

Workgroup: Networking Working Group
Internet-Draft:
draft-ppsenak-lsr-igp-ureach-prefix-
announce-04

Published: 22 August 2023

Intended Status: Standards Track

Expires: 23 February 2024

Authors: P. Psenak, Ed. C. Filsfils S. Litkowski
Cisco Systems Cisco Systems Cisco Systems
D. Voyer A. Dhamija S. Hegde
Bell Canada Rakuten Juniper Networks, Inc.
G. Van de Velde G. Mishra
Nokia Verizon Inc.

IGP Unreachable Prefix Announcement

Abstract

In the presence of summarization, there is a need to signal loss of reachability to an individual prefix covered by the summary in order to enable fast convergence away from paths to the node which owns the prefix which is no longer reachable.

This document describes how to use the existing protocol mechanisms in IS-IS and OSPF, together with the two new flags, to advertise such prefix reachability loss.

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.

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 23 February 2024.

Copyright Notice

Copyright (c) 2023 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. Supporting UPA in IS-IS](#)
 - [2.1. Advertisement of UPA in IS-IS](#)
 - [2.2. Propagation of UPA in IS-IS](#)
- [3. Supporting UPA in OSPF](#)
 - [3.1. Advertisement of UPA in OSPF](#)
 - [3.2. Propagation of UPA in OSPF](#)
- [4. Generation of the UPA](#)
- [5. Signaling UPA](#)
 - [5.1. Signaling UPA in IS-IS](#)
 - [5.2. Signaling UPA in OSPFv2](#)
 - [5.3. Signaling UPA in OSPFv3](#)
 - [5.4. Treatment of the U-Flag and UP-Flag](#)
- [6. Deployment Considerations for UPA](#)
- [7. Processing of the UPA](#)
- [8. IANA Considerations](#)
 - [8.1. IS-IS Prefix Attribute Flags Sub-TLV](#)
 - [8.2. OSPFv2 Extended Prefix TLV Flags](#)
 - [8.3. OSPFv3 Extended-LSA Sub-TLVs](#)
 - [8.4. OSPFv3 Extended Prefix TLV Flags](#)
- [9. Security Considerations](#)
- [10. Acknowledgements](#)
- [11. References](#)
 - [11.1. Normative References](#)
 - [11.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Link-state IGP protocols like IS-IS and OSPF are primarily used to distribute routing information between routers belonging to a single Autonomous System (AS) and to calculate the reachability for IPv4 or IPv6 prefixes advertised by the individual nodes inside the AS. Each node advertises the state of its local adjacencies, connected prefixes, capabilities, etc. The collection of these states from all the routers inside the area form a link-state database (LSDB) that describes the topology of the area and holds additional state information about the prefixes, router capabilities, etc.

The growth of networks running a link-state routing protocol results in the addition of more state which leads to scalability and convergence challenges. The organization of networks into levels/areas and IGP domains helps limit the scope of link-state information within certain boundaries. However, the state related to prefix reachability often requires propagation across a multi-area/level and/or multi-domain IGP network. Techniques such as summarization have been used traditionally to address the scale challenges associated with advertising prefix state outside of the local area/domain. However, this results in suppression of the individual prefix state that is useful for triggering fast-convergence mechanisms outside of the IGPs - e.g., BGP PIC Edge [I-D.ietf-rtgwg-bgp-pic].

This document defines two new flags in IS-IS and OSPF. These flags, together with the existing protocol mechanisms, provide the support for the necessary functionality. The functionality being described is called Unreachable Prefix Announcement (UPA).

Similarly, when an egress router needs to be taken out for maintenance, the traffic is drained from the node before taking it down. This is typically achieved by setting OL-bit together with using max-metric for all prefixes advertised by the node in IS-IS, or by setting max-metric on all-links and prefixes advertised by the node in OSPF. When prefixes from such node are summarized by the ABR/ASBR, nodes outside of the area or domain are unaware of such prefixes becoming unreachable. This document proposes protocol extensions to carry information about such prefixes in a backward compatible manner.

2. Supporting UPA in IS-IS

[[RFC5305](#)] defines the encoding for advertising IPv4 prefixes using 4 octets of metric information. Section 4 specifies:

"If a prefix is advertised with a metric larger than MAX_PATH_METRIC (0xFE000000, see paragraph 3.0), this prefix MUST NOT be considered

during the normal SPF computation. This allows advertisement of a prefix for purposes other than building the normal IP routing table.
"

Similarly, [[RFC5308](#)] defines the encoding for advertising IPv6 prefixes using 4 octets of metric information. Section 2 states:

"...if a prefix is advertised with a metric larger than MAX_V6_PATH_METRIC (0xFE000000), this prefix MUST NOT be considered during the normal Shortest Path First (SPF) computation. This will allow advertisement of a prefix for purposes other than building the normal IPv6 routing table."

This functionality can be used to advertise a prefix (IPv4 or IPv6) in a manner which indicates that reachability has been lost - and to do so without requiring all nodes in the network to be upgraded to support the functionality.

2.1. Advertisement of UPA in IS-IS

Existing nodes in a network that do not support UPA will not use UPAs during the route calculation, but will continue to flood them. This allows flooding of such advertisements to occur without the need to upgrade all nodes in a network.

Recognition of the advertisement as UPA is only required on routers which have a use case for this information. Area Border Routers (ABRs), which would be responsible for propagating UPA advertisements into other areas would need to recognize such advertisements.

As per the definitions referenced in the preceding section, any prefix advertisement with a metric value greater than 0xFE000000 can be used for purposes other than normal routing calculations. Such an advertisement can be interpreted by the receiver as a UPA.

Optionally, an implementation may use local configuration to limit the set of metric values which will be interpreted as UPA. The only restriction is that such values MUST be greater than 0xFE000000.

UPA in IS-IS is supported for all IS-IS Sub-TLVs Advertising Prefix Reachability, e.g., SRv6 Locator

[[I-D.ietf-lsr-isis-srv6-extensions](#)], Extended IP reachability [[RFC5305](#)], MT IP Reach [[RFC5120](#)], IPv6 IP Reach [[RFC5308](#)], MT IPv6 IP Reach [[RFC5120](#)], IPv4 Algorithm Prefix Reachability TLV [[I-D.ietf-lsr-ip-flexalgo](#)], and IPv6 Algorithm Prefix Reachability TLV [[I-D.ietf-lsr-ip-flexalgo](#)]

2.2. Propagation of UPA in IS-IS

IS-IS allows propagation of IP prefixes in both directions between level 1 and level 2. For reachable prefixes this is only done if the prefix is reachable in source level - e.g., the prefix needs to be reachable in level 1 to be propagated to level 2 and vice versa. Such requirement of reachability MUST NOT be applied for UPAs, as they are propagating unreachability.

IS-IS L1/L2 routers may wish to advertise received UPAs into other areas (upwards and/or downwards). When propagating UPAs the original metric value MUST be preserved. The cost to reach the originator of the received UPA MUST NOT be considered when readvertising the UPA.

3. Supporting UPA in OSPF

[[RFC2328](#)] Appendix B defines the following architectural constant for OSPF:

"LSInfinity The metric value indicating that the destination described by an LSA is unreachable. Used in summary-LSAs and AS-external-LSAs as an alternative to premature aging (see Section 14.1). It is defined to be the 24-bit binary value of all ones: 0xffffffff."

[[RFC5340](#)] Appendix B states:

"Architectural constants for the OSPF protocol are defined in Appendix B of OSPFv2."

indicating that these same constants are applicable to OSPFv3.

[[RFC2328](#)] section 14.1. also describes the usage of LSInfinity as a way to indicate loss of prefix reachability:

"Premature aging can also be used when, for example, one of the router's previously advertised external routes is no longer reachable. In this circumstance, the router can flush its AS-external-LSA from the routing domain via premature aging. This procedure is preferable to the alternative, which is to originate a new LSA for the destination specifying a metric of LSInfinity."

In addition, NU-bit is defined for OSPFv3 [[RFC2328](#)]. Prefixes having the NU-bit set in their PrefixOptions field should be ignored by the route calculation.

UPA in OSPFv2 is supported for OSPFv2 Summary-LSA [[RFC2328](#)], AS-external-LSAs [[RFC2328](#)], NSSA AS-external LSA [[RFC3101](#)], and OSPFv2 Extended Prefix TLV [[I-D.ietf-lsr-ip-flexalgo](#)].

UPA in OSPFv3 is supported for Inter-Area-Prefix-LSA [[RFC5340](#)], AS-External-LSA [[RFC5340](#)], NSSA-LSA [[RFC5340](#)], E-Inter-Area-Prefix-LSA [[RFC8362](#)], E-AS-External-LSA [[RFC8362](#)], E-Type-7-LSA [[RFC8362](#)], and SRV6 Locator LSA [[I-D.ietf-lsr-ospfv3-srv6-extensions](#)].

3.1. Advertisement of UPA in OSPF

Using the existing mechanism already defined in the standards, as described in previous section, an advertisement of the inter-area or external prefix inside OSPF or OSPFv3 LSA that has the age set to value lower than MaxAge and metric set to LSInfinity can be interpreted by the receiver as a UPA.

Existing nodes in a network which receive UPA advertisements will propagate it following existing standard procedures defined by OSPF.

OSPF Area Border Routers (ABRs), which would be responsible for propagating UPA advertisements into other areas would need to recognize such advertisements.

3.2. Propagation of UPA in OSPF

Advertising prefix reachability between OSPF areas assumes prefix reachability in a source area. Such requirement of reachability MUST not be applied for UPAs, as they are propagating unreachability.

OSPF ABRs may wish to advertise received UPAs into other connected areas. When doing so, the original LSInfinity metric value in UPA MUST be preserved. The cost to reach the originator of the received UPA MUST NOT be considered when readvertising the UPA to connected areas.

4. Generation of the UPA

UPA MAY be generated by the ABR or ASBR that is performing the summarization, when all of the following conditions are met:

- reachability of a prefix that was reachable earlier was lost
- a summary address which covers the prefix is being advertised by the ABR/ASBR

Implementations are free to limit the UPA generation to specific prefixes, e.g. host prefixes, SRv6 locators, or similar. Such filtering is optional and MAY be controlled via configuration.

5. Signaling UPA

In IS-IS a prefix can be advertised with metric higher than 0xFE000000, in OSPF with metric LSInfinity, or in OSPFv3 with NU-bit

set in PrefixOptions, for various reasons. Even though in all cases the treatment of such metric, or NU-bit, is specified for IS-IS, OSPF and OSPFv3, having an explicit way to signal that the prefix was advertised in order to signal unreachability is required to distinguish it from other cases where the prefix with such metric is advertised.

5.1. Signaling UPA in IS-IS

Two new bits in the IPv4/IPv6 Extended Reachability Attribute Flags [[RFC7794](#)] are defined:

U-Flag: - Unreachable Prefix Flag (Bit TBD). When set, it indicates that the prefix is unreachable due to the unplanned reason.

UP-Flag: - Unreachable Planned Prefix Flag (Bit TBD). When set, it indicates that the prefix is unreachable due to the planned reason, e.g., planned maintenance.

The prefix that is advertised with U-Flag or UP-Flag MUST have the metric set to a value larger than 0xFE000000. If the prefix metric is less than or equal 0xFE000000, both of these flags MUST be ignored.

5.2. Signaling UPA in OSPFv2

Two new bits in Flags field of the OSPFv2 Extended Prefix TLV [[RFC7684](#)] are defined:

U-Flag: - Unreachable Prefix Flag (Bit TBD). When set, it indicates that the prefix is unreachable due to the unplanned reason.

UP-Flag: - Unreachable Planned Prefix Flag (Bit TBD). When set, it indicates that the prefix is unreachable due to the planned reason, e.g., planned maintenance.


The prefix that is advertised with U-Flag or UP-Flag MUST have the metric set to a value LSInfinity. If the prefix metric is not equal to LSInfinity, both of these flags MUST be ignored. For default algorithm 0 prefixes with U-Flag or UP-Flag, it is therefore REQUIRED to advertise the unreachable prefix in the base OSPFv2 LSA - e.g., OSPFv2 Summary-LSA [[RFC2328](#)], or AS-external-LSAs [[RFC2328](#)], or NSSA AS-external LSA [[RFC3101](#)].

5.3. Signaling UPA in OSPFv3

A new Prefix Attribute Flags sub-TLV is defined as a sub-TLV of the following OSPFv3 TLVs as defined in [[RFC8362](#)]:

Inter-Area Prefix TLV

It MUST appear only once in the parent TLV and has the following format:



where:

Type: TBD

Length: variable, dependent on the size of the Prefix Attribute Flags. MUST be a multiple of 4 octets.

Prefix Attribute Flags: individual bits representing the flags.

Bits in Prefix Attribute Flags are defined/sent starting with Bit 0 defined below. Additional bit definitions that may be defined in the future SHOULD be assigned in ascending bit order so as to minimize the number of bits that will need to be transmitted.

Undefined bits MUST be transmitted as 0 and MUST be ignored on receipt.

Bits that are NOT transmitted MUST be treated as if they are set to 0 on receipt.

Two new bits in OSPFv3 Prefix Attribute Flags are defined:

U-Flag: - Unreachable Prefix Flag (Bit TBD). When set, it indicates that the prefix became unreachable due to the unplanned reason.

UP-Flag: - Unreachable Planned Prefix Flag (Bit TBD). When set, it indicates that the prefix became unreachable due to the planned reason, e.g., planned maintenance.

The prefix that is advertised with U-Flag or UP-flag MUST have the metric set to a value LSInfinity. For default algorithm 0 prefixes,

the LSInfinity MUST be set in the parent TLV. For IP Algorithm Prefixes [[I-D.ietf-lsr-ip-flexalgo](#)], the LSInfinity MUST be set in OSPFv3 IP Algorithm Prefix Reachability sub-TLV. If the prefix metric is not equal to LSInfinity, both of these flags MUST be ignored.

The prefix that is advertised with U-Flag or UP-Flag MUST have the NU-bit set in the PrefixOptions of the parent TLV. If the NU-bit in PrefixOptions of the parent TLV is not set, both of these flags MUST be ignored.

5.4. Treatment of the U-Flag and UP-Flag

The setting of the U-Flag or the UP-Flag signals that the prefix is unreachable. They constitute the UPA signals. Treatment of these flags on the receiver is optional and the usage of them is outside of scope of this document.

6. Deployment Considerations for UPA

The intent of UPA is to provide an event driven signal of the transition of a destination from reachable to unreachable. It is not intended to advertise a persistent state. UPA advertisements should therefore be withdrawn after a modest amount of time, that would provides sufficient time for UPA to be flooded network-wide and acted upon by receiving nodes, but limits the presence of UPA in the network to a short time period. The time the UPA is kept in the network SHOULD also reflect the intended use-case for which the UPA was advertised.

As UPA advertisements in IS-IS are advertised in existing Link State PDUs (LSPs) and the unit of flooding in IS-IS is an LSP, it is recommended that, when possible, UPAs are advertised in LSPs dedicated to this type of advertisement. This will minimize the number of LSPs which need to be updated when UPAs are advertised and withdrawn.

In OSPF and OSPFv3, each inter-area and external prefix is advertised in it's own LSA, so the above optimisation does not apply to OSPF.

It is also recommended that implementations limit the number of UPA advertisements which can be originated at a given time.

UPA is not meant to address an area/domain partition. When an area partitions, while multiple ABRs advertise the same summary, each of the ABRs can only reach portion of the summarized prefix. As a result, depending on which ABR the traffic is using to enter a partitioned area, the traffic could be dropped or be delivered to its final destination. UPA does not make the problem of an area

partition any worse. In case of an area partition each of the ABRs will generate UPAs for the destinations for which the reachability was lost locally. As the UPA propagates to the nodes outside of a partitioned area, it may result in such nodes picking an alternative egress node for the traffic, if such alternate egress node exists. If such alternate egress node resides outside of a partitioned area, traffic will be restored. If such alternate egress node resides in a partitioned area and is covered by the summary, the traffic will be dropped if it enters a partitioned area via the ABR that can not reach the alternate egress node - resulting in similar behavior as without the UPA. Above is similarly applicable to a domain partition.

7. Processing of the UPA

Processing of the received UPAs is optional and SHOULD be controlled by the configuration at the receiver. The receiver itself, based on its configuration, decides what the UPA will be used for and what applications, if any, will be notified when UPA is received.

8. IANA Considerations

8.1. IS-IS Prefix Attribute Flags Sub-TLV

This document adds two new bits in the "IS-IS Bit Values for Prefix Attribute Flags Sub-TLV" registry:

Bit #: TBD

Description: U-Flag

Reference: This document ([Section 5.1](#)).

Bit #: TDB

Description: UP-Flag

Reference: This document ([Section 5.1](#)).

8.2. OSPFv2 Extended Prefix TLV Flags

This document adds two new bits in the "OSPFv2 Extended Prefix TLV Flags" registry:

Bit #: TBD

Description: U-Flag

Reference: This document ([Section 5.2](#)).

Bit #: TDB

Description: UP-Flag

Reference: This document ([Section 5.2](#)).

8.3. OSPFv3 Extended-LSA Sub-TLVs

This document adds new Sub-TLV in the "OSPFv3 Extended-LSA Sub-TLVs" registry:

Value: TBD

Description: Prefix Attribute Flags sub-TLV

Reference: This document ([Section 5.3](#)).

The L2BM column is set to X.

8.4. OSPFv3 Extended Prefix TLV Flags

This document defines a new "OSPFv3 Extended Prefix TLV Flags" registry under "Open Shortest Path First v3 (OSPFv3) Parameters".

This document adds two new bits in the newly defined "OSPFv3 Extended Prefix TLV Flags" registry:

Bit #: TBD

Description: U-Flag

Reference: This document ([Section 5.3](#)).

Bit #: TDB

Description: UP-Flag

Reference: This document ([Section 5.3](#)).

9. Security Considerations

The use of UPAs introduces the possibility that an attacker could inject a false, but apparently valid, UPA. However, the risk of this

occurring is no greater than the risk today of an attacker injecting any other type of false advertisement .

The risks can be reduced by the use of existing security extensions as described in [RFC5304] and [RFC5310] for IS-IS, in [RFC2328] and [RFC7474] for OSPFv2, and in [RFC5340] and [RFC4552] for OSPFv3.

10. Acknowledgements

The authors would like to thank Kamran Raza, Michael MacKenzie and Luay Jalil for their contribution and support of the overall solution proposed in this document.

11. References

11.1. Normative References

- [ISO10589] ISO, "Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", November 2002.
- [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>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC3101] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, DOI 10.17487/RFC3101, January 2003, <<https://www.rfc-editor.org/info/rfc3101>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [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>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.

[RFC5305]

Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

[RFC5308]

Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.

[RFC5310]

Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.

[RFC5340]

Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.

[RFC7474]

Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <<https://www.rfc-editor.org/info/rfc7474>>.

[RFC7684]

Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.

[RFC7794]

Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability", RFC 7794, DOI 10.17487/RFC7794, March 2016, <<https://www.rfc-editor.org/info/rfc7794>>.

[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>>.

[RFC8362]

Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.

11.2. Informative References

[I-D.ietf-lsr-ip-flexalgo] Britto, W., Hegde, S., Kaneriy, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithms (Flex-Algorithm) In IP Networks", Work in Progress, Internet-Draft, draft-ietf-lsr-ip-flexalgo-17,

24 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-ip-flexalgo-17>>.

[I-D.ietf-lsr-isis-srv6-extensions]

Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-srv6-extensions-19, 14 November 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-srv6-extensions-19>>.

[I-D.ietf-lsr-ospfv3-srv6-extensions] Li, Z., Hu, Z., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-ospfv3-srv6-extensions-15, 21 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-ospfv3-srv6-extensions-15>>.

Authors' Addresses

Peter Psenak (editor)
Cisco Systems
Pribinova Street 10
Bratislava 81109
Slovakia

Email: ppsenak@cisco.com

Clarence Filsfils
Cisco Systems
Brussels
Belgium

Email: cfilsfil@cisco.com

Stephane Litkowski
Cisco Systems
La Rigourdiere
Cesson Sevine
France

Email: slitkows@cisco.com

Daniel Voyer
Bell Canada

Email: daniel.voyer@bell.ca

Amit Dhamija
Rakuten

Email: amit.dhamija@rakuten.com

Shraddha Hegde
Juniper Networks, Inc.
Embassy Business Park
Bangalore, KA
560093
India

Email: shraddha@juniper.net

Gunter Van de Velde
Nokia
Antwerp
Belgium

Email: gunter.van_de_velde@nokia.com

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com