

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 11 August 2024

G. Fioccola
T. Zhou
Huawei
M. Cociglio
Telecom Italia
G. Mishra
Verizon Inc.
X. Wang
Ruijie
G. Zhang
China Mobile
8 February 2024

Application of the Alternate Marking Method to the Segment Routing
Header
draft-fz-spring-srv6-alt-mark-08

Abstract

The Alternate Marking Method is a passive performance measurement method based on marking consecutive batches of packets, which can be used to measure packet loss, latency, and jitter of live traffic. This method requires a packet marking method so that packet flows can be distinguished and identified.

A mechanism to carry suitable packet marking in the Hop-by-Hop Header and the Destination Options Header of an IPv6 packet is described in RFC 9343 and is also applicable to Segment Routing for IPv6 (SRv6).

This document describes an alternative approach that uses a new TLV in the Segment Routing Header (SRH) of an SRv6 packet. This approach has been implemented and has potential scaling and simplification benefits over the technique described in RFC 9343.

This protocol extension has been developed outside the IETF and is published here to guide implementation, ensure interoperability among implementations, and enable wide-scale deployment to determine the potential benefits of this approach.

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 11 August 2024.

Copyright Notice

Copyright (c) 2024 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.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Application of the Alternate Marking to SRv6	4
2.1. Controlled Domain	4
3. Definition of the SRH AltMark TLV	5
3.1. The Flow Monitoring Identification (FlowMonID)	6
3.2. Optional Extended Data Fields for Enhanced Alternate Marking	7
4. Use of the SRH AltMark TLV	10
5. Observations on RFC 9343	11
6. SRH AltMark TLV Compatibility	13
7. Implementation Overview	14
8. Security Considerations	15
9. IANA Considerations	15
10. Acknowledgements	15
11. Contributors	15
12. References	16
12.1. Normative References	16
12.2. Informative References	16
Authors' Addresses	18

1. Introduction

[RFC9341] and [RFC9342] describe a passive performance measurement method, which can be used to measure packet loss, latency and jitter on live traffic. Since this method is based on marking consecutive batches of packets, the method is often referred as the Alternate Marking Method.

The Alternate Marking Method requires a marking field so that packet flows can be distinguished and identified. An IETF standards track solution is described in [RFC9343] which analyzes the possible implementation options for the application of the Alternate Marking Method in an IPv6 domain and defines how the marking field can be encoded in a new TLV that is carried in the Option Headers (both Hop-by-hop or Destination) of IPv6 packets for to achieve Alternate Marking in an IPv6 domain. That solution is equally applicable to Segment Routing for IPv6 (SRv6) networks [RFC8402].

This document describes an alternative approach that encodes the marking field in a new TLV carried in the Segment Routing Header (SRH) [RFC8754] of an SRv6 packet. This approach is specific to SRv6 networks and does not apply in native IPv6 networks, but it has potential scaling and simplification benefits over the technique described in [RFC9343]. Indeed, the rationale is to place an information related to an SRv6 path directly inside the SRH. It has been implemented taking into account that SR nodes are supposed to support fast parsing and processing of the SRH, while the SR nodes may not handle properly Destination Options, as described in [RFC9098] and [I-D.ietf-6man-eh-limits].

This protocol extension has been developed outside the IETF and is published here to guide implementation, ensure interoperability among implementations, and enable wide-scale deployment. See Section 7 for more details .

1.1. 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.

2. Application of the Alternate Marking to SRv6

[RFC9341] and [RFC9342] defines the Alternate Marking Method which employs one or two marking bits inside the packet header to perform network measurements. The Flow Monitoring Identification field (FlowMonID), as introduced in Section 3, is used to identify monitored flows and aids the optimization of implementation and scaling of the Alternate Marking Method.

Note that the Flow Label field of the IPv6 Header [RFC8200] is also used to identify packet flows, but its usage is somewhat different from the FlowMonID. While the FlowMonID is used to identify the monitored flow, the Flow Label is used for application services, such as load-balancing, equal cost multi-path (ECMP), and QoS. Reusing the Flow Label field for identifying monitored flows is ruled out because setting it for flow monitoring purposes might cause changes to the application intent and the forwarding behaviour. Conversely, the Flow Label might be changed by normal processing along the packet's path, and this would break the measurement/monitoring task.

An important point that will also be discussed in this document is the uniqueness of the FlowMonID and how to allow disambiguation of the FlowMonID in case of collision.

Section 2.1 highlights an important requirement for the application of the Alternate Marking to IPv6 and SRv6. The concept of the Controlled Domain is explained as an essential precondition.

2.1. Controlled Domain

[RFC8799] introduces the concept of specific limited domain solutions and notes application of the Alternate Marking Method as an example.

Despite the flexibility of IPv6, when innovative applications are proposed they are often applied within controlled domains to help constrain the domain-wide policies, options supported, the style of network management, and security requirements. This is also the case for the application of the Alternate Marking Method to SRv6.

Therefore, the application of the Alternate Marking Method to SRv6 MUST be deployed only within a controlled domain. Implementations MUST reject or discard packets that carry Alternate Marking data (using the new SRH TLV) that attempt to enter the controlled domain, and SHOULD prevent packets carrying Alternate Marking data from leaving the controlled domains.

For SRv6, the controlled domain corresponds to an SR domain, as defined in [RFC8402]. [RFC9343] introduces the Alternate-Marking measurement domain that can overlap with the controlled domain or may be a subset of the controlled domain. Therefore, it is also possible to enter the controlled domain with an SRH already in place and add the Alternate Marking data to the SRH.

3. Definition of the SRH AltMark TLV

The AltMark SRH TLV is defined to carry the data fields associated with the Alternate Marking Method. The TLV has some initial fields that are always present, and further extension fields that are present when Enhanced Alternate Marking is in use.

Figure 1 shows the format of the AltMark TLV.

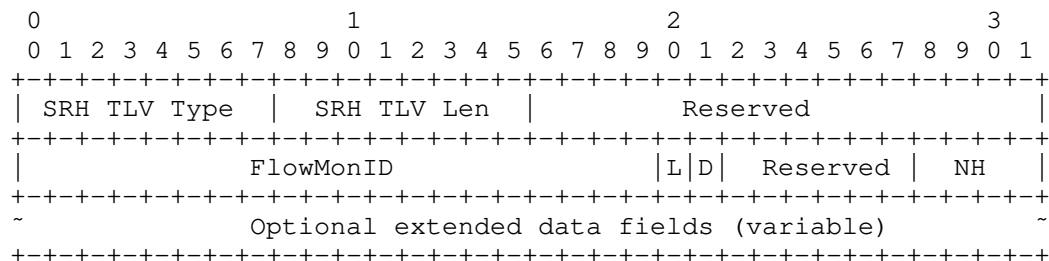


Figure 1: AltMark: SRH TLV for alternate marking

The fields of this TLV are as follows:

- * **SRH TLV Type:** 8 bit identifier of the Alternate Marking SRH TLV. The value for this field is taken from the range 124-126. That is, it is an Experimental code point that indicates a TLV that does not change en route. Deployments of implementations of this document must coordinate the value used by all implementations participating in the deployment. Thus, implementations **MUST** make this value configurable. Further, deployments must carefully consider any other implementations running in the network to avoid clashes with other SRH TLVs.
- * **SRH TLV Len:** The length of the Data Fields of this TLV in bytes. This is set to 6 when Enhanced Alternate Marking is not in use.
- * **FlowMonID:** 20 bits unsigned integer. The FlowMon identifier is described in Section 3.1.

- * **Reserved:** Reserved for future use. These bits MUST be set to zero on transmission and ignored on receipt.
- * **L:** Loss flag as defined in [RFC9343].
- * **D:** Delay flag as defined in [RFC9343].
- * **NH:** The NH (NextHeader) field is used to indicate extended data fields are present to support Enhanced Alternate Marking as follows:
 - NextHeader value of 0x0 means that there is no extended data field attached.
 - NextHeader values of 0x1-0x8 are reserved for further usage.
 - NextHeader value of 0x9 indicates the extended data fields are present as described in Section 3.2.
 - NextHeader values of 0xA-0xF are reserved for further usage.
- * Optional extended data fields may be present according to the setting of the NH field and as described in Section 3.2.

3.1. The Flow Monitoring Identification (FlowMonID)

The Flow Monitoring Identification (FlowMonID) is required for three reasons:

First, it helps to reduce the per node configuration. Otherwise, each node needs to configure an access-control list (ACL) for each of the monitored flows. Moreover, using a flow identifier allows a flexible granularity for the flow definition.

Second, it simplifies the handling of counters. Hardware processing of flow tuples (and ACL matching) is challenging and often incurs into performance issues, especially on tunnel interfaces.

Third, it eases the data export encapsulation and correlation for the collectors.

The FlowMonID field is used to uniquely identify a monitored flow within the controlled measurement domain. The field is set at the entry node to the domain. The FlowMonID can be assigned by a central controller or algorithmically generated by the domain entry node. The latter approach cannot guarantee the uniqueness of FlowMonID, but it may be preferred for a network where the conflict probability is small due to the large FlowMonID space.

It is important to note that if the 20 bit FlowMonID is set by the domain entry nodes, there is a chance of collision even when the values are chosen using a pseudo-random algorithm. In these cases a value may be not be sufficient to uniquely identify a monitored flow. In such cases the packets need to be tagged with additional flow information to allow disambiguation. Such additional tagging is carried in the extended data fields described in Section 3.2.

3.2. Optional Extended Data Fields for Enhanced Alternate Marking

The optional extended data fields to support Enhanced Alternate Marking are illustrated in Figure 2. They are present when the NH field of the AltMark TLV is set to 0x9.




Figure 2: Optional Extended Data Fields for Enhanced Alternate Marking

The extended data fields are as follows:

- * FlowMonID Ext - 20 bits unsigned integer. This is used to extend the FlowMonID in order to reduce the conflict when random allocation is applied. The disambiguation of the FlowMonID field is discussed in IPv6 AltMark Option [RFC9343].
- * Four bit-flags indicate special-purpose usage.

M bit: Measurement mode. If M=0, it indicates that it is for

hop-by-hop monitoring. If M=1, it indicates that it is for end-to-end monitoring.

F bit: Fragmentation. If F=1, it indicates that the original packet is fragmented, therefore it is necessary to only count a single packet, ignoring all the following fragments with F set to 1.

W bit: Flow direction identification. This flag is used if backward direction flow monitoring is requested to be set up automatically. If W=1, it indicates that the flow direction is forward. If W=0, it indicates that the flow direction is backward.

R bit: Reserved. This bit MUST be set to zero and ignored on receipt.

- * Len - Length. Indicates the length of the extended data fields for enhanced alternate marking. It includes all of the fields shown in Figure 2 including any meta data that is present.
- * Rsvd - Reserved for further use. These bits MUST be set to zero on transmission and ignored on receipt.
- * MetaInfo - A 16-bit Bitmap to indicate more meta data attached in the Optional MetaData field for enhanced functions. More than one bit may be set, in which case the additional meta data is present in the order that the bits are set. MetaInfo bits are numbered from 0 as the most significant bit. Three bits and associated meta data are defined as follows:

bit 0: If set to 1, it indicates that a 6 byte Timestamp is present as shown in Figure 3.




Figure 3: The Timestamp Extended Data Field

This Timestamp can be filled by the encapsulation node, and is taken all the way to the decapsulation node so that all the intermediate nodes can compare it against their local time, and measure the one way delay. The timestamp consists of two fields:

Timestamp(s) is a 16 bit integer that carries the number of seconds.

Timestamp(ns) is a 32 bit integer that carries the number of nanoseconds.

bit 1: If set to 1, it indicates that control information to set up the backward direction flow monitoring based on the trigger packet is present as shown in Figure 4.




Figure 4: Control Information for Backward Direction Flow Monitoring

The control information includes several fields and flags to match in order to set up the backward direction:

DIP Mask: The length of the destination IP prefix used to match the flow.

SIP Mask: The length of the source IP prefix used to match the flow.

P bit: If set to 1, it indicates to match the flow using the protocol identifier in the trigger packet.

I bit: If set to 1, it indicates to match the source port.

O bit: If set to 1, it indicates to match the destination port.

V bit: If set to 1, the node will automatically set up reverse direction monitoring, and allocate a FlowMonID.

S bit: If set to 1, it indicates to match the DSCP.

T bit: Used to control the scope of tunnel measurement. T=1 means measure between Network-to-Network Interfaces (i.e., NNI to NNI). T=0 means measure between User-to-Network Interfaces (i.e., UNI to UNI).

Period: Indicates the alternate marking period counted in seconds.

bit 2: If set to 1, it indicates a 4 byte sequence number is present as shown in Figure 5.




Figure 5: Sequence Number Data Field

The unique Sequence Number can be used to detect the out-of-order packets, in addition to enabling packet loss measurement. Moreover, the Sequence Number can be used together with the latency measurement, to access per packet timestamps.

4. Use of the SRH AltMark TLV

SRv6 leverages the IPv6 Segment Routing Header (SRH). The SRH can carry TLVs as described in [RFC8754]. This document defines the SRH AltMark TLV (see Section 3) to carry Alternative Marking data fields for use in SRv6 networks.

Assuming that the measurement domain overlaps with the SR controlled domain, the procedure for AltMark data encapsulation in the SRv6 SRH is summarized as follows:

- * Ingress SR Node: As part of the SRH encapsulation, the Ingress SR Node of an SR domain or an SR Policy [RFC9256] that supports the mechanisms defined in this document and that wishes to perform the Alternate Marking Method adds the AltMark TLV in the SRH of the data packets.
- * Intermediate SR Node: The Intermediate SR Node is any node receiving an IPv6 packet where the destination address of that packet is a local Segment Identifier (SID). If an Intermediate SR Node is not capable of processing AltMark TLV, it simply ignores it according to the processing rules of [RFC8754]. If an

Intermediate SR Node is capable of processing AltMark TLV, it checks if SRH AltMark TLV is present in the packet and processes it.

- * Egress SR Node: The Egress SR Node is the last node in the segment list of the SRH. The processing of AltMark TLV at the Egress SR Node is similar to the processing of AltMark TLV at the Intermediate SR Nodes.

The use of the AltMark TLV may be combined with the network programming capability of SRv6 ([RFC8986]). Specifically, the ability for an SRv6 endpoint to determine whether to process or ignore some specific SRH TLVs (such as the AltMark TLV) may be based on the SID function associated with the SID advertised by an Intermediate or Egress SR Node and used in the Destination Address field of the SRv6 packet. When a packet is addressed to a SID which does not support the Alternate Marking functionality, the receiving node does not have to look for or process the SRH AltMark TLV and can simply ignore it. This also enables collection of Alternate Marking data only from the supporting segment endpoints.

5. Observations on RFC 9343

Like any other IPv6 use case, Hop-by-Hop and Destination Options can also be used when the SRH is present. As specified in [RFC8200], the Hop-by-Hop Options Header is used to carry optional information that needs be examined at every hop along the path, while the Destination Options Header is used to carry optional information that needs be examined only by the packet's destination node(s).

When a Routing Header exists, the Destination Options before the Routing Header is "for options to be processed by the first destination that appears in the IPv6 Destination Address field plus subsequent destinations listed in the Routing header", while the Destination Options after the Routing Header is "for options to be processed only by the final destination of the packet". Because the SRH is a Routing Header, Destination Options present in the IPv6 packet before the SRH header are processed by destination indicated in the SRH's route list. As specified in [RFC8754], SR segment endpoint nodes process the local SID corresponding to the packet destination address. Then, the destination address is updated according to the segment list. The SRH TLV provides metadata for segment processing, while processing the SID, if the node is locally configured to do so. From the aspect of processing function, both the Destination Options Header before SRH and the SRH TLV are processed at the node being indicated in the destination address field of the IPv6 header.

The distinction between the approaches is most notable for SRv6 packets that traverse a network where the paths between sequential segment end points include multiple hops. If the Hop-by-Hop Option is used, then every hop along the path will process the AltMark data. If the Destination Option positioned before the SRH is used, or the SRH AltMark TLV is used, then only the segment end points will process the AltMark data.

Thus, the Alternate Marking Method can be achieved in two ways in an SRv6 network: either using the mechanism defined in this document, or using Destination Option preceding the SRH to carry AltMark data fields as described in [RFC9343]. These solutions can co-exist according to the current specifications, which raises an issue in deployments.

But, it is to be noted that the approach with the Destination Option requires two IPv6 extension headers and this can have operational implications, as described in [RFC9098] and [I-D.ietf-6man-eh-limits]. It may, therefore, be desirable to choose the use of the SRH AltMark TLV, as described in this document, in order to limit the number of extension headers present in the packets. [I-D.peng-v6ops-eh-deployment-considerations] also analyzes the issues with the extension headers, and aims to provide deployment guidance when IPv6 options are used.

Further, there is a simplification in placing all information related to an SRv6 path inside the SRH, and that includes the Alternate Marking Method information associated with that path. Additionally, it is likely that SR nodes support fast parsing and processing of the SRH, while it is possible that SR nodes may be explicitly configured to not handle Destination Options for security and legacy reasons.

From a device prospective, SRH TLV and Destination Options are generally two functional modules in the forwarding plane. The difference is that SRH and SRH TLV are integrated modules, while Destination Option is a general IPv6 functional module. Supporting two modules (SRH and Destination Options) at the same time may not be optimal and may consume resources.

Moreover, this document also introduces in Section 3.2 extended data fields, which are not defined in [RFC9343], to support additional telemetry requirements. In particular, [I-D.ietf-opsawg-ipfix-on-path-telemetry] introduces new IP Flow Information Export (IPFIX) information elements to expose the On-Path Telemetry measured delay. It defines how the timestamp can be encoded in the encapsulation node and be read at the intermediate and decapsulation node to calculate the on-path delay. Therefore, the implementation of the SRH AltMark TLV, as defined in this document, is also correlated with the implementation of [I-D.ietf-opsawg-ipfix-on-path-telemetry].

For all these reasons, the preferred solution for Alternate Marking Method in SR networks can be the SRH AltMark TLV as defined in this document, while the solution for other IPv6 networks is as described in [RFC9343]. This document does not change or invalidate any procedures defined in [RFC9343].

As noted in Section 3, implementations of this document must use a code point chosen from the Experimental range. Such implementations should make it possible for the operator to configure the value used in a deployment such that it is possible to conduct multiple implementations within the same network.

6. SRH AltMark TLV Compatibility

As highlighted in the previous section, the use of the Destination Option to carry the AltMark data preceding the SRH is equivalent to the SRH AltMark TLV. Therefore, it is important to analyze what happens when both the SRH AltMark TLV and the Destination Option are used, and how that would impact processing and complexity. There are significant benefits to having only one solution for any problem. It simplifies implementation and makes deployment less complicated, reducing configuration and interoperability issues.

It is worth mentioning that the SRH AltMark TLV and the the Destination Option carrying AltMark data can coexist without problems. If both are present, the only issue could be the duplication of information but this will not affect in any way the device and the network services. The security requirement of controlled domain applies to both this document and [RFC9343], and it also confines this duplication to a single service provider networks. However, duplication of the same information in different places should be avoided and this document recommends the use of SRH TLV to carry SRv6 related information.

How the Alternate Marking Method is applied in a specific controlled domain also involves the specific capabilities of the devices in the network. The use of the SRH AltMark TLV should be evaluated before supporting the Alternate Marking Method capability. It is recommended to employ capability advertisement mechanisms which can be utilized for this purpose. The choice between SRH TLV and Destination Option can be up to the network operator, depending on the service requirements and network device characteristics.

7. Implementation Overview

This document describes a protocol extension built on existing technology and using an Experimental code point. The purpose is to determine the practicality and optimality of the protocol extension, in particular in consideration of implementations that cannot support multiple IPv6 extension headers in the same packet, or which do not support Destination Option Header processing, or which process the Destination Option Header on the slow path.

The deployment should determine whether the protocol extensions defined achieve the desired function and can be supported in the presence of normal SRv6 processing especially in regard to concerns about SRH size and the potential complexity of SRH TLV processing. In particular, it needs to verify the ability to support SR network programming support of SID function control of the support or non-support of the AltMark TLV.

It is anticipated that the implementation with the AltMark TLV will be contained within single service provider networks in keeping with the normal constraints of an SR Domain, and also in keeping with the normal limits in sharing performance and monitoring data collected on the path of packets in the network. The scope of the deployment may depend on the availability of implementations and the willingness of operators to deploy it on live networks. Other implementers and depolyers are invited to share their experiences with the authors of this document.

The results of this implementation will be collected and shared with the IETF SPRING working group (possibly as an Internet-Draft) to help forward the discussions that will determine the correct development of Alternate marking Method solutions in SRv6 networks. It is expected that a first set of results will be made available within two years of the publication of this document as an RFC.

8. Security Considerations

The security considerations of SRv6 are discussed in [RFC8754] and [RFC8986], and the security considerations of Alternate Marking in general and its application to IPv6 are discussed in [RFC9341] and [RFC9343].

[RFC9343] analyzes different security concerns and related solutions. These aspects are valid and applicable also to this document. In particular the fundamental security requirement is that Alternate Marking MUST only be applied in a limited domain, as also mentioned in [RFC8799] and Section 2.1.

Alternate Marking is a feature applied to a trusted domain, where one or several operators decide on leveraging and configuring Alternate Marking according to their needs. Additionally, operators need to properly secure the Alternate Marking domain to avoid malicious configuration and attacks, which could include injecting malicious packets into a domain. So the implementation of Alternate Marking is applied within a controlled domain where the network nodes are locally administered and where packets containing the AltMark TLV are prevented from entering or leaving the domain. A limited administrative domain provides the network administrator with the means to select, monitor and control the access to the network.

9. IANA Considerations

This document makes no requests for IANA actions. The code point used is taken from an Experimental range, must be agreed between implementations within any individual deployment, and is not to be reported in any published specification. [RFC8754] allows for SRH TLV code points for experimentation and testing. It could be possible to reserve some code point values for specific behaviors, such as the implementation described in this document, but this is out of scope for this document.

10. Acknowledgements

The authors would like to thank Adrian Farrel and Haoyu Song for the precious comments and suggestions.

11. Contributors

The following people provided relevant contributions to this document:

Massimo Nilo
Telecom Italia
Email: massimo.nilo@telecomitalia.it

Fabrizio Milan
Telecom Italia
Email: fabrizio.milan@telecomitalia.it

Fabio Bulgarella
Telecom Italia
Email: fabio.bulgarella@guest.telecomitalia.it

12. References

12.1. 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/info/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/info/rfc8174>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [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/info/rfc9341>>.
- [RFC9342] Fioccola, G., Ed., Cociglio, M., Sapio, A., Sisto, R., and T. Zhou, "Clustered Alternate-Marking Method", RFC 9342, DOI 10.17487/RFC9342, December 2022, <<https://www.rfc-editor.org/info/rfc9342>>.
- [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/info/rfc9343>>.

12.2. Informative References

- [I-D.ietf-6man-eh-limits]
Herbert, T., "Limits on Sending and Processing IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-ietf-6man-eh-limits-12, 18 December 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-eh-limits-12>>.
- [I-D.ietf-opsawg-ipfix-on-path-telemetry]
Graf, T., Claise, B., and A. H. Feng, "Export of On-Path Delay in IPFIX", Work in Progress, Internet-Draft, draft-ietf-opsawg-ipfix-on-path-telemetry-06, 14 January 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ipfix-on-path-telemetry-06>>.
- [I-D.peng-v6ops-eh-deployment-considerations]
Peng, S., Fioccola, G., and J. Dong, "Deployment considerations of IPv6 packets with options", Work in Progress, Internet-Draft, draft-peng-v6ops-eh-deployment-considerations-00, 13 March 2023, <<https://datatracker.ietf.org/doc/html/draft-peng-v6ops-eh-deployment-considerations-00>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [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/info/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/info/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/info/rfc8986>>.
- [RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. Liu, "Operational Implications of IPv6 Packets with Extension Headers", RFC 9098, DOI 10.17487/RFC9098, September 2021, <<https://www.rfc-editor.org/info/rfc9098>>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

Authors' Addresses

Giuseppe Fioccola
Huawei
Palazzo Verrocchio, Centro Direzionale Milano 2
20054 Segrate (Milan)
Italy
Email: giuseppe.fioccola@huawei.com

Tianran Zhou
Huawei
156 Beiqing Rd.
Beijing
100095
China
Email: zhoutianran@huawei.com

Mauro Cociglio
Telecom Italia
Email: mauro.cociglio@outlook.com

Gyan S. Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com

Xuewei Wang
Ruijie
Email: wangxuewei1@ruijie.com.cn

Geng Zhang
China Mobile
Email: zhanggeng@chinamobile.com