Minimum Path MTU Hop-by-Hop Option", RFC 9268, DOI 10.17487/RFC9268, August 2022, <https://www.rfc-editor.org/rfc/rfc9268>.

[RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <https://www.rfc-editor.org/rfc/rfc9326>.

[RFC9341] Fioccola, G., Ed., Cociglio, M., Mirsky, G., Mizrahi, T., and T. Zhou, "Alternate-Marking Method", RFC 9341, DOI 10.17487/RFC9341, December 2022, <https://www.rfc-editor.org/rfc/rfc9341>.

[RFC9343] Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate-Marking Method", RFC 9343, DOI 10.17487/RFC9343, December 2022, <https://www.rfc-editor.org/rfc/rfc9343>.

[RFC9486] Bhandari, S., Ed. and F. Brockners, Ed., "IPv6 Options for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9486, DOI 10.17487/RFC9486, September 2023, <https://www.rfc-editor.org/rfc/rfc9486>.

[RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025, <https://www.rfc-editor.org/rfc/rfc9800>.

[RFC9837] Bonica, R., Li, X., Farrel, A., Kamite, Y., and L. Jalil, "The IPv6 VPN Service Destination Option", RFC 9837, DOI 10.17487/RFC9837, August 2025, <https://www.rfc-editor.org/rfc/rfc9837>.

[salsano25-eipioam]
Salsano, S., Mayer, A., Sidoretti, G., Loreti, P., Bracciale, L., ElBakoury, H., and D. Lopez, "Integrating Extensible In-Band Processing (EIP) into the IOAM Framework: A Unified Approach to In-Packet Telemetry and Metadata", IEEE CSCN 2025 Conference, 2025.

Acknowledgments

TODO acknowledgements.

Authors' Addresses

Stefano Salsano
Univ. of Rome Tor Vergata / CNIT
Email: stefano.salsano@uniroma2.it

Hesham ElBakoury
Consultant
Email: helbakoury@gmail.com

Diego R. Lopez
Telefonica, I+D
Email: diego.r.lopez@telefonica.com