
Stream: Internet Engineering Task Force (IETF)
RFC: [9658](#)
Updates: [7307](#)
Category: Standards Track
Published: September 2024
ISSN: 2070-1721
Authors: IJ. Wijnands M. Mishra, Ed. K. Raza Z. Zhang
Individual Cisco Systems, Inc. Cisco Systems, Inc. Juniper Networks

A. Gulko
Edward Jones

RFC 9658

Multipoint LDP Extensions for Multi-Topology Routing

Abstract

Multi-Topology Routing (MTR) is a technology that enables service differentiation within an IP network. The Flexible Algorithm (FA) is another mechanism for creating a sub-topology within a topology using defined topology constraints and computation algorithms. In order to deploy Multipoint LDP (mLDP) in a network that supports MTR, FA, or other methods of signaling non-default IGP Algorithms (IPAs), mLDP is required to become topology and algorithm aware. This document specifies extensions to mLDP to support the use of MTR/IPAs such that, when building multipoint Label Switched Paths (LSPs), it can follow a particular topology and algorithm. This document updates RFC 7307 by allocating eight bits from a previously reserved field to be used as the "IPA" field.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9658>.

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. 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	3
2. Terminology	4
2.1. Abbreviations	4
2.2. Specification of Requirements	4
3. MT-Scoped mLDP FECs	4
3.1. MP FEC Extensions for MT	5
3.1.1. MP FEC Element	5
3.1.2. MT IP Address Families	5
3.1.3. MT MP FEC Element	6
3.2. Topology IDs	7
4. MT Multipoint Capability	8
5. MT Applicability on FEC-Based Features	9
5.1. Typed Wildcard MP FEC Elements	9
5.2. End-of-LIB	9
6. Topology-Scoped Signaling and Forwarding	10
6.1. Upstream LSR Selection	10
6.2. Downstream Forwarding Interface Selection	10
7. LSP Ping Extensions	10
8. Security Considerations	11

9. IANA Considerations	11
10. References	11
10.1. Normative References	11
10.2. Informative References	12
Contributors	13
Acknowledgments	13
Authors' Addresses	13

1. Introduction

Multi-Topology Routing (MTR) is a technology that enables service differentiation within an IP network. IGPs (e.g., OSPF and IS-IS) and LDP have already been extended to support MTR. To support MTR, an IGP maintains distinct IP topologies referred to as "Multi-Topologies" (or "MTs"), and computes and installs routes specific to each topology. OSPF extensions (see [RFC4915]) and IS-IS extensions (see [RFC5120]) specify the MT extensions under respective IGPs. To support IGP MT, similar LDP extensions (see [RFC7307]) have been specified to make LDP be MT aware and to be able to set up unicast Label Switched Paths (LSPs) along IGP MT routing paths.

A more lightweight mechanism to define constraint-based topologies is the Flexible Algorithm (FA) (see [RFC9350]). The FA is another mechanism for creating a sub-topology within a topology using defined topology constraints and computation algorithms. This can be done within an MTR topology or the default topology. An instance of such a sub-topology is identified by a 1-octet value (Flex-Algorithm) as documented in [RFC9350]. At the time of writing, an FA is a mechanism to create a sub-topology; in the future, different algorithms might be defined for this purpose. Therefore, in the remainder of this document, we'll refer to this as the "IGP Algorithm" or "IPA". The "IPA" field (see Sections 3.1.2 and 5.1) is an 8-bit identifier for the algorithm. The permissible values are tracked in the "IGP Algorithm Types" registry [IANA-IGP-ALGO-TYPES].

Throughout this document, the term "Flexible Algorithm" (or "FA") shall denote the process of generating a sub-topology and signaling it through the IGP. However, it is essential to note that the procedures outlined in this document are not exclusively applicable to the FA: they are extendable to any non-default algorithm as well.

"Multipoint LDP" (or "mLDP") refers to extensions in LDP to set up multipoint LSPs (i.e., point-to-multipoint (P2MP) or multipoint-to-multipoint (MP2MP) LSPs) by means of a set of extensions and procedures defined in [RFC6388]. In order to deploy mLDP in a network that supports MTR and the FA, mLDP is required to become topology and algorithm aware. This document specifies extensions to mLDP to support the use of MTR/IPAs such that, when building multipoint LSPs, it can follow a particular topology and algorithm. Therefore, the identifier for the particular topology to be used by mLDP has to become a 2-tuple {MTR Topology Id, IPA}.

2. Terminology

2.1. Abbreviations

FA: Flexible Algorithm

FEC: Forwarding Equivalence Class

IGP: Interior Gateway Protocol

IPA: IGP Algorithm

LDP: Label Distribution Protocol

LSP: Label Switched Path

mLDP: Multipoint LDP

MP: Multipoint

MP2MP: Multipoint-to-Multipoint

MT: Multi-Topology

MT-ID: Multi-Topology Identifier

MTR: Multi-Topology Routing

MVPN: Multicast VPN in [Section 2.3](#) of [\[RFC6513\]](#)

P2MP Point-to-Multipoint

PMSI Provider Multicast Service Interfaces [\[RFC6513\]](#)

2.2. Specification of Requirements

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

3. MT-Scoped mLDP FECs

As defined in [\[RFC7307\]](#), an MPLS Multi-Topology Identifier (MT-ID) is used to associate an LSP with a certain MTR topology. In the context of MP LSPs, this identifier is part of the mLDP FEC encoding; this is so that LDP peers are able to set up an MP LSP via their own defined MTR policy. In order to avoid conflicting MTR policies for the same mLDP FEC, the MT-ID needs to be a part of the FEC. This ensures that different MT-ID values will result in unique MP-LSP FEC elements.

The same applies to the IPA. The IPA needs to be encoded as part of the mLDP FEC to create unique MP LSPs. The IPA is also used to signal to the mLDP (hop-by-hop) which algorithm needs to be used to create the MP LSP.

Since the MT-ID and IPA are part of the FEC, they apply to all the LDP messages that potentially include an mLDP FEC element.

3.1. MP FEC Extensions for MT

The following subsections define the extensions to bind an mLDP FEC to a topology. These mLDP MT extensions reuse some of the extensions specified in [RFC7307].

3.1.1. MP FEC Element

The base mLDP specification ([RFC6388]) defines the MP FEC Element as follows:

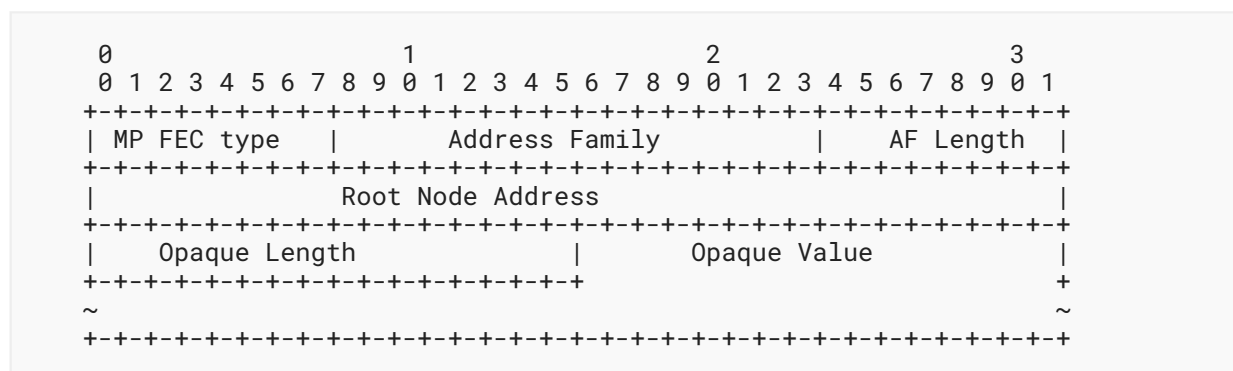


Figure 1: MP FEC Element Format

Where the "Root Node Address" field encoding is defined according to the given "Address Family" field with its length (in octets) specified by the "AF Length" field.

To extend MP FEC elements for MT, the {MT-ID, IPA} tuple is relevant in the context of the root address of the MP LSP. This tuple determines the (sub-)topology in which the root address needs to be resolved. As the {MT-ID, IPA} tuple should be considered part of the mLDP FEC, it is most naturally encoded as part of the root address.

3.1.2. MT IP Address Families

[RFC7307] specifies new address families, named "MT IP" and "MT IPv6," to allow for the specification of an IP prefix within a topology scope. In addition to using these address families for mLDP, 8 bits of the 16-bit "Reserved" field that was described in RFC 7307 are utilized to encode the IPA. The resulting format of the data associated with these new address families is as follows:




Figure 2: Modified Data Format for MT IP Address Families

Where:

IPv4 Address and IPv6 Address: An IP address corresponding to the "MT IP" and "MT IPv6" address families, respectively.

IPA: The IGP Algorithm.

Reserved: This 8-bit field **MUST** be zero on transmission and **MUST** be ignored on receipt.

3.1.3. MT MP FEC Element

When using the extended "MT IP" address family, the resulting MT-Scoped MP FEC element should be encoded as follows:

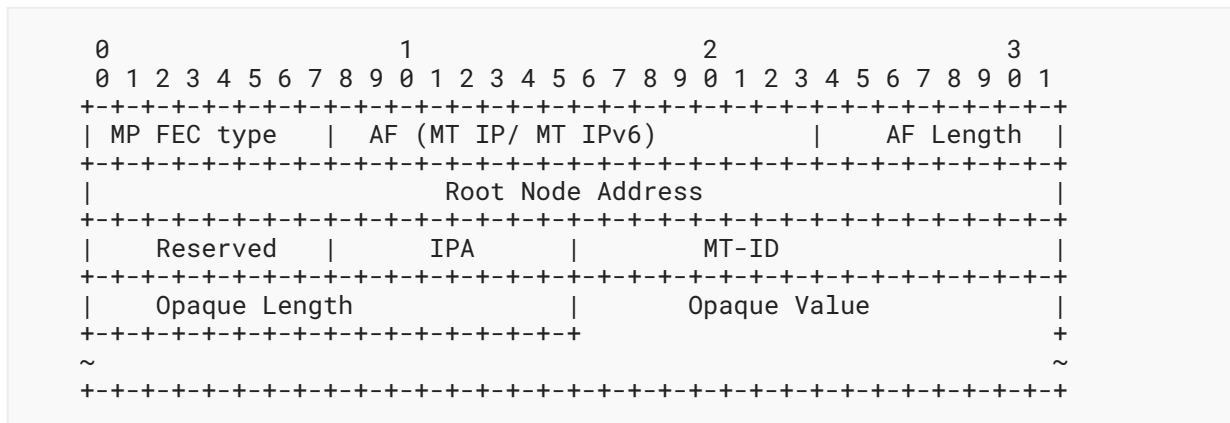


Figure 3: Data Format for an IP MT-Scoped MP FEC Element

In the context of this document, the applicable LDP FECs for MT mLDP ([RFC6388]) include:

- MP FEC Elements:
 - P2MP (type 0x6)
 - MP2MP-up (type 0x7)
 - MP2MP-down (type 0x8)
- Typed Wildcard FEC Element (type 0x5 defined in [RFC5918])

In the case of the Typed Wildcard FEC Element, the FEC Element type **MUST** be one of the MP FECs listed above.

This specification allows the use of topology-scoped mLDP FECs in LDP labels and notification messages, as applicable.

[RFC6514] defines the PMSI tunnel attribute for MVPN and specifies that:

- when the Tunnel Type is set to mLDP P2MP LSP, the Tunnel Identifier is a P2MP FEC Element, and
- when the Tunnel Type is set to mLDP MP2MP LSP, the Tunnel Identifier is an MP2MP FEC Element.

When the extension defined in this specification is in use, the IP MT-Scoped MP FEC Element form of the respective FEC elements **MUST** be used in these two cases.

3.2. Topology IDs

This document assumes the same definitions and procedures associated with MPLS MT-ID as specified in [RFC7307].

4. MT Multipoint Capability

The "MT Multipoint" capability is a new LDP capability, defined in accordance with the LDP capability definition guidelines outlined in [RFC5561]. An mLDP speaker advertises this capability to its peers to announce its support for MTR and the procedures specified in this document. This capability **MAY** be sent either in an Initialization message at session establishment or dynamically during the session's lifetime via a Capability message, provided that the "Dynamic Announcement" capability from [RFC5561] has been successfully negotiated with the peer.

The format of this capability is as follows:

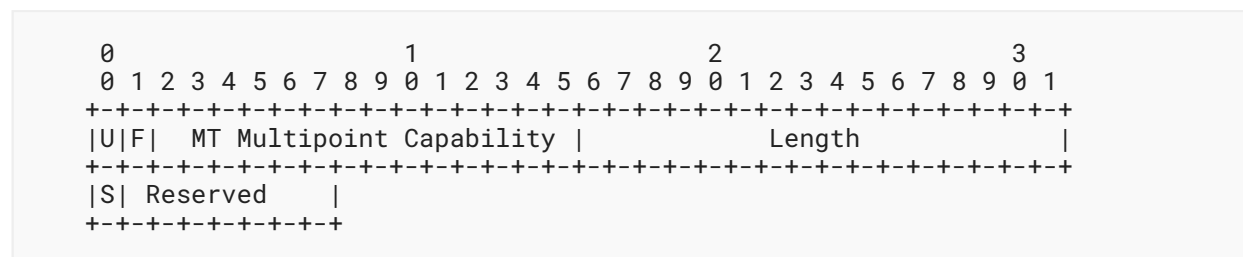


Figure 4: Data Format for the MT Multipoint Capability TLV

Where:

U and F bits: **MUST** be 1 and 0, respectively, as per [Section 3](#) of [RFC5561].

MT Multipoint Capability: The TLV type.

Length: This field specifies the length of the TLV in octets. The value of this field **MUST** be 1, as there is no capability-specific data [RFC5561] following the TLV.

S bit: Set to 1 to announce and 0 to withdraw the capability (as per [RFC5561]).

An mLDP speaker that has successfully advertised and negotiated the "MT Multipoint" capability **MUST** support the following:

1. Topology-scoped mLDP FECs in LDP messages ([Section 3.1](#))
2. Topology-scoped mLDP forwarding setup ([Section 6](#))

5. MT Applicability on FEC-Based Features

5.1. Typed Wildcard MP FEC Elements

[RFC5918] extends the base LDP and defines the Typed Wildcard FEC Element framework. A Typed Wildcard FEC element can be used in any LDP message to specify a wildcard operation for a given type of FEC.

The MT extensions defined in this document do not require any extension to procedures for support of the Typed Wildcard FEC Element [RFC5918], and these procedures apply as is to Multipoint MT FEC wildcarding. Similar to the Typed Wildcard MT Prefix FEC Element, as defined in [RFC7307], the MT extensions allow the use of "MT IP" or "MT IPv6" in the "Address Family" field of the Typed Wildcard MP FEC element. This is done in order to use wildcard operations for MP FECs in the context of a given (sub-)topology as identified by the "MT-ID" and "IPA" fields.

This document defines the following format and encoding for a Typed Wildcard MP FEC element:

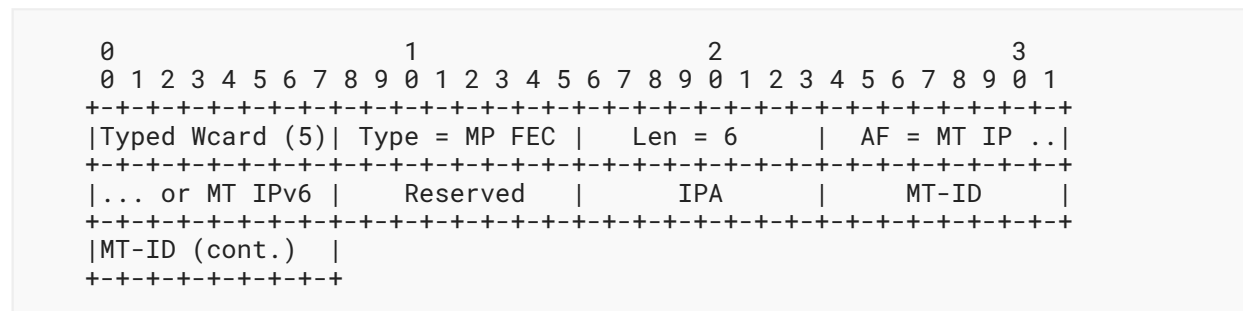


Figure 5: Data Format for the Typed Wildcard MT MP FEC Element

Where:

Type: One of the MP FEC Element types (P2MP, MP2MP-up, or MP2MP-down)

MT-ID: MPLS MT-ID

IPA: The IGP Algorithm

The defined format allows a Label Switching Router (LSR) to perform wildcard MP FEC operations under the scope of a (sub-)topology.

5.2. End-of-LIB

[RFC5919] specifies extensions and procedures that allow an LDP speaker to signal its End-of-LIB (Label Information Base) for a given FEC type to a peer. By leveraging the End-of-LIB message, LDP ensures that label distribution remains consistent and reliable, even during network

disruptions or maintenance activities. The MT extensions for MP FEC do not require any modifications to these procedures and apply as they are to MT MP FEC elements. Consequently, an MT mLDP speaker **MAY** signal its convergence per (sub-)topology using the MT Typed Wildcard MP FEC element.

6. Topology-Scoped Signaling and Forwarding

Since the {MT-ID, IPA} tuple is part of an mLDP FEC, there is no need to support the concept of multiple (sub-)topology forwarding tables in mLDP. Each MP LSP will be unique due to the tuple being part of the FEC. There is also no need to have specific label forwarding tables per topology, and each MP LSP will have its own unique local label in the table. However, in order to implement MTR in an mLDP network, the selection procedures for an upstream LSR and a downstream forwarding interface need to be changed.

6.1. Upstream LSR Selection

The procedures described in [Section 2.4.1.1](#) of [\[RFC6388\]](#) depend on the best path to reach the root. When the {MT-ID, IPA} tuple is signaled as part of the FEC, the tuple is also used to select the (sub-)topology that must be used to find the best path to the root address. Using the next-hop from this best path, an LDP peer is selected following the procedures defined in [\[RFC6388\]](#).

6.2. Downstream Forwarding Interface Selection

[Section 2.4.1.2](#) of [\[RFC6388\]](#) describes the procedures for how a downstream forwarding interface is selected. In these procedures, any interface leading to the downstream LDP neighbor can be considered to be a candidate forwarding interface. When the {MT-ID, IPA} tuple is part of the FEC, this is no longer true. An interface must only be selected if it is part of the same (sub-)topology that was signaled in the mLDP FEC element. Besides this restriction, the other procedures in [\[RFC6388\]](#) apply.

7. LSP Ping Extensions

[\[RFC6425\]](#) defines procedures to detect data plane failures in multipoint MPLS LSPs. [Section 3.1.2](#) of [\[RFC6425\]](#) defines new sub-types and sub-TLVs for Multipoint LDP FECs to be sent in the "Target FEC Stack" TLV of an MPLS Echo Request message [\[RFC8029\]](#).

To support LSP ping for MT MP LSPs, this document uses existing sub-types "P2MP LDP FEC Stack" and "MP2MP LDP FEC Stack" defined in [\[RFC6425\]](#). The LSP ping extension is to specify "MT IP" or "MT IPv6" in the "Address Family" field, set the "Address Length" field to 8 (for MT IP) or 20 (for MT IPv6), and encode the sub-TLV with additional {MT-ID, IPA} information as an extension to the "Root LSR Address" field. The resultant format of sub-TLV is as follows:




Figure 6: Multipoint LDP FEC Stack Sub-TLV Format for MT

The rules and procedures of using this new sub-TLV in an MPLS Echo Request message are the same as defined for the P2MP/MP2MP LDP FEC Stack sub-TLV in [RFC6425]. The only difference is that the "Root LSR Address" field is now (sub-)topology scoped.

8. Security Considerations

This extension to mLDP does not introduce any new security considerations beyond what is already applied to the base LDP specification [RFC5036], the LDP extensions for Multi-Topology specification [RFC7307], the base mLDP specification [RFC6388], and the MPLS security framework specification [RFC5920].

9. IANA Considerations

This document defines a new LDP capability parameter TLV called the "MT Multipoint Capability". IANA has assigned the value 0x0510 from the "TLV Type Name Space" registry in the "Label Distribution Protocol (LDP) Parameters" group as the new code point.

Value	Description	Reference	Notes/Registration Date
0x0510	MT Multipoint Capability	RFC 9658	

Table 1: MT Multipoint Capability

10. References

10.1. 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>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <<https://www.rfc-editor.org/info/rfc4915>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <<https://www.rfc-editor.org/info/rfc6388>>.
- [RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", RFC 6425, DOI 10.17487/RFC6425, November 2011, <<https://www.rfc-editor.org/info/rfc6425>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [RFC7307] Zhao, Q., Raza, K., Zhou, C., Fang, L., Li, L., and D. King, "LDP Extensions for Multi-Topology", RFC 7307, DOI 10.17487/RFC7307, July 2014, <<https://www.rfc-editor.org/info/rfc7307>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [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>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.

10.2. Informative References

- [IANA-IGP-ALGO-TYPES] IANA, "IGP Algorithm Types", <<https://www.iana.org/assignments/igp-parameters>>.

- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, DOI 10.17487/RFC5561, July 2009, <<https://www.rfc-editor.org/info/rfc5561>>.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", RFC 5918, DOI 10.17487/RFC5918, August 2010, <<https://www.rfc-editor.org/info/rfc5918>>.
- [RFC5919] Asati, R., Mohapatra, P., Chen, E., and B. Thomas, "Signaling LDP Label Advertisement Completion", RFC 5919, DOI 10.17487/RFC5919, August 2010, <<https://www.rfc-editor.org/info/rfc5919>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.

Contributors

Anuj Budhiraja
Cisco Systems

Acknowledgments

The authors would like to acknowledge Eric Rosen for his input on this specification.

Authors' Addresses

Ijsbrand Wijnands
Individual
Email: ice@braindump.be

Mankamana Mishra (EDITOR)
Cisco Systems, Inc.
821 Alder Drive
Milpitas, CA 95035
United States of America
Email: mankamis@cisco.com

Kamran Raza
Cisco Systems, Inc.
2000 Innovation Drive
Kanata ON K2K-3E8
Canada
Email: skraza@cisco.com

Zhaohui Zhang

Juniper Networks

10 Technology Park Dr.

Westford, MA 01886

United States of America

Email: zzhang@juniper.net**Arkadiy Gulko**

Edward Jones Wealth Management

United States of America

Email: Arkadiy.gulko@edwardjones.com