

Internet Engineering Task Force (IETF)
Request for Comments: 9929
Category: Standards Track
ISSN: 2070-1721

P. Psenak, Ed.
C. Filsfils
D. Voyer
Cisco Systems
S. Hegde
HPE
G. Mishra
Verizon Inc.
February 2026

IGP Unreachable Prefix Announcement

Abstract

Summarization is often used in multi-area or multi-domain networks to improve network efficiency and scalability. With summarization in place, there is a need to signal loss of reachability to an individual prefix covered by the summary. This enables fast convergence by steering traffic, when applicable, away from the node which owns the prefix and is no longer reachable.

This document specifies protocol mechanisms in IS-IS and OSPF, together with two new flags, to advertise such prefix reachability loss.

The term "OSPF" in this document is used to refer to both OSPFv2 and OSPFv3.

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/rfc9929>.

Copyright Notice

Copyright (c) 2026 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
 - 1.1. Requirements Language

- 2. Generation of the UPA
 - 3. Supporting UPA in IS-IS
 - 3.1. Advertisement of UPA in IS-IS
 - 3.2. Signaling UPA in IS-IS
 - 3.3. Propagation of UPA in IS-IS
 - 4. Supporting UPA in OSPF
 - 4.1. Advertisement of UPA in OSPF
 - 4.2. Signaling UPA in OSPF
 - 4.2.1. Signaling UPA in OSPFv2
 - 4.2.2. Signaling UPA in OSPFv3
 - 4.3. Propagation of UPA in OSPF
 - 5. Processing of the UPA
 - 6. Area and Domain Partition
 - 7. IANA Considerations
 - 7.1. IS-IS Prefix Attribute Flags Sub-TLV
 - 7.2. OSPFv2 and OSPFv3 Prefix Extended Flags
 - 8. Security Considerations
 - 9. References
 - 9.1. Normative References
 - 9.2. Informative References
- Acknowledgements
Contributors
Authors' Addresses

1. Introduction

Link-state Interior Gateway Protocols (IGPs) like Intermediate System to Intermediate System (IS-IS) [ISO10589], Open Shortest Path First version 2 (OSPFv2) [RFC2328], and Open Shortest Path First version 3 (OSPFv3) [RFC5340] 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. IGP summarization is a network engineering technique for combining multiple smaller, contiguous IP networks into a single, larger summary route. Techniques such as summarization have been used to address the scaling 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., Border Gateway Protocol (BGP) Prefix-Independent Convergence (PIC) [BGP-PIC].

Similarly, when a router needs to be taken out of service for maintenance, the traffic is drained from the node before taking it down. This is typically achieved by setting the OVERLOAD bit together with using a high metric for all prefixes advertised by the node in IS-IS. The mechanisms available for that purpose are (in OSPFv2) using the cost of MaxLinkMetric for all non-stub links in the router-LSA [RFC6987] or using the H-bit [RFC8770], or (in OSPFv3 [RFC5340]) using the R-bit.

When prefixes from such nodes are summarized by an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR), nodes outside of

the area or domain are unaware of these summarized prefixes becoming unreachable. This document proposes protocol extensions to carry information about such prefixes in a backward-compatible manner.

This document does not define how to advertise a prefix that is not reachable for routing. That has been defined for IS-IS in [RFC5305] and [RFC5308], for OSPFv2 in [RFC2328], and for OSPFv3 in [RFC5340].

This document defines a method to signal a specific reason for which the prefix was advertised with the metric that excludes it from the route calculation. This is done to distinguish it from any other possible cases, where such metric advertisement may be used.

IGPs typically only advertise the reachability of the prefix. A prefix that was previously advertised as reachable is made unreachable just by withdrawing the previous advertisement of the prefix. Some of the use cases mentioned earlier in this section require that unreachability be signaled for a prefix for which the reachability was not explicitly signaled previously, because it was covered by the reachability of the summary prefix.

This document defines two new flags in IS-IS, OSPFv2, and OSPFv3. These flags provide the support for advertising prefix unreachability, together with the reason for which the unreachability is advertised. The functionality being described is called Unreachable Prefix Announcement (UPA).

This document also defines how the UPA is propagated across IS-IS levels and OSPF areas.

The term "OSPF" in this document is used to cover both OSPFv2 and OSPFv3 protocols.

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. Generation of the UPA

UPA MAY be generated by an ABR or ASBR for a prefix that is summarized by the summary prefix originated by an ABR or ASBR in the following cases:

1. A prefix that was reachable earlier becomes unreachable.
2. For any of the planned maintenance cases:
 - * the node originating the prefix is signaling the overload state in IS-IS, or the H-bit in OSPFv2 [RFC8770], or the R-bit in OSPFv3 [RFC5340], or
 - * the metric to reach the prefix from an ABR or ASBR crosses the configured threshold.

Generation as well as propagation of the UPA at an ABR or ASBR is optional and SHOULD be controlled by a configuration knob. It SHOULD be disabled by default.

Implementations MAY limit the UPA generation as well as propagation to specific prefixes, e.g. host prefixes, Segment Routing over IPv6 (SRv6) locators, or similar. Such filtering is optional and SHOULD be controlled via configuration.

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.

ABR or ASBR MUST withdraw the previously advertised UPA when the reason for which the UPA was generated ceases, e.g., prefix reachability was restored or its metric has changed such that it is below a configured threshold value.

Even if the reasons persist, UPA advertisements SHOULD be withdrawn after some amount of time, that would provide sufficient time for UPA to be flooded network-wide and acted upon by receiving nodes, but limits the presence of UPA in the network. The time the UPA is kept in the network SHOULD also reflect the intended use case for which the UPA was advertised. Not withdrawing the UPA would result in stale information being kept in the link state databases of all routers in the area.

Implementations SHOULD provide a configuration option to specify the UPA lifetime at the originating ABR or ASBR.

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 that need to be updated when UPAs are advertised and withdrawn.

In OSPFv2 and OSPFv3, each inter-area and external prefix is advertised in its own LSA, so the above consideration does not apply to OSPFv2 and OSPFv3.

It is also RECOMMENDED that implementations limit the number of UPA advertisements that can be originated at a given time to limit the number of UPAs present in the network at any given point of time. UPA implementations SHOULD provide a configuration option to limit the number of such UPAs.

3. Supporting UPA in IS-IS

[RFC5305] defines the encoding for advertising IPv4 prefixes using 4 octets of metric information, and Section 4 of [RFC5305] 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 and Section 2 of [RFC5308] 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 that indicates that reachability has been lost -- and to do so without requiring all nodes in the network to be upgraded to support the functionality.

3.1. 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 within the area. This allows flooding of such advertisements to occur without the need to upgrade all nodes in a network to support this specification.

Those ABRs or ASBRs that are responsible for propagating UPA advertisements into other areas or domains are also expected to recognize UPA advertisements.

As per the definitions referenced in Section 3, any prefix advertisement with a metric value greater than 0xFE000000 can be used for purposes other than normal routing calculations. Such a metric MUST be used when advertising UPA in IS-IS.

[RFC7370] introduced the "IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability" registry, which lists TLVs for advertising different types of prefix reachability. (The list at the time of publication of this document is below.) UPA in IS-IS is supported for prefixes advertised in all such TLVs identified by that registry, for example:

- * SRv6 Locator [RFC9352]
- * Extended IP reachability [RFC5305]
- * Multi-Topology (MT) IP Reach [RFC5120]
- * IPv6 IP Reach [RFC5308]
- * MT IPv6 IP Reach [RFC5120]
- * IPv4 Algorithm Prefix Reachability TLV [RFC9502]
- * IPv6 Algorithm Prefix Reachability TLV [RFC9502]

3.2. Signaling UPA in IS-IS

In IS-IS, a prefix can be advertised with a metric higher than 0xFE000000, for various reasons. While IS-IS specifies the treatment of such metrics, explicit signaling is required to distinguish UPA from other high-metric advertisements.

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

U-Flag: Unreachable Prefix Flag (bit 5). When set, it indicates that the prefix is unreachable.

UP-Flag: Unreachable Planned Prefix Flag (bit 6). When set, this flag indicates that the prefix is unreachable due to a planned event (e.g., planned maintenance).

The originating node MUST NOT set the UP-flag without setting the U-flag.

The receiving node MUST ignore the UP-flag in the advertisement if the U-flag is not set.

The prefix that is advertised with the U-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.

3.3. Propagation of UPA in IS-IS

IS-IS L1/L2 routers, which would be responsible for propagating UPA

advertisements between levels, need to recognize such advertisements. Failure to do so would prevent UPA from reaching the routers in the remote areas.

IS-IS allows propagation of IP prefixes in both directions between level 1 and level 2. Propagation is only done if the prefix is reachable in the source level, i.e., the prefix is only propagated from a level in which the prefix is reachable. 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.

4. Supporting UPA in OSPF

Appendix B of [RFC2328] defines the following architectural constant for OSPFv2:

```
| 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.
```

Appendix B of [RFC5340] 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, the NU-bit is defined for OSPFv3 [RFC5340]. Prefixes having the NU-bit set in their PrefixOptions field are not included in the routing calculation.

UPA in OSPFv2 is supported for prefix reachability advertised via OSPFv2 Summary-LSA [RFC2328], AS-external-LSAs [RFC2328], Not-So-Stubby Area (NSSA) AS-external-LSA [RFC3101], and OSPFv2 IP Algorithm Prefix Reachability Sub-TLV [RFC9502].

UPA in OSPFv3 is supported for prefix reachability advertised via OSPFv3 E-Inter-Area-Prefix-LSA [RFC8362], E-AS-External-LSA [RFC8362], E-Type-7-LSA [RFC8362], and SRv6 Locator LSA [RFC9513].

For prefix reachability advertised via Inter-Area-Prefix-LSA [RFC5340], AS-External-LSA [RFC5340], NSSA-LSA [RFC5340], UPA is signaled using their corresponding extended LSAs. This requires support of the OSPFv3 Extended LSAs in a sparse mode as specified in Section 6.2 of [RFC8362].

4.1. Advertisement of UPA in OSPF

If an ABR or ASBR advertises UPA in an advertisement of an inter-area or external prefix inside OSPFv2 or OSPFv3, then it MUST set the age to a value lower than MaxAge and set the metric to LSInfinity.

UPA flooding inside the area follows the existing standard procedures defined by OSPFv2 [RFC2328] and OSPFv3 [RFC5340].

4.2. Signaling UPA in OSPF

A prefix can be advertised (in OSPFv2) with the LSInfinity metric or (in OSPFv3) with the NU-bit set in PrefixOptions, for various reasons. While OSPFv2 and OSPFv3 specify the treatment of the LSInfinity metric and the NU-bit, explicit signaling is required to distinguish UPA from other scenarios using the LSInfinity metric or NU-bit.

OSPFv2 and OSPFv3 Prefix Extended Flags Sub-TLVs been defined in [RFC9792] for advertising additional prefix attribute flags in OSPFv2 and OSPFv3.

Two new bits in the Prefix Attribute Flags Sub-TLV are defined:

U-Flag: Unreachable Prefix Flag (bit 0). When set, it indicates that the prefix is unreachable.

UP-Flag: Unreachable Planned Prefix Flag (bit 1). When set, this flag indicates that the prefix is unreachable due to a planned event (e.g., planned maintenance).

The originating node MUST NOT set the UP-flag without setting the U-flag.

The receiving node MUST ignore the UP-flag in the advertisement if the U-flag is not set.

4.2.1. Signaling UPA in OSPFv2

The OSPFv2 Prefix Extended Flags Sub-TLV [RFC9792] is a sub-TLV of the OSPFv2 Extended Prefix TLV [RFC7684].

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. Therefore, for the prefixes in default algorithm 0 that are advertised with U-Flag, or Up-Flag, it is REQUIRED to advertise the unreachable prefix in the base OSPFv2 LSA, e.g., e.g., OSPFv2 Summary-LSA [RFC2328], or AS-external-LSAs [RFC2328], or NSSA AS-external LSA [RFC3101].

4.2.2. Signaling UPA in OSPFv3

OSPFv3 Prefix Extended Flags Sub-TLV is defined as a sub-TLV of the following OSPFv3 TLVs that are defined in [RFC8362]:

- * Intra-Area-Prefix TLV
- * Inter-Area-Prefix TLV
- * External-Prefix TLV

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 [RFC9502], 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.

4.3. Propagation of UPA in OSPF

OSPF ABRs, which would be responsible for propagating UPA advertisements into other areas, need to recognize such advertisements. Failure to do so would prevent UPA from reaching the routers in the remote areas.

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

OSPF ABRs or ASBRs MAY advertise received UPAs between connected areas or domains. 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.

5. 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. Usage of the UPA at the receiver is outside of the scope of this document.

As an example, UPA may be used to trigger BGP PIC Edge at the receiving router [BGP-PIC].

Applications using the UPA cannot use the absence of the UPA to infer that the reachability of the prefix is back. They must rely on their own mechanisms to verify the reachability of the remote endpoints.

6. Area and Domain Partition

UPA is not meant to address an area/domain partition. When an area or domain partitions, while multiple ABRs or ASBRs advertise the same summary, each of them can only reach a portion of the summarized prefix. As a result, depending on which ABR or ASBR the traffic is using to enter a partitioned area, the traffic could be either dropped or delivered to its final destination. UPA does not make the problem of an area partition any worse. In case of an area partition, each ABR or ASBR will generate UPAs for the destinations for which the reachability was lost locally. As the UPA propagates to the nodes outside a partitioned area, it may result in such nodes picking an alternative egress node for the traffic, if such a node exists. If such an alternative egress node resides outside a partitioned area, traffic will be restored. If such an alternative 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 an ABR or ASBR that cannot reach that node. This will result in similar behavior as without the UPA. The above statements are also applicable to a domain partition.

7. IANA Considerations

7.1. 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 #: 5
Name: U-Flag
Reference: RFC 9929 (Section 3.2)

Bit #: 6
Name: UP-Flag
Reference: RFC 9929 (Section 3.2)

7.2. OSPFv2 and OSPFv3 Prefix Extended Flags

This document adds two new bits in the "OSPFv2 Prefix Extended Flags" and "OSPFv3 Prefix Extended Flags" registries:

Bit: 0
Description: U-Flag
Reference: RFC 9929 (Section 4.2)

Bit: 1
Description: UP-Flag
Reference: RFC 9929 (Section 4.2)

8. 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], [RFC5310], and [RFC7794] for IS-IS.
- * [RFC2328], [RFC7474], and [RFC7684] for OSPFv2.
- * [RFC5340], [RFC4552], and [RFC8362] for OSPFv3.

9. References

9.1. Normative References

- [ISO10589] ISO/IEC, "Information technology -- Telecommunications and information exchange between systems -- 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)", ISO/IEC 10589:2002, November 2002, <<https://www.iso.org/en/contents/data/standard/03/09/30932.html>>.
- [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>>.
- [RFC6987] Retana, A., Nguyen, L., Zinin, A., White, R., and D. McPherson, "OSPF Stub Router Advertisement", RFC 6987, DOI 10.17487/RFC6987, September 2013, <<https://www.rfc-editor.org/info/rfc6987>>.
- [RFC7370] Ginsberg, L., "Updates to the IS-IS TLV Codepoints Registry", RFC 7370, DOI 10.17487/RFC7370, September 2014, <<https://www.rfc-editor.org/info/rfc7370>>.
- [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>>.
- [RFC8770] Patel, K., Pillay-Esnault, P., Bhardwaj, M., and S. Bayraktar, "Host Router Support for OSPFv2", RFC 8770, DOI 10.17487/RFC8770, April 2020, <<https://www.rfc-editor.org/info/rfc8770>>.

- [RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.
- [RFC9502] Britto, W., Hegde, S., Kaneriy, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithm in IP Networks", RFC 9502, DOI 10.17487/RFC9502, November 2023, <<https://www.rfc-editor.org/info/rfc9502>>.
- [RFC9513] Li, Z., Hu, Z., Talaulikar, K., Ed., and P. Psenak, "OSPFv3 Extensions for Segment Routing over IPv6 (SRv6)", RFC 9513, DOI 10.17487/RFC9513, December 2023, <<https://www.rfc-editor.org/info/rfc9513>>.
- [RFC9792] Chen, R., Zhao, D., Psenak, P., Talaulikar, K., and L. Gong, "Prefix Flag Extension for OSPFv2 and OSPFv3", RFC 9792, DOI 10.17487/RFC9792, June 2025, <<https://www.rfc-editor.org/info/rfc9792>>.

9.2. Informative References

- [BGP-PIC] Bashandy, A., Ed., Filsfils, C., Mohapatra, P., and Y. Qu, "BGP Prefix Independent Convergence", Work in Progress, Internet-Draft, draft-ietf-rtgwg-bgp-pic-23, 15 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-bgp-pic-23>>.

Acknowledgements

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

Contributors

The following people contributed to the content of this document and should be considered coauthors:

Stephane Litkowski
Email: slitkows@cisco.com

Amit Dhamija
Email: amitd@arrcus.com

Gunter Van de Velde
Email: gunter.van_de_velde@nokia.com

The following people contributed to the problem statement and the solution requirement discussion:

Aijun Wang
Email: wangaj3@chinatelecom.cn

Zhibo Hu
Email: huzhibo@huawei.com

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

Daniel Voyer
Cisco Systems
Email: davoyer@cisco.com

Shraddha Hegde
HPE
Email: shraddha.hegde@hpe.com

Gyan Mishra
Verizon Inc.
Email: hayabusagsm@gmail.com