

Routing Area Working Group
Internet-Draft
Updates: [5286](#) (if approved)
Intended status: Standards Track
Expires: May 25, 2019

P. Sarkar, Ed.
Arrcus, Inc.
U. Chunduri, Ed.
Huawei USA
S. Hegde
Juniper Networks, Inc.
J. Tantsura
Apstra, Inc.
H. Gredler
RtBrick, Inc.
November 21, 2018

Loop-Free Alternates selection for Multi-Homed Prefixes
draft-ietf-rtgwg-multihomed-prefix-lfa-09

Abstract

Deployment experience gained from implementing algorithms to determine Loop-Free Alternates (LFAs) for multi-homed prefixes has revealed some avenues for potential improvement. This document provides explicit inequalities that can be used to evaluate neighbors as a potential alternates for multi-homed prefixes. It also provides detailed criteria for evaluating potential alternates for external prefixes advertised by OSPF ASBRs. This documents updates and expands some of the "Routing Aspects" as specified in [Section 6 of RFC 5286](#).

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 RFC8174 \[RFC2119\] RFC8174 \[RFC8174\]](#) 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 May 25, 2019.

Copyright Notice

Copyright (c) 2018 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	3
2.	LFA inequalities for MHPs	4
3.	LFA selection for the multi-homed prefixes	5
3.1.	Improved coverage with simplified approach to MHPs	7
3.2.	IS-IS ATT Bit considerations	8
4.	LFA selection for the multi-homed external prefixes	9
4.1.	IS-IS	9
4.2.	OSPF	9
4.2.1.	Rules to select alternate ASBR	9
4.2.1.1.	Multiple ASBRs belonging different area	11
4.2.1.2.	Type 1 and Type 2 costs	11
4.2.1.3.	RFC1583compatibility is set to enabled	11
4.2.1.4.	Type 7 routes	11
4.2.2.	Inequalities to be applied for alternate ASBR selection	12
4.2.2.1.	Forwarding address set to non-zero value	12
4.2.2.2.	ASBRs advertising type1 and type2 cost	13
5.	LFA Extended Procedures	13
5.1.	Links with IGP MAX_METRIC	13
5.2.	Multi Topology Considerations	14
6.	IANA Considerations	15
7.	Acknowledgements	15
8.	Contributing Authors	15
9.	Security Considerations	16
10.	References	16

10.1.	Normative References	16
10.2.	Informative References	16
	Authors' Addresses	18

[1.](#) Introduction

A framework for the development of IP fast-reroute mechanisms is detailed in [\[RFC5714\]](#). The use of LFAs for IP Fast Reroute is specified in [\[RFC5286\]](#). If a prefix is advertised by more than one router that prefix is called as 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 a 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, additional considerations to [\[RFC5286\]](#), to address a few coverage and operational observations. These observations are in the area of handling IS-IS attach (ATT) bit in Level-1 (L1) area, links provisioned with MAX_METRIC (see [Section 5.1](#)) for traffic engineering (TE) purposes and in the area of Multi Topology (MT) IGP deployments. These are elaborated in detail in [Section 3.2](#) and [Section 5](#).

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

[1.1.](#) Acronyms

AF	-	Address Family
ATT	-	IS-IS Attach Bit
ECMP	-	Equal Cost Multi Path
IGP	-	Interior Gateway Protocol

IS-IS - Intermediate System to Intermediate System

LFA - Loop-Free Alternate

LSP - IS-IS Link State PDU

OSPF - Open Shortest Path First

MHP - Multi-homed Prefix

MT - Multi Topology

SPF - Shortest Path First

2. LFA inequalities for MHPs

This document proposes the following set of LFA inequalities for selecting the most appropriate LFAs for MHPs. $D_{\text{opt}}(X,Y)$ terminology is defined in [[RFC5714](#)], which is nothing but the metric sum of the shortest path from X to Y and $\text{Cost}(X,Y)$ 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 $D_{\text{opt}}(N,P) = \text{Min} (D_{\text{opt}}(N,PO_i) + \text{Cost}(PO_i,P))$ over all PO_i

Link-Protection:

A neighbor N can provide a loop-free alternate (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-Protection + Downstream-paths-only:

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-Protection:

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 multi-homed prefix being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next-hop on 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) - Cost of reaching the prefix P from prefix originating node X.
- D_{opt}(X,Y) - Distance on the shortest path from node X to node Y.

Figure 1: LFA inequalities for MHPs

3. LFA selection for the multi-homed prefixes

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

Link-Protection :

=====

1. if, in addition to being an alternate neighbor, N is also a prefix-originator of P,
 - 1.a. Select N as a 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 the N as the alternate neighbor.
 - 2.a. If LFA inequality condition is met, select N as a LFA for prefix P.
 - 2.b. Else, N is not a LFA for prefix P.

Link-Protection + Downstream-paths-only :

=====

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

Node-Protection :

=====

1. if, in addition to being an alternate neighbor, N is also a prefix-originator of P,
 - 1.a. Select N as a 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 the N as the alternate neighbor.
 - 2.a. If LFA inequality condition is met, select N as a LFA for prefix P.
 - 2.b. Else, N is not a LFA for prefix P.

Figure 2: Rules for selecting LFA for MHPs

In case an alternate neighbor N is also one of the prefix-originators of prefix P, N being a prefix-originator it is guaranteed that N will not loop back packets destined for prefix P to computing router S. So N MUST be chosen as a valid LFA for prefix P, without evaluating any of the inequalities in Figure 1 as long as 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-only LFA inequality specified in Figure 1 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, as mentioned in Figure 1, once for each remote node that originated the prefix. In case 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 the later sections of this document.

3.1. Improved coverage with simplified approach to MHPs

LFA base specification [\[RFC5286\] Section 6.1](#) 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 and the same has been elaborated with appropriate inequalities in the above section. However, [\[RFC5286\] Section 6.1](#) 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 equal cost multi path (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 through the below example network presented in Figure 3. The approach specified here may also be applicable for handling default routes as explained in [Section 3.2](#).




Figure 3: MHP with same ECMP Next-hop

In the above network a prefix P, is advertised from both Node E and Node F. With simplified approach taken as specified in [\[RFC5286\] Section 6.1](#), prefix P will get only link protection 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 and here it is A.

However, in the below example network presented in Figure 4, prefix P has an ECMP through both node E and node F with cost 20. Though it has 2 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 corresponding LFAs of each primary next-hop calculated for the router advertising the same respectively. In the below diagram that would be node E's and node F's LFA i.e., node N1 and node N2 respectively.

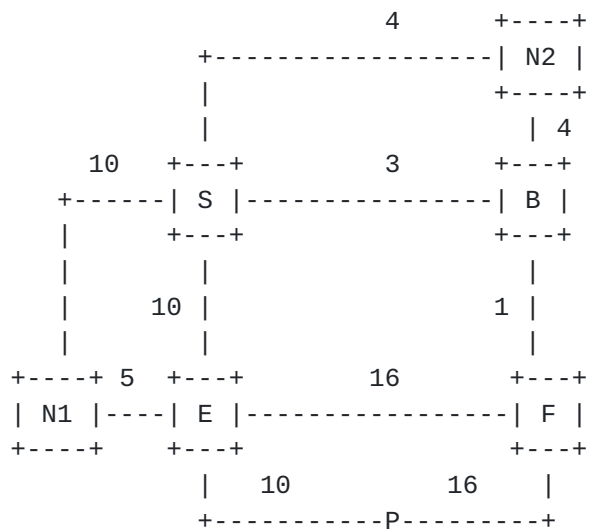


Figure 4: MHP with different ECMP Next-hops

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

3.2. IS-IS ATT Bit considerations

Per [RFC1195] a default route needs to be added in Level1 (L1) router to the closest reachable Level1/Level2 (L1/L2) router in the network advertising 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 base LFA specification [RFC5286] does not specify any procedure for computing LFA for a default route in IS-IS L1 area. This document specifies, a node can consider a default route is being advertised from the border L1/L2

router where 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 default route as this would be similar to ECMP MHP example as described in [Section 3.1](#). Considerations as specified in [Section 3](#) and [Section 3.1](#) are applicable for default routes, if the default route is considered as 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 the multi-homed external prefixes](#)

Redistribution of external routes into IGP is required in case of two different networks getting merged into one or during protocol migrations. External routes could be distributed into an IGP domain via multiple nodes to avoid a single point of failure.

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 ASBR 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 same to the multi-homed internal prefixes. Inequalities described in [Section 2](#) would also apply to multi-homed external prefixes.

[4.2. OSPF](#)

Loop Free Alternates [\[RFC5286\]](#) describes mechanisms to apply inequalities to find the loop free alternate neighbor. For the selection of alternate ASBR for LFA consideration, additional rules have to be applied in selecting the alternate ASBR due to the external route calculation rules imposed by [\[RFC2328\]](#).

This document defines inequalities specifically for the alternate loop-free ASBR evaluation, based on those in [\[RFC5286\]](#).

[4.2.1. Rules to select alternate ASBR](#)

The process to select an alternate ASBR is best explained using the rules below. The below process is applied when primary ASBR for the

concerned prefix is chosen and there is an alternate ASBR originating same prefix.

1. If RFC1583 Compatibility is disabled

- 1a. if primary ASBR and alternate ASBR belong to intra-area non-backbone go to step 2.
- 1b. If primary ASBR and alternate ASBR belong to intra-area backbone and/or inter-area path go to step 2.
- 1c. for other paths, skip this alternate ASBR and consider next ASBR.

2. Compare cost types (type 1/type 2) advertised by alternate ASBR and by the primary ASBR

- 2a. If not the same type skip alternate ASBR and consider next ASBR.
- 2b. If same proceed to step 3.

3. If cost types are type 1, compare costs advertised by alternate ASBR and by the primary ASBR

- 3a. If costs are the same then program ECMP Fast ReRoute (FRR) and return.
- 3b. else go to step 5..

4 If cost types are type 2, compare costs advertised by alternate ASBR and by the primary ASBR

- 4a. If costs are different, skip alternate ASBR and consider next ASBR.
- 4b. If cost are the same, proceed to step 4c to compare cost to reach ASBR/forwarding address.
- 4c. If cost to reach ASBR/forwarding address are also same program ECMP FRR and return.
- 4d. If cost to reach ASBR/forwarding address are different go to step 5.

5. If route type (type 5/type 7)

- 5a. If route type is same, check if the route p-bit and the forwarding address field for routes from both ASBRs match. If p-bit and forwarding address matches proceed to step 6.
If not, skip this alternate ASBR and consider next ASBR.
- 5b. If route type is not same, skip this alternate ASBR and consider next alternate ASBR.

6. Apply inequality on the alternate ASBR.

Figure 5: Rules for selecting alternate ASBR in OSPF

4.2.1.1. Multiple ASBRs belonging different area

When "RFC1583compatibility" is set to disabled, OSPF [[RFC2328](#)] defines certain rules of preference to choose the ASBRs. While selecting 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 area advertising the same prefix, pruning rules as defined in [\[RFC2328\] section 16.4](#) are applied. The alternate ASBRs pruned using above 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 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 same type 2 cost, the alternate ASBRs are considered along with their cost to reach ASBR/forwarding address for evaluation. If the two ASBRs have same type 2 cost as well as same cost to reach ASBR, ECMP FRR is programmed. When there are multiple ASBRs advertising same type 2 cost for the prefix, primary Autonomous System (AS) external route calculation as described in [\[RFC2328\] section 16.4.1](#) selects the route with lowest type 2 cost. ASBRs advertising different type 2 cost (higher cost) are not considered for LFA evaluation. Alternate ASBRs advertising type 2 cost for the prefix but are not chosen as primary due to 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 equal type 2 cost ASBRs.

4.2.1.3. RFC1583compatibility is set to enabled

When RFC1583Compatibility is set to enabled, multiple ASBRs belonging to different area advertising 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 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 same type. Among Type 7 routes, routes with p-bit and forwarding address set have higher preference than routes without these attributes. Alternate

ASBRs selected for LFA comparison should have same p-bit and forwarding address attributes.

4.2.2. Inequalities to be applied for alternate ASBR selection

The alternate ASBRs selected using above mechanism described in [Section 4.2.1](#), are evaluated for Loop free criteria using below inequalities.

4.2.2.1. Forwarding address set to non-zero value

Similar to inequalities as defined in Figure 1, the following inequalities are defined when forwarding address is a non-zero value.

Link-Protection:

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

Link-Protection + Downstream-paths-only:

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

Node-Protection:

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

Where,

- P - The multi-homed prefix being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next-hop on 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) - External cost for Y as advertised by X
- F_opt(X,Y) - Distance on the shortest path from node X to Forwarding address specified by ASBR Y.
- D_opt(X,Y) - Distance on the shortest path from node X to node Y.

Figure 6: LFA inequality definition when forwarding address is non-zero

4.2.2.2. ASBRs advertising type1 and type2 cost

Similar to inequalities as defined in Figure 1, the following inequalities are defined for type1 and type2 cost.

Link-Protection:

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

Link-Protection + Downstream-paths-only:

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

Node-Protection:

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

Where,

- P - The multi-homed prefix being evaluated for computing alternates
- S - The computing router
- N - The alternate router being evaluated
- E - The primary next-hop on 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) - External cost for Y as advertised by X.
- D_opt(X,Y) - Distance on the shortest path from node X to node Y.

Figure 7: LFA inequality definition for type1 and type2 cost

5. LFA Extended Procedures

This section explains the additional considerations in various aspects as listed below to the base LFA specification [[RFC5286](#)].

5.1. Links with IGP MAX_METRIC

[Section 3.5](#) and 3.6 of [[RFC5286](#)] describe procedures for excluding nodes and links from use in alternate paths based on the maximum link metric. If these procedures are strictly followed, there are situations, as described below, where the only potential alternate available which satisfies the basic loop-free condition will not be considered as alternative. This document refers the maximum link metric in IGP as the MAX_METRIC. MAX_METRIC is defined for IS-IS in

[RFC5305], where it is called as "maximum link metric" and defined for OSPF in [RFC6987], where it is called as "MaxLinkMetric".




Figure 8: Link with IGP MAX_METRIC

In the simple example network, all the link costs have a cost of 10 in both directions, except for the link between S and N2. The S-N2 link has a cost of 10 in the forward direction i.e., from S to N2, and a cost of MAX_METRIC ($0xffffffff / 2^{24} - 1$ for IS-IS and $0xffff$ for OSPF) in the reverse direction i.e., from N2 to S for a specific end-to-end Traffic Engineering (TE) requirement of the operator. At node S, D1 is reachable through N1 with cost 20, and D2 is reachable through N2 with cost 20. Even though neighbor N2 satisfies basic loop-free condition (inequality 1 of [RFC5286]) for D1, S's neighbor N2 could be excluded as a potential alternative because of the current exclusions as specified in [section 3.5](#) and 3.6 procedure of [RFC5286]. But, as the primary traffic destined to D2 continues to use the link and hence irrespective of the reverse metric in this case, same link MAY be used as a potential LFA for D1.

Alternatively, 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. Multi Topology Considerations

[Section 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 (MT) 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 which is done per MT ID. The eligible-set for each MT ID is determined by the presence of IGP adjacency from 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 topology" is used with MT-ID #0 [\[RFC5286\]](#) 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 same address families provisioned on all the links and all the nodes of the network with MT-ID #0. Here with single decision process both IPv4 and IPv6 next-hops are computed for all the prefixes in the network and similarly with one LFA computation from all eligible neighbors per [\[RFC5286\]](#), all potential alternatives can be computed.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

Authors acknowledge Alia Atlas and Salih K A for their useful feedback and inputs. 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 Gen-art review. Thanks to Alvaro Retena, Adam Roach, Ben Campbell, Benjamin Kaduk and sponsoring Routing Area Director Martin Vigoureux for providing detailed feedback and suggestions.

8. Contributing Authors

The following people contributed substantially to the content of this document and should be considered co-authors.

Chris Bowers
Juniper Networks, Inc.
1194 N. Mathilda Ave,
Sunnyvale, CA 94089, USA

Email: cbowers@juniper.net

Bruno Decraene
Orange,
France

Email: bruno.decraene@orange.com

9. 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 KARP. This document does not change any of the discussed protocol specifications [RFC1195] [RFC5120] [RFC2328] [RFC5838], and the security considerations of the LFA base specification [RFC5286] therefore continue to apply.

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>>.
- [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>>.

10.2. 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>>.

Authors' Addresses

Pushpasis Sarkar (editor)
Arrcus, Inc.

Email: pushpasis.ietf@gmail.com

Uma Chunduri (editor)
Huawei USA
2330 Central Expressway
Santa Clara, CA 95050
USA

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

