

Internet Engineering Task Force (IETF)
Request for Comments: 8518
Updates: 5286
Category: Standards Track
ISSN: 2070-1721

P. Sarkar, Ed.
Arcus, Inc.
U. Chunduri, Ed.
Huawei USA
S. Hegde
Juniper Networks, Inc.
J. Tantsura
Apstra, Inc.
H. Gredler
RtBrick, Inc.
March 2019

Selection of Loop-Free Alternates for Multi-Homed Prefixes

Abstract

Deployment experience gained from implementing algorithms to determine Loop-Free Alternates (LFAs) for multi-homed prefixes (MHPs) has revealed some avenues for potential improvement. This document provides explicit inequalities that can be used to evaluate neighbors as potential alternates for MHPs. It also provides detailed criteria for evaluating potential alternates for external prefixes advertised by OSPF ASBRs. This document updates Section 6 of RFC 5286 by expanding some of the routing aspects.

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/rfc8518>.

Copyright Notice

Copyright (c) 2019 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Acronyms	4
1.2. Requirements Language	4
2. LFA Inequalities for MHPs	4
3. LFA Selection for MHPs	6
3.1. Improved Coverage with Simplified Approach to MHPs	7
3.2. IS-IS ATT Bit Considerations	9
4. LFA Selection for Multi-Homed External Prefixes	10
4.1. IS-IS	10
4.2. OSPF	10
4.2.1. Rules to Select Alternate ASBRs	10
4.2.1.1. Multiple ASBRs Belonging to Different Areas ..	12
4.2.1.2. Type 1 and Type 2 Costs	12
4.2.1.3. RFC1583Compatibility is Set to "Enabled"	12
4.2.1.4. Type 7 Routes	13
4.2.2. Inequalities to Be Applied for Alternate ASBR Selection	13
4.2.2.1. Forwarding Address Set to Non-zero Value	13
4.2.2.2. ASBRs Advertising Type 1 and Type 2 Costs	14
5. LFA Extended Procedures	15
5.1. Links with IGP MAX_METRIC	15
5.2. MT Considerations	16
6. IANA Considerations	16
7. Security Considerations	17
8. References	17
8.1. Normative References	17
8.2. Informative References	17
Acknowledgements	19
Contributors	19
Authors' Addresses	20

1. Introduction

A framework for the development of IP Fast Reroute (FRR) mechanisms is detailed in [RFC5714]. The use of LFAs for IP FRR is specified in [RFC5286]. If a prefix is advertised by more than one router, that prefix is called a "multi-homed prefix (MHP)". MHPs generally occur for prefixes obtained from outside the routing domain by multiple routers, for subnets on links where the subnet is announced from multiple ends of the link, and for prefixes advertised by multiple routers to provide resiliency.

Section 6.1 of [RFC5286] describes a method to determine LFAs for MHPs. This document describes a procedure using explicit inequalities that can be used by a computing router to evaluate a neighbor as a potential alternate for an MHP. The results obtained are equivalent to those obtained using the method described in Section 6.1 of [RFC5286].

Section 6.3 of [RFC5286] discusses complications associated with computing LFAs for MHPs in OSPF. This document provides detailed criteria for evaluating potential alternates for external prefixes advertised by OSPF ASBRs, as well as explicit inequalities.

This document also provides clarifications and additional considerations to [RFC5286] to address a few coverage and operational observations. These observations are concerned with 1) the IS-IS ATT (attach) bit in the Level 1 (L1) area, 2) links provisioned with MAX_METRIC (see Section 5.1) for traffic engineering (TE) purposes, and 3) multi-topology (MT) IGP deployments. These are elaborated in detail in Sections 3.2 and 5.

This specification uses the same terminology introduced in [RFC5714] to represent LFA and builds on the notation for inequalities used in [RFC5286] to compute LFAs for MHPs.

1.1. Acronyms

AF	-	Address Family
ATT	-	IS-IS Attach Bit
ECMP	-	Equal-Cost Multipath
FRR	-	Fast Reroute
IGP	-	Interior Gateway Protocol
IS-IS	-	Intermediate System to Intermediate System
LFA	-	Loop-Free Alternate
LSP	-	Link State PDU (IS-IS)
MHP	-	Multi-Homed Prefix
MT	-	Multi-Topology
OSPF	-	Open Shortest Path First
SPF	-	Shortest Path First

1.2. 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. LFA Inequalities for MHPs

This document proposes the following set of LFA inequalities for selecting the most appropriate LFAs for MHPs. $Distance_opt(X,Y)$ (called " $D_opt(X,Y)$ " in this document) is defined in [RFC5714] and is nothing but the metric sum of the shortest path from X to Y. $Cost(X,Y)$, introduced in this document, is defined as the metric value of prefix Y from the prefix advertising node X. These LFAs can be derived from the inequalities in [RFC5286] combined with the observation that $D_opt(N,P) = \text{Min} (D_opt(N,PO_i) + Cost(PO_i,P))$ over all PO_i .

Link-Protecting LFAs:

A neighbor N can provide an LFA if and only if

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, S) + D_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Link-Protecting + Downstream-paths-only LFAs:

A subset of loop-free alternates are downstream paths that must meet a more restrictive condition that is applicable to more complex failure scenarios.

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Node-Protecting LFAs:

For an alternate next hop N to protect against node failure of a primary neighbor E for MHP P, N must be loop-free with respect to both E and MHP P. In other words, N's path to MHP P must not go through E (where N is the neighbor providing a loop-free alternate).

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, E) + D_{\text{opt}}(E, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Where:

- P - The MHP being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next hop on the shortest path from S to prefix P
- PO_i - The specific prefix-originating router being evaluated
- PO_{best} - The prefix-originating router on the shortest path from the computing router S to prefix P
- Cost(X,P) - The cost of reaching the prefix P from prefix originating node X
- D_{opt}(X,Y) - The distance on the shortest path from node X to node Y

3. LFA Selection for MHPs

To compute a valid LFA for a given MHP *P*, a computing router *S* MUST, for each alternate neighbor *N*, follow one of the appropriate procedures below once for each remote node that originated the prefix *P*.

Link-Protecting LFAs:

1. If, in addition to being an alternate neighbor, *N* is also a prefix originator of *P*,
 - A. Select *N* as an LFA for prefix *P* (irrespective of the metric advertised by *N* for the prefix *P*).
2. Else, evaluate the link-protecting LFA inequality for *P* with *N* as the alternate neighbor.
 - A. If the LFA inequality condition is met, select *N* as an LFA for prefix *P*.
 - B. Else, *N* is not an LFA for prefix *P*.

Link-Protecting + Downstream-paths-only LFAs:

1. Evaluate the link-protecting + downstream-paths-only LFA inequality for *P* with *N* as the alternate neighbor.
 - A. If the LFA inequality condition is met, select *N* as an LFA for prefix *P*.
 - B. Else, *N* is not an LFA for prefix *P*.

Node-Protecting LFAs:

1. If, in addition to being an alternate neighbor, *N* is also a prefix originator of *P*,
 - A. Select *N* as an LFA for prefix *P* (irrespective of the metric advertised by *N* for the prefix *P*).
2. Else, evaluate the appropriate node-protecting LFA inequality for *P* with *N* as the alternate neighbor.
 - A. If the LFA inequality condition is met, select *N* as an LFA for prefix *P*.
 - B. Else, *N* is not an LFA for prefix *P*.

If an alternate neighbor N is also one of the prefix originators of prefix P, it is guaranteed that N will not loop back packets destined for prefix P to computing router S. Therefore, N MUST be chosen as a valid LFA for prefix P without evaluating any of the inequalities in Section 2 as long as a downstream-paths-only LFA is not desired. To ensure such a neighbor N also provides a downstream-paths-only LFA, router S MUST also evaluate the downstream-paths-only LFA inequality specified in Section 2 for neighbor N and ensure router N satisfies the inequality.

However, if N is not a prefix originator of P, the computing router MUST evaluate one of the corresponding LFA inequalities defined in Section 2 once for each remote node that originated the prefix. If the inequality is satisfied by the neighbor N, router S MUST choose neighbor N as one of the valid LFAs for the prefix P.

For more specific rules, please refer to Section 4.

3.1. Improved Coverage with Simplified Approach to MHPs

Section 6.1 of the LFA base specification [RFC5286] recommends that a router computes the alternate next hop for an IGP MHP by considering alternate paths via all routers that have announced that prefix. The same has been elaborated with appropriate inequalities in the previous section. However, Section 6.1 of [RFC5286] also allows for the router to simplify the MHP calculation by assuming that the MHP is solely attached to the router that was its pre-failure optimal point of attachment, at the expense of potentially lower coverage. If an implementation chooses to simplify the MHP calculation by assuming that the MHP is solely attached to the router that was its pre-failure optimal point of attachment, the procedure described in this memo can potentially improve coverage for ECMP MHPs without incurring extra computational cost.

This document improves the above approach to provide loop-free alternatives without any additional cost for ECMP MHPs as described in the example network presented in Figure 1. The approach specified here may also be applicable for handling default routes as explained in Section 3.2.

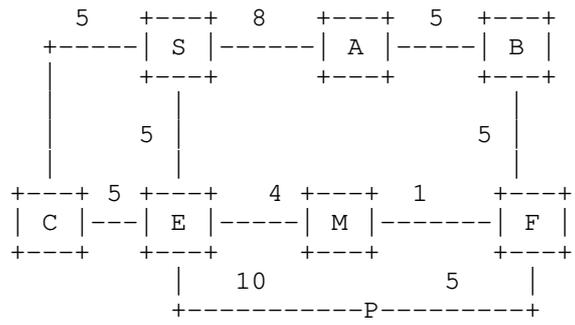


Figure 1: MHP with Same ECMP Next Hop

In Figure 1, a prefix P is advertised from both node E and node F. With a simplified approach taken as specified in Section 6.1 of [RFC5286], prefix P will get only a link-protecting LFA through the neighbor C while a node-protection path is available through neighbor A. In this scenario, E and F both are pre-failure optimal points of attachment and share the same primary next hop. Hence, an implementation MAY compare the kind of protection A provides to F (link and node protection) with the kind of protection C provides to E (link protection) and inherit the better alternative to prefix P. In this case, the better alternative is A.

However, in the example network presented in Figure 2, prefix P has an ECMP through both node E and node F with cost 20. Though it has two pre-failure optimal points of attachment, the primary next hop to each pre-failure optimal point of attachment is different. In this case, prefix P MUST inherit the corresponding LFAs of each primary next hop calculated for the router advertising the same. In Figure 2, that would be the LFA for node E and node F, i.e., node N1 and node N2, respectively.

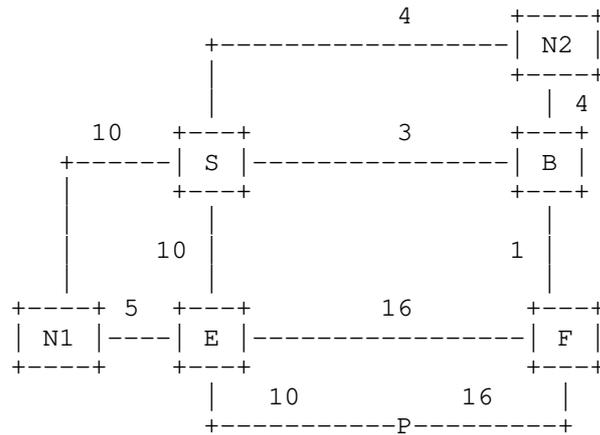


Figure 2: MHP with Different ECMP Next Hops

In summary, if there are multiple pre-failure points of attachment for an MHP, and the primary next hop of an MHP is the same as that of the primary next hop of the router that was the pre-failure optimal point of attachment, an implementation MAY provide a better protection to the MHP without incurring any additional computation cost.

3.2. IS-IS ATT Bit Considerations

Per [RFC1195], a default route needs to be added in the Level 1 (L1) router to the closest reachable Level 1 / Level 2 (L1/L2) router in the network advertising the ATT (attach) bit in its LSP-0 fragment. All L1 routers in the area would do this during the decision process with the next hop of the default route set to the adjacent router through which the closest L1/L2 router is reachable. The LFA base specification [RFC5286] does not specify any procedure for computing LFA for a default route in the IS-IS L1 area. This document specifies that a node can consider a default route is being advertised from the border L1/L2 router where the ATT bit is set and can do LFA computation for that default route. But, when multiple ECMP L1/L2 routers are reachable in an L1 area, corresponding best LFAs SHOULD be computed for each primary next hop associated with the default route as this would be similar to the ECMP MHP example described in Section 3.1. Considerations specified in Sections 3 and 3.1 are applicable for default routes if the default route is considered an ECMP MHP. Note that this document doesn't alter any ECMP handling rules or computation of LFAs for ECMP in general as laid out in [RFC5286].

4. LFA Selection for Multi-Homed External Prefixes

Redistribution of external routes into IGP is required 1) when two different networks get merged into one or 2) during protocol migrations.

During LFA calculation, alternate LFA next hops to reach the best ASBR could be used as LFA for the routes redistributed via that ASBR. When there is no LFA available to the best ASBR, it may be desirable to consider the other ASBRs (referred to as "alternate ASBRs" hereafter) redistributing the external routes for LFA selection as defined in [RFC5286] and leverage the advantage of having multiple redistributing nodes in the network.

4.1. IS-IS

LFA evaluation for multi-homed external prefixes in IS-IS is the same as the multi-homed internal prefixes. Inequalities described in Section 2 would also apply to multi-homed external prefixes.

4.2. OSPF

The LFA base specification [RFC5286] describes mechanisms to apply inequalities to find the loop-free alternate neighbor. Additional rules have to be applied in selecting the alternate ASBR for LFA consideration due to the external route calculation rules imposed by [RFC2328].

This document defines inequalities specifically for alternate loop-free ASBR evaluation. These inequalities are based on those in [RFC5286].

4.2.1. Rules to Select Alternate ASBRs

The process to select an alternate ASBR is best explained using the rules below. The process below is applied when a primary ASBR for the concerned prefix is chosen and there is an alternate ASBR originating the same prefix.

1. If RFC1583Compatibility is disabled:
 - A. If primary ASBR and alternate ASBR belong to intra-area non-backbone, go to step 2.
 - B. If primary ASBR and alternate ASBR belong to intra-area backbone and/or inter-area path, go to step 2.

- C. For other paths, skip this alternate ASBR and consider next ASBR.
2. Compare cost types (type 1 / type 2) advertised by alternate ASBR and primary ASBR:
 - A. If not the same type, skip alternate ASBR and consider next ASBR.
 - B. If the same, proceed to step 3.
3. If cost types are type 1, compare costs advertised by alternate ASBR and primary ASBR:
 - A. If costs are the same, then program ECMP FRR and return.
 - B. Else, go to step 5.
4. If cost types are type 2, compare costs advertised by alternate ASBR and primary ASBR:
 - A. If costs are different, skip alternate ASBR and consider next ASBR.
 - B. If costs are the same, proceed to step 4C to compare costs to reach ASBR/forwarding address.
 - C. If costs to reach ASBR/forwarding address are also the same, program ECMP FRR and return.
 - D. If costs to reach ASBR/forwarding address are different, go to step 5.
5. Compare route types (type 5 and type 7) for alternate ASBR and primary ASBR:
 - A. If route types are the same, check if route p-bit and forwarding address field for routes from both ASBRs match. If p-bit and forwarding address match, proceed to step 6. If not, skip this alternate ASBR and consider next ASBR.
 - B. If route types are not the same, skip this alternate ASBR and consider next alternate ASBR.
6. Apply inequality on alternate ASBR.

4.2.1.1. Multiple ASBRs Belonging to Different Areas

When RFC1583Compatibility is set to "disabled", OSPF [RFC2328] defines certain rules of preference to choose the ASBRs. While selecting an alternate ASBR for loop evaluation for LFA, these rules should be applied to ensure that the alternate neighbor does not cause looping.

When there are multiple ASBRs belonging to different areas advertising the same prefix, pruning rules as defined in Section 16.4 of [RFC2328] are applied. The alternate ASBRs pruned using these rules are not considered for LFA evaluation.

4.2.1.2. Type 1 and Type 2 Costs

If there are multiple ASBRs not pruned via the rules described in Section 4.2.1.1, the cost type advertised by the ASBRs is compared. ASBRs advertising type 1 costs are preferred, and the type 2 costs are pruned. If two ASBRs advertise the same type 2 cost, the alternate ASBRs are considered along with their cost to reach the ASBR/forwarding address for evaluation. If the two ASBRs have the same type 2 cost as well as the same cost to reach the ASBR, ECMP FRR is programmed. When there are multiple ASBRs advertising the same type 2 cost for the prefix, primary Autonomous System (AS) external route calculation, as described in Section 16.4.1 of [RFC2328], selects the route with the lowest type 2 cost. ASBRs advertising a different type 2 cost (higher cost) are not considered for LFA evaluation. Alternate ASBRs advertising a type 2 cost for the prefix but not chosen as primary due to a higher cost to reach ASBR are considered for LFA evaluation. The inequalities for evaluating alternate ASBR for type 1 and type 2 costs are same, as the alternate ASBRs with different type 2 costs are pruned and the evaluation is based on ASBRs with equal type 2 costs.

4.2.1.3. RFC1583Compatibility is Set to "Enabled"

When RFC1583Compatibility is set to "enabled", multiple ASBRs belonging to different areas advertising the same prefix are chosen based on cost and hence are valid alternate ASBRs for the LFA evaluation. The inequalities described in Section 4.2.2 are applicable based on forwarding address and cost type advertised in the external Link State Advertisement (LSA).

4.2.1.4. Type 7 Routes

Type 5 routes always get preference over type 7, and the alternate ASBRs chosen for LFA calculation should belong to the same type. Among type 7 routes, routes with the p-bit and forwarding address set have a higher preference than routes without these attributes. Alternate ASBRs selected for LFA comparison should have the same p-bit and forwarding address attributes.

4.2.2. Inequalities to Be Applied for Alternate ASBR Selection

The alternate ASBRs selected using the mechanism described in Section 4.2.1 are evaluated for loop-free criteria using the inequalities below.

4.2.2.1. Forwarding Address Set to Non-zero Value

Similar to the inequalities defined in Section 2, the following inequalities are defined when the forwarding address is a non-zero value.

Link-Protecting LFAs:

$$F_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, S) + F_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Link-Protecting + Downstream-paths-only LFAs:

$$F_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < F_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Node-Protecting LFAs:

$$F_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, E) + F_{\text{opt}}(E, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Where:

- P - The MHP being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next hop on the shortest path from S to prefix P
- PO_i - The specific prefix-originating router being evaluated

- PO_best - The prefix-originating router on the shortest path from the computing router S to prefix P
- Cost(X,Y) - The external cost for Y as advertised by X
- F_opt(X,Y) - The distance on the shortest path from node X to the forwarding address specified by ASBR Y
- D_opt(X,Y) - The distance on the shortest path from node X to node Y

4.2.2.2. ASBRs Advertising Type 1 and Type 2 Costs

Similar to the inequalities defined in Section 2, the following inequalities are defined for type 1 and type 2 costs.

Link-Protecting LFAs:

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, S) + D_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Link-Protecting + Downstream-paths-only LFAs:

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(S, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Node-Protecting LFAs:

$$D_{\text{opt}}(N, PO_i) + \text{Cost}(PO_i, P) < D_{\text{opt}}(N, E) + D_{\text{opt}}(E, PO_{\text{best}}) + \text{Cost}(PO_{\text{best}}, P)$$

Where:

- P - The MHP being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next hop on the shortest path from S to prefix P
- PO_i - The specific prefix-originating router being evaluated
- PO_best - The prefix-originating router on the shortest path from the computing router S to prefix P

N2 could be excluded as a potential alternative because of the current exclusions specified in Sections 3.5 and 3.6 of [RFC5286]. But, the primary traffic destined to D2 continues to use the link; hence, irrespective of the reverse metric in this case, the same link MAY be used as a potential LFA for D1.

Alternatively, the reverse metric of the link MAY be configured with `MAX_METRIC-1` so that the link can be used as an alternative while meeting the operator's TE requirements and without having to update the router to fix this particular issue.

5.2. MT Considerations

Sections 6.2 and 6.3.2 of [RFC5286] state that multi-topology OSPF and IS-IS are out of scope for that specification. This memo clarifies and describes the applicability.

In multi-topology IGP deployments, for each MT-ID, a separate shortest path tree (SPT) is built with topology-specific adjacencies so the LFA principles laid out in [RFC5286] are actually applicable for MT IS-IS [RFC5120] LFA SPF. The primary difference in this case is identifying the eligible set of neighbors for each LFA computation; this is done per MT-ID. The eligible set for each MT-ID is determined by the presence of IGP adjacency from the source to the neighboring node on that MT-ID apart from the administrative restrictions and other checks laid out in [RFC5286]. The same is also applicable for MT-OSPF [RFC4915] or different AFs in multi-instance OSPFv3 [RFC5838].

However, for MT IS-IS, if a "standard unicast topology" is used with MT-ID #0 [RFC5120] and both IPv4 [RFC5305] and IPv6 routes/AFs [RFC5308] are present, then the condition of network congruency is applicable for LFA computation as well. Network congruency here refers to having the same address families provisioned on all the links and all the nodes of the network with MT-ID #0. Here, with a single-decision process, both IPv4 and IPv6 next hops are computed for all the prefixes in the network. Similarly, with one LFA computation from all eligible neighbors per [RFC5286], all potential alternatives can be computed.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

The existing OSPF security considerations continue to apply, as do the recommended manual key management mechanisms specified in [RFC7474]. The existing security considerations for IS-IS also continue to apply, as specified in [RFC5304] and [RFC5310] and extended by [RFC7645] for Keying and Authentication for Routing Protocols (KARP). This document does not change any of the discussed protocol specifications (i.e., [RFC1195], [RFC5120], [RFC2328], and [RFC5838]); therefore, the security considerations of the LFA base specification [RFC5286] continue to apply.

8. References

8.1. 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>>.
- [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <<https://www.rfc-editor.org/info/rfc5286>>.
- [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>>.

8.2. Informative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [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>>.
- [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>>.
- [RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.
- [RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, April 2010, <<https://www.rfc-editor.org/info/rfc5838>>.
- [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>>.
- [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>>.
- [RFC7645] Chunduri, U., Tian, A., and W. Lu, "The Keying and Authentication for Routing Protocol (KARP) IS-IS Security Analysis", RFC 7645, DOI 10.17487/RFC7645, September 2015, <<https://www.rfc-editor.org/info/rfc7645>>.

Acknowledgements

The authors acknowledge Alia Atlas and Salih K.A. for their useful feedback and input. Thanks to Stewart Bryant for being Document Shepherd and providing detailed review comments. Thanks to Elwyn Davies for reviewing and providing feedback as part of the Gen-ART review. Thanks to Alvaro Retana, Adam Roach, Ben Campbell, Benjamin Kaduk, and sponsoring Routing Area Director Martin Vigoureux for providing detailed feedback and suggestions.

Contributors

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

Chris Bowers
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States of America

Email: cbowers@juniper.net

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Authors' Addresses

Pushpasis Sarkar (editor)
Arccus, Inc.

Email: pushpasis.ietf@gmail.com

Uma Chunduri (editor)
Huawei USA
2330 Central Expressway
Santa Clara, CA 95050
United States of America

Email: uma.chunduri@huawei.com

Shraddha Hegde
Juniper Networks, Inc.
Electra, Exora Business Park
Bangalore, KA 560103
India

Email: shraddha@juniper.net

Jeff Tantsura
Apstra, Inc.

Email: jefftant.ietf@gmail.com

Hannes Gredler
RtBrick, Inc.

Email: hannes@rtbrick.com

