

IPPM
Internet-Draft
Intended status: Informational
Expires: 31 August 2025

G. Fioccola
K. Zhu
Huawei
T. Graf
Swisscom
M. Nilo
Telecom Italia
L. Zhang
China Mobile
27 February 2025

Alternate Marking Deployment Framework
draft-ietf-ippm-alt-mark-deployment-03

Abstract

This document provides a framework for Alternate Marking deployment and includes considerations and guidance for the deployment of the methodology.

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 31 August 2025.

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
1.1. Requirements Language	3
1.2. Terminology	3
2. Alternate Marking Deployment Domain	4
3. Alternate Marking Measurement Nodes	5
4. Type of Measurements	6
5. Configuration	7
6. Data Export	8
6.1. IPFIX	9
6.2. YANG Push	9
6.3. Hybrid Two-Step	10
7. Encapsulations	10
7.1. IPv6	10
7.2. SRv6	10
7.3. BIER	10
7.4. MPLS	10
7.5. SFC	10
7.6. NVO3	11
7.7. Enhanced capabilities	11
8. Implementation Observations	11
9. Security Considerations	11
10. IANA Considerations	12
11. Acknowledgements	12
12. Contributors	12
13. References	12
13.1. Normative References	12
13.2. Informative References	13
Authors' Addresses	17

1. Introduction

The Alternate Marking [RFC9341] and Multipoint Alternate Marking [RFC9342] define the Alternate Marking technique that is a hybrid performance measurement method, per [RFC7799] classification of measurement methods. This method is based on marking consecutive batches of packets and it can be used to measure packet loss, latency, and jitter on live traffic.

The first experiments on Alternate-Marking are described in [RFC8321] and [RFC8889].

According to the definitions of [RFC7799], the Alternate-Marking Method can be classified as Hybrid Type I. Indeed, Alternate Marking can be implemented by using reserved bits in the protocol header, and the change in value of these marking bits at the source node is formally considered a modification of the stream of interest.

This document complements [RFC9341] and [RFC9342] as it explains the mechanisms that can be used to manage and deploy the method.

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.

1.2. Terminology

Abbreviations used in this document:

AltMark: Alternate-Marking

NMS: Network Management System

IPv6: Internet Protocol version 6

SRv6: Segment Routing over IPv6 dataplane

BIER: Bit Index Explicit Replication

MPLS: Multi-Protocol Label Switching

SFC: Service Function Chaining

NVO3: Network Virtualization Overlays

IPFIX: IP Flow Information Export

YANG: Yet Another Next Generation

PCEP: Path Computation Element Communication Protocol

BGP: Border Gateway Protocol

2. Alternate Marking Deployment Domain

The Alternate Marking Method MUST be deployed in a controlled domain for security and compatibility reasons. In this regard, [RFC8799] reports further examples of specific limited domain solutions. It is not common that the user traffic originates and terminates within the controlled domain. For this reason, it will typically only be applicable in an overlay network, where user traffic is encapsulated at one domain border, decapsulated at the other domain border and the encapsulation incorporates the relevant extension header for Alternate Marking. This requirement also implies that an implementation MUST filter packets that carry Alternate Marking data and are entering or leaving the controlled domain.

A controlled domain is a managed network where it is required to select, monitor and control the access to the network by enforcing policies at the domain boundaries in order to discard undesired external packets entering the domain and check the internal packets leaving the domain. It does not necessarily mean that a controlled domain is a single administrative domain or a single organization. A controlled domain can correspond to a single administrative domain or can be composed by multiple administrative domains under a defined network management. Indeed, some scenarios may imply that the Alternate Marking Method involves more than one domain, but in these cases, it is RECOMMENDED that the multiple domains create a whole controlled domain while traversing the external domain by employing IPsec authentication and encryption or other VPN technology that provides full packet confidentiality and integrity protection. In a few words, it must be possible to control the domain boundaries and eventually use specific precautions if the traffic traverse the Internet.

The Alternate Marking measurement domain can overlap with the controlled domain or may be a subset of the controlled domain. The typical scenarios for the application of the Alternate Marking Method depend on the controlled domain boundaries, in particular:

the user equipment can be the starting or ending node, only in case it is fully managed and if it belongs to the controlled domain. In this case the user generated packets contain the Alternate Marking data. But, in practice, this is not common due to the fact that the user equipment cannot be totally secured in the majority of cases.

the CPE (Customer Premises Equipment) or the PE (Provider Edge) routers are most likely to be the starting or ending nodes since they can be border routers of the controlled domain. For instance, the CPE, which connects the user's premises with the

service provider's network, belongs to a controlled domain only if it is managed by the service provider and if additional security measures are taken to keep it trustworthy. Typically the CPE or the PE can encapsulate a received packet in an outer header which contains the Alternate Marking data. They can also be able to filter and drop packets from outside of the domain with inconsistent fields to make effective the relevant security rules at the domain boundaries, for example a simple security check can be to insert the Alternate Marking data if and only if the destination is within the controlled domain.

3. Alternate Marking Measurement Nodes

An Alternate-Marking Domain consists of marking nodes, unmarking nodes, and transit nodes.

A marking node, also called encapsulating node, incorporates the AltMark Data Fields into packets in order to enable Alternate-Marking. If the Alternate-Marking method is enabled for a selected flow of the traffic, the encapsulating node is responsible for applying the AltMark functionality to the selected flow and to take initial timestamps and packet counters.

A transit node only reads AltMark Data Fields in order to take timestamps and packet counters.

An unmarking node, also called decapsulating node, reads AltMark Data Fields in order to take final timestamps and packet counters and then removes any AltMark Option from packets.

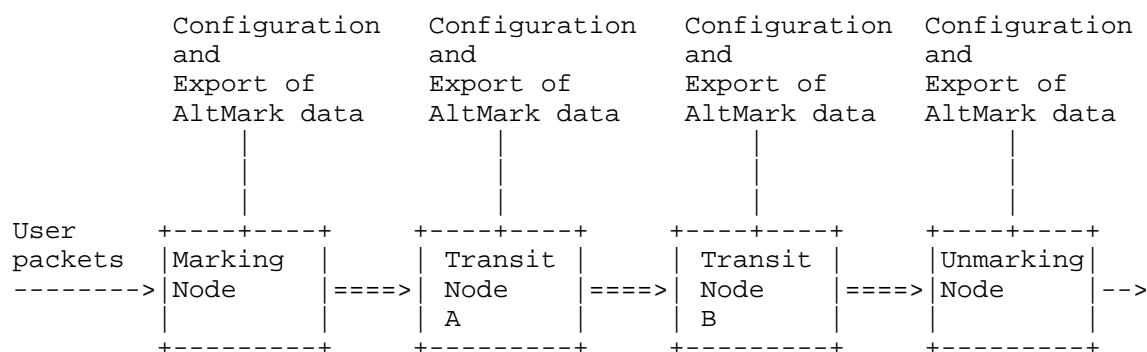


Figure 1: Roles of Alternate-Marking Nodes

4. Type of Measurements

The methodology described in the previous sections can be applied to various performance measurement problems. The only requirement is to select and mark the flow to be monitored; in this way, packets are batched by the sender, and each batch is alternately marked such that it can be easily recognized by the receiver.

Either one or two flag bits might be available for marking in different deployments:

One flag: packet loss measurement **MUST** be done as described in Section 3.1 of [RFC9341], while delay measurement **MUST** be done according to the single-marking method described in Section 3.2.1 of [RFC9341]. Mean delay (Section 3.2.1.1 of [RFC9341]) **MAY** also be used but it could imply more computational load.

Two flags: packet loss measurement **MUST** be done as described in Section 3.1 of [RFC9341], while delay measurement **MUST** be done according to double-marking method Section 3.2.2 of [RFC9341]. In this case single-marking **MAY** also be used in combination with double-marking and the two approaches provide slightly different pieces of information that can be combined to have a more robust data set.

There are some operational guidelines to consider for the purpose of deciding to follow the recommendations above and use one or two flags.

The Alternate-Marking method utilizes specific flags in the packet header, so an important factor is the number of flags available for the implementation. Indeed, if there is only one flag available there is no other way, while if two flags are available the option with two flags is certainly more complete.

The duration of the Alternate-Marking period affects the frequency of the measurement and this is a parameter that can be decided on the basis of the required temporal sampling. But it cannot be freely chosen, as explained in Section 5 of [RFC9341].

The Alternate-Marking methodologies enable packet loss, delay and delay variation calculation, but in accordance with the method used (e.g. single-marking or double-marking), there is different kind of information that can be derived. For example, to get more statistics of extent data, the option with two flags is desirable. For this reason, the type of data needed in the specific scenario is an additional element to take into account.

The Alternate-Marking methods imply different computational load depending on the method employed. Therefore, the available computational resources on the measurement points can also influence the choice. As an example, mean delay calculation may require more processing and it may not be the best option to minimize the computational load.

A deployment of the Alternate-Marking Method should also take into account how to handle and recognize marked and unmarked traffic. Since Alternate-Marking normally employs a marking field which is dedicated, reserved, and included in a protocol extension, the measurement points can learn whether the measurement is activated or not by checking if the specific extension is included or not within the packets.

5. Configuration

The YANG model can be used for the definition of the AltMark data sent over network management protocols such as the NETCONF and RESTCONF. They can be used for configuring Alternate-Marking in network nodes that support it. An example of the Alternate-Marking YANG model is defined in [I-D.ietf-ippm-alt-mark-yang].

There are also other control plane mechanisms to advertise and activate AltMark capabilities, using PCEP or BGP:
[I-D.ietf-idr-sr-policy-ifit], [I-D.ietf-idr-bgp-ifit-capabilities],
[I-D.ietf-pce-pcep-ifit].

These mechanisms can be used to signal and configure the parameters to identify the flow to monitor both in case of point-to-point flow and multipoint-to-multipoint flow. Indeed, the selection of the identification fields directly affects the type of paths that the flow would follow in the network. As an example, for IPv6 the setting of the Flow Monitoring Identification (FlowMonID) is used in combination with source and destination addresses to identify a flow, as described in Section 5.3 of [RFC9343], and it can be pseudo-randomly generated by the source node or assigned by the central controller.

Additionally, other parameters are essential for the activation of the AltMark methodology: the choice between end-to-end or hop-by-hop measurement, the choice between the methods with one flag or two flags and the duration of the Alternate-Marking period which affects the measurement frequency (longer the duration of the block, the less frequently the measurement can be taken).

6. Data Export

Each packet marked for Alternate-Marking, as for example the AltMark IPv6 option type defined in Section 3.1 of [RFC9343] or the Segment Routing TLV Type as defined in Section 3.1 of [I-D.fz-spring-srv6-alt-mark] MUST be copied to the IPFIX or YANG push metering process depending which Network Telemetry [RFC9232] protocol is used to export the data.

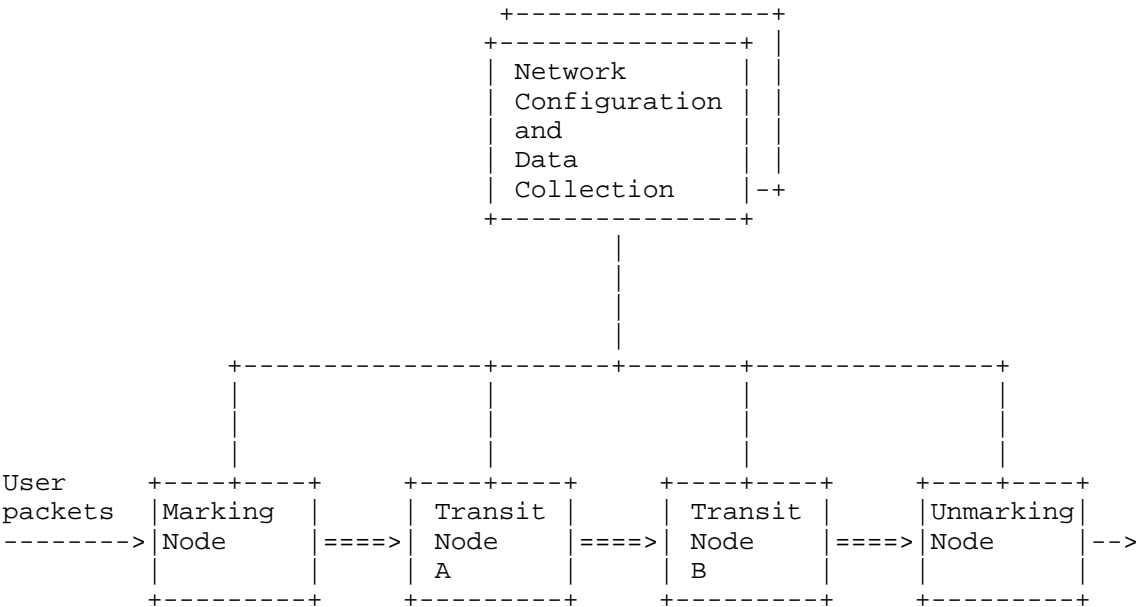


Figure 2: Alternate-Marking Framework with Configuration and Data Export

When data is collected packet counts and timestamps are reported to the collector, but a certain synchronization mechanism is required to ensure that the collected data is correlated. Therefore, the Period Number (PN) can be used to help to determine the packet counts related to the same block of markers, or the timestamps related to the same marked packet. The PN is generated each time a node reads the packet counts or timestamps, and is associated with each packet count and timestamp reported. The assumption is that the nodes are time synchronized as described in [RFC9341] and [RFC9342]. The PN can be calculated as the modulo of the local time (when the counts or timestamps are read) and the interval of the marking time period.

6.1. IPFIX

The new Information Elements (IEs) to export Alternate Marking measurement data are specified in [I-D.ietf-opsawg-ipfix-alt-mark].

For IPFIX [RFC7011], the data decomposition can be achieved on the Alternate-Marking-aware node exporting the data or on the data collection. When decomposed at the data collection, the headers, as example the IPv6 options type header described in Section 3.1 of [RFC9343] or the Segment Routing header TLV as described in Section 3.1 of [I-D.fz-spring-srv6-alt-mark] containing the FlowMonID, Loss and Delay flag are being exposed as part of `ipPayloadPacketSection(IE314)`, defined in Section 4.2 of [RFC7133]. When being decomposed on the Alternate-Marking-aware node, new IPFIX entities for FlowMonID, Loss and Delay flag are needed so that the data can now be aggregated according to section 5 of [RFC7015]. FlowMonID, Loss and Delay flag are Flow Key fields. The IPFIX entities, which are of interest to describe the relationship to the forwarding topology and the control-plane are further described in [I-D.ietf-opsawg-ipfix-alt-mark].

To calculate loss, the packet count can be done with `octetDeltaCount(IE1)` or `packetDeltaCount(IE2)`. And to calculate delay, either `flowStartSeconds(IE150)`, `flowStartMilliseconds(IE152)`, `flowStartMicroseconds(IE154)` or `flowStartNanoseconds(IE156)`, can be used depending on timestamp granularity requirements. It is also possible to use `flowEndSeconds(IE151)`, `flowEndMilliseconds(IE153)`, `flowEndMicroseconds(IE155)` or `flowEndNanoseconds(IE157)`.

6.2. YANG Push

For YANG Push [RFC8639], periodical subscription as defined in Section 3.1 of [RFC8641] is used to subscribe data. Decomposition is done on the Alternate-Marking-aware node publishing the data. The YANG module contains FlowMonID as key, Loss and Delay flag, ingress and egress interface `ifIndex` [RFC2863], octet delta count describing the amount of observed packets within a flow to measure loss, and flow start timestamp describing the first packet observed for measuring delay as leafes. [I-D.fz-ippm-on-path-telemetry-yang] introduces a YANG data model for monitoring Alternate-Marking telemetry data.

Since the amount of observed data could overwhelm a route processor on a network node, publishing data from network processors as specified in [I-D.ietf-netconf-distributed-notif] is advised.

6.3. Hybrid Two-Step

Another option for collecting AltMark data can be the Hybrid Two-Step (HTS) approach, defined in [I-D.ietf-ippm-hybrid-two-step]. It is a method of telemetry collection that separates the act of measuring the performance metric from collecting this information. A packet in the flow to which the AltMark method is applied can be used as an HTS Trigger, while the HTS Follow-up packet is then generated to collect the measurement data from the nodes, as further detailed in [I-D.ietf-ippm-hybrid-two-step].

7. Encapsulations

7.1. IPv6

The Alternate-Marking encapsulation for IPv6 is defined in [RFC9343], which also discusses deployment considerations for IPv6 networks.

The IPv6 AltMark Option [RFC9343] applies the Alternate Marking Method to IPv6, and defines an Extension Header Option to encode the Alternate Marking Method for both the Hop-by-Hop Options Header and the Destination Options Header.

Other alternatives for the application of the Alternate-Marking Method to IPv6 are introduced in [I-D.wang-ippm-ipv6-flow-measurement] and [I-D.wang-ippm-ipv6-distributed-flow-measurement].

7.2. SRv6

The Alternate-Marking encapsulation for SRv6 is discussed in IPv6 AltMark Option [RFC9343] and [I-D.fz-spring-srv6-alt-mark].

7.3. BIER

The Alternate-Marking encapsulation for BIER is introduced in [I-D.ietf-bier-pmmm-oam].

7.4. MPLS

The Alternate-Marking application to MPLS has some options, as introduced in [RFC9571], [RFC9714], [I-D.cx-mpls-mna-inband-pm].

7.5. SFC

The Alternate-Marking encapsulation for SFC is introduced in [I-D.mfm-ippm-sfc-nsh-pmamm].

7.6. NVO3

The Alternate-Marking encapsulation for NVO3 is introduced in [I-D.fmm-nvo3-pm-alt-mark].

7.7. Enhanced capabilities

[I-D.zhou-ippm-enhanced-alternate-marking] defines extended data fields for the AltMark Option and provides enhanced capabilities to overcome some challenges and enable future proof applications.

It is worth mentioning that the enhanced capabilities are intended for further use and are optional.

8. Implementation Observations

In a controlled domain, the nodes may support the AltMark specific encapsulation and this also depends on the implementation. If a node is configured to read the AltMark option, the measurement is done on that node, otherwise it is simply not considered in the measurement.

Assuming that the measurement domain overlaps with the controlled domain, the procedure for AltMark data encapsulation can be summarized as follows:

- * Ingress Marking Node: the Ingress Node of a controlled domain that supports the Alternate Marking Method adds the AltMark data in the the data packets.
- * Intermediate Transit Node: if an Intermediate Node is not capable of processing the AltMark data, it simply ignores it. If an Intermediate Node is capable of processing the AltMark data, it processes it.
- * Egress Unmarking Node: The Egress Node is the last node of the controlled domain. The processing if the AltMark data is similar to the processing at the Intermediate Nodes. The only difference is that it needs to remove the AltMark data from the the data packets.

9. Security Considerations

Alternate Marking [RFC9341] and Multipoint Alternate Marking [RFC9342] analyze 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 specific limited domain, as also mentioned in [RFC8799].

10. IANA Considerations

This document has no request to IANA.

11. Acknowledgements

The authors of this document would like to thank Greg Mirsky, Chongfeng Xie and Xiao Min for their comments and reviews.

12. Contributors

Tianran Zhou
Huawei
Email: zhoutianran@huawei.com

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

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

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

13. References

13.1. 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>>.
- [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>>.

13.2. Informative References

- [I-D.cx-mpls-mna-inband-pm]
Cheng, W., Min, X., Gandhi, R., and G. Mirsky, "MNA for Performance Measurement with Alternate Marking Method", Work in Progress, Internet-Draft, draft-cx-mpls-mna-inband-pm-05, 21 October 2024, <<https://datatracker.ietf.org/doc/html/draft-cx-mpls-mna-inband-pm-05>>.
- [I-D.fmm-nvo3-pm-alt-mark]
Fioccola, G., Mirsky, G., and T. Mizrahi, "Performance Measurement (PM) with Alternate Marking in Network Virtualization Overlays (NVO3)", Work in Progress, Internet-Draft, draft-fmm-nvo3-pm-alt-mark-03, 23 October 2018, <<https://datatracker.ietf.org/doc/html/draft-fmm-nvo3-pm-alt-mark-03>>.
- [I-D.fz-ippm-on-path-telemetry-yang]
Fioccola, G. and T. Zhou, "On-path Telemetry YANG Data Model", Work in Progress, Internet-Draft, draft-fz-ippm-on-path-telemetry-yang-01, 20 December 2024, <<https://datatracker.ietf.org/doc/html/draft-fz-ippm-on-path-telemetry-yang-01>>.
- [I-D.fz-spring-srv6-alt-mark]
Fioccola, G., Zhou, T., Cociglio, M., Mishra, G. S., wang, X., and G. Zhang, "Application of the Alternate Marking Method to the Segment Routing Header", Work in Progress, Internet-Draft, draft-fz-spring-srv6-alt-mark-10, 3 February 2025, <<https://datatracker.ietf.org/doc/html/draft-fz-spring-srv6-alt-mark-10>>.

`[I-D.ietf-bier-pmmm-oam]`

Mirsky, G., Zheng, L., Chen, M., and G. Fioccola,
"Performance Measurement (PM) with Marking Method in Bit
Index Explicit Replication (BIER) Layer", Work in
Progress, Internet-Draft, draft-ietf-bier-pmmm-oam-16, 7
November 2024, <[https://datatracker.ietf.org/doc/html/
draft-ietf-bier-pmmm-oam-16](https://datatracker.ietf.org/doc/html/draft-ietf-bier-pmmm-oam-16)>.

`[I-D.ietf-idr-bgp-ifit-capabilities]`

Fioccola, G., Pang, R., Wang, S., Decraene, B., Zhuang,
S., and H. Wang, "Advertising In-situ Flow Information
Telemetry (IFIT) Capabilities in BGP", Work in Progress,
Internet-Draft, draft-ietf-idr-bgp-ifit-capabilities-06,
21 October 2024, <[https://datatracker.ietf.org/doc/html/
draft-ietf-idr-bgp-ifit-capabilities-06](https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ifit-capabilities-06)>.

`[I-D.ietf-idr-sr-policy-ifit]`

Qin, F., Yuan, H., Yang, S., Zhou, T., and G. Fioccola,
"BGP SR Policy Extensions to Enable IFIT", Work in
Progress, Internet-Draft, draft-ietf-idr-sr-policy-ifit-
09, 18 October 2024,
<[https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-
policy-ifit-09](https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-ifit-09)>.

`[I-D.ietf-ippm-alt-mark-yang]`

Graf, T., Wang, M., Fioccola, G., Zhou, T., and X. Min, "A
YANG Data Model for the Alternate Marking Method", Work in
Progress, Internet-Draft, draft-ietf-ippm-alt-mark-yang-
00, 8 January 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-ippm-
alt-mark-yang-00](https://datatracker.ietf.org/doc/html/draft-ietf-ippm-alt-mark-yang-00)>.

`[I-D.ietf-ippm-hybrid-two-step]`

Mirsky, G., Lingqiang, W., Zhui, G., Song, H., and P.
Thubert, "Hybrid Two-Step Performance Measurement Method",
Work in Progress, Internet-Draft, draft-ietf-ippm-hybrid-
two-step-04, 25 February 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-ippm-
hybrid-two-step-04](https://datatracker.ietf.org/doc/html/draft-ietf-ippm-hybrid-two-step-04)>.

`[I-D.ietf-netconf-distributed-notif]`

Zhou, T., Zheng, G., Voit, E., Graf, T., and P. Francois,
"Subscription to Distributed Notifications", Work in
Progress, Internet-Draft, draft-ietf-netconf-distributed-
notif-11, 27 February 2025,
<[https://datatracker.ietf.org/api/v1/doc/document/draft-
ietf-netconf-distributed-notif/](https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-netconf-distributed-notif/)>.

[I-D.ietf-opsawg-ipfix-alt-mark]

Graf, T., Fioccola, G., Zhou, T., Cociglio, M., and M. Nilo, "IP Flow Information Export (IPFIX) Alternate-Marking Information Elements", Work in Progress, Internet-Draft, draft-ietf-opsawg-ipfix-alt-mark-02, 20 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ipfix-alt-mark-02>>.

[I-D.ietf-opsawg-ipfix-srv6-srh]

Graf, T., Claise, B., and P. Francois, "Export of Segment Routing over IPv6 Information in IP Flow Information Export (IPFIX)", Work in Progress, Internet-Draft, draft-ietf-opsawg-ipfix-srv6-srh-14, 25 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ipfix-srv6-srh-14>>.

[I-D.ietf-pce-pcep-ifit]

Yuan, H., 邇矩岌闕」, Yang, P., Li, W., and G. Fioccola, "Path Computation Element Communication Protocol (PCEP) Extensions to Enable IFIT", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-ifit-06, 3 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-ifit-06>>.

[I-D.mfm-ippm-sfc-nsh-pmamm]

Mirsky, G., Fioccola, G., and T. Mizrahi, "Performance Measurement (PM) with Alternate Marking Method in Service Function Chaining (SFC) Network Service Header (NSH) Domain", Work in Progress, Internet-Draft, draft-mfm-ippm-sfc-nsh-pmamm-00, 1 April 2022, <<https://datatracker.ietf.org/doc/html/draft-mfm-ippm-sfc-nsh-pmamm-00>>.

[I-D.wang-ippm-ipv6-distributed-flow-measurement]

Wang, H., Li, J., Lin, C., Min, X., and G. Mirsky, "Distributed Flow Measurement in IPv6", Work in Progress, Internet-Draft, draft-wang-ippm-ipv6-distributed-flow-measurement-06, 29 December 2024, <<https://datatracker.ietf.org/doc/html/draft-wang-ippm-ipv6-distributed-flow-measurement-06>>.

[I-D.wang-ippm-ipv6-flow-measurement]

Wang, H., Liu, Y., Lin, C., Min, X., Mirsky, G., and J. Li, "Flow Measurement in IPv6 Network", Work in Progress, Internet-Draft, draft-wang-ippm-ipv6-flow-measurement-08, 29 December 2024, <<https://datatracker.ietf.org/doc/html/draft-wang-ippm-ipv6-flow-measurement-08>>.

- [I-D.zhou-ippm-enhanced-alternate-marking]
Zhou, T., Fioccola, G., Liu, Y., Cociglio, M., Pang, R., Xiong, L., Lee, S., and W. Li, "Enhanced Alternate Marking Method", Work in Progress, Internet-Draft, draft-zhou-ippm-enhanced-alternate-marking-16, 27 November 2024, <<https://datatracker.ietf.org/doc/html/draft-zhou-ippm-enhanced-alternate-marking-16>>.
- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, DOI 10.17487/RFC2863, June 2000, <<https://www.rfc-editor.org/info/rfc2863>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC7015] Trammell, B., Wagner, A., and B. Claise, "Flow Aggregation for the IP Flow Information Export (IPFIX) Protocol", RFC 7015, DOI 10.17487/RFC7015, September 2013, <<https://www.rfc-editor.org/info/rfc7015>>.
- [RFC7133] Kashima, S., Kobayashi, A., Ed., and P. Aitken, "Information Elements for Data Link Layer Traffic Measurement", RFC 7133, DOI 10.17487/RFC7133, May 2014, <<https://www.rfc-editor.org/info/rfc7133>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.
- [RFC8639] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Subscription to YANG Notifications", RFC 8639, DOI 10.17487/RFC8639, September 2019, <<https://www.rfc-editor.org/info/rfc8639>>.
- [RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.

- [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>>.
- [RFC8889] Fioccola, G., Ed., Cociglio, M., Sapio, A., and R. Sisto, "Multipoint Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8889, DOI 10.17487/RFC8889, August 2020, <<https://www.rfc-editor.org/info/rfc8889>>.
- [RFC9232] Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Network Telemetry Framework", RFC 9232, DOI 10.17487/RFC9232, May 2022, <<https://www.rfc-editor.org/info/rfc9232>>.
- [RFC9571] Bryant, S., Ed., Swallow, G., Chen, M., Fioccola, G., and G. Mirsky, "Extension of RFC 6374-Based Performance Measurement Using Synonymous Flow Labels", RFC 9571, DOI 10.17487/RFC9571, May 2024, <<https://www.rfc-editor.org/info/rfc9571>>.
- [RFC9714] Cheng, W., Ed., Min, X., Ed., Zhou, T., Dai, J., and Y. Peleg, "Encapsulation for MPLS Performance Measurement with the Alternate-Marking Method", RFC 9714, DOI 10.17487/RFC9714, February 2025, <<https://www.rfc-editor.org/info/rfc9714>>.

Authors' Addresses

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

Keyi Zhu
Huawei
156 Beiqing Rd.
Beijing
100095
China
Email: zhukeyi@huawei.com

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland
Email: thomas.graf@swisscom.com

Massimo Nilo
Telecom Italia
Via Reiss Romoli, 274
10148 Torino
Italy
Email: massimo.nilo@telecomitalia.it

Lin Zhang
China Mobile
Email: zhanglin1@cmdi.chinamobile.com