

Routing Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 10, 2014

P. Sarkar, Ed.
H. Gredler
S. Hegde
H. Raghuvver
Juniper Networks, Inc.
July 09, 2013

Node Protecting Remote-LFA and Manageability
draft-psarkar-rtgwg-rlfa-node-protection-00

Abstract

The loop-free alternates computed following the current Remote-LFA [I-D.ietf-rtgwg-remote-lfa] specification guarantees only link-protection. The resulting Remote-LFA nexthops (also called PQ-nodes), may not guarantee node-protection for all destinations being protected by it.

This document describes procedures for determining if a given PQ-node provides node-protection for a specific destination or not. The document also shows how the same procedure can be utilised for collection of complete characteristics for alternate paths. Knowledge about the characteristics of all alternate path is precursory to apply operator defined policy for eliminating paths not fitting constraints.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [RFC2119].

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 <http://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 January 10, 2014.

Copyright Notice

Copyright (c) 2013 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 (<http://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	2
2. Node Protection with Remote-LFA	3
2.1. The Problem	3
2.2. The Solution	4
3. Manageability of Remote-LFA Alternate Paths	6
3.1. The Problem	6
3.2. The Solution	7
4. Prior Solutions	7
4.1. Advantage of this Proposal	7
5. Acknowledgements	8
6. IANA Considerations	8
7. Security Considerations	8
8. References	8
8.1. Normative References	8
8.2. Informative References	8
Authors' Addresses	9

1. Introduction

The Remote-LFA [I-D.ietf-rtgwg-remote-lfa] specification provides loop-free alternates that gaurantees only link-protection. The resulting Remote-LFA alternate nexthops (also referred to as the PQ-nodes) may not provide node-protection for all destinations covered by the same, in case of failure of the primary nexthop node. Neither does the specification provide a means to determine the same.

Also, the LFA Manageability [I-D.ietf-rtgwg-lfa-manageability] document, requires a computing router to find all possible (including all possible Remote-LFA) alternate nexthops, collect the complete set

of path characteristics for each alternate path, run a alternate-selection policy (configured by the operator), and find the best alternate path. This will require the Remote-LFA implementation to gather all the required path characteristics along each link on the entire Remote-LFA alternate path.

With current LFA [RFC5286] and Remote-LFA implementations, the forward SPF (and reverse SPF) is run on the computing router and its immediate 1-hop routers as the roots. While that enables computation of path attributes (e.g. SRLG, Admin-groups) first alternate path segment from the computing router PQ-node, there is no means for the computing router to gather any path attributes for the path segment from the PQ-node to destination. Consecutively any policy-based selection of alternate paths will consider only the path attributes from the computing router up until the PQ-node.

This document describes a procedure for determining node-protection with Remote-LFA. The same procedure are also extended for collection of complete set of path attributes, enabling more accurate policy-based selection for alternate paths obtained with Remote-LFA.

2. Node Protection with Remote-LFA

2.1. The Problem

To better illustrate the problem and the solution proposed in this document the following topology diagram from the Remote-LFA [I-D.ietf-rtgwg-remote-lfa], draft is being re-used with slight modification.




Figure 1: Sample Ring Topology

In the above topology, for all (non-ECMP) destinations reachable via the S-E link there is no standard LFA alternate. As per the Remote-LFA [I-D.ietf-rtgwg-remote-lfa] alternate specifications node C being the PQ-node for the S-E link provides nexthop for all the above destinations. Table 1 below, shows all possible primary and Remote-LFA alternate paths for each destination.

Destination	Primary Path	Remote-LFA Backup Path
D	S->E->D	S=>A=>B=>C->D
E	S->E	S=>A=>B=>C->D->E
F	S->E->F	S=>A=>B=>C->D->E->F
G	S->E->D->G	S=>A=>B=>C->D->G

Table 1: Backup paths with Remote-LFA

A closer look at Table 1 shows that, while the PQ-node C provides link-protection for all the destinations, it does not provide node-protection for destinations E and F. In the event of the node-failure on primary nexthop E, the alternate path from Remote-LFA nexthop C to E and F also becomes unavailable. So for a Remote-LFA nexthop to provide node-protection for a given destination, it is mandatory that, the shortest path from the given PQ-node to the given destination must not traverse the primary nexthop.

2.2. The Solution

This document proposes an additional forward SPF computation for each of the PQ-nodes, as a mechanism to provide node-protection with remote LFA. In case, a alternate selection policy has been configured, the mechanism proposed, shall also provide a means to collect complete path attributes for the alternate path via a Remote-LFA nexthop to a given destination.

The additional forward SPF computation for each PQ-node, shall help determine, if a given primary nexthop node is on shortest paths from a given PQ-node to any given destination or not. To determine if a given PQ-node provides node-protecting alternate for a given destination, the primary nexthop node should not be on any of the shortest paths from the given PQ-node to the given destination.

Some SPF implementations may produce a list of links and nodes traversed on the shortest path(s) from a given root to others. In such implementations, running a forward SPF rooted at a given PQ-node will produce a list of nodes (one per each destination), on one or more shortest path(s) from the PQ-node to other destinations in the network. To determine whether a PQ-node provides node-protection for a given destination or not, the list of nodes computed from forward SPF run on the PQ-node, for the given destination, should be inspected. In case the list contains the primary nexthop node, the PQ-node does not provide node-protection. Else, the PQ-node guarantees node-protecting alternate for the given destination. Below is an illustration of the mechanism with the topology in Figure 1.

Destination	PQ-node	Shortest Path(PQ-node to Dest)	Link-Protection	Node-Protection
D	C	C->D	Yes	Yes
E	C	C->D->E	Yes	No
F	C	C->D->E->F	Yes	No
G	C	C->D->G	Yes	Yes

Table 2: Types of protection with Remote-LFA

As seen in the above example while C is node-protecting Remote-LFA nexthop for D and G, it is not so for E and F, since the primary nexthop E is in the shortest path from C to E and F.

Alternatively, an implementation may also run the node-protection condition from the LFA [RFC5286] specification with slight modification as shown in Figure 2 below. PQ-nodes that does not qualify the condition for a given destination, does not guarantee node-protection for the same.

$$D_{\text{opt}}(\text{Npq}, \text{Dst}) < D_{\text{opt}}(\text{Npq}, \text{Np}) + \text{Distance}_{\text{opt}}(\text{Np}, \text{Dst})$$

- $D_{\text{opt}}(X, Y)$: Distance on most optimum path from X to Y.
- Npq : The PQ-node being considered.
- Dst : The destination being protected.
- Np : The primary nexthop node on shortest path from computing router to destination.

Figure 2: Node-Protection Condition for Remote-LFA

All of the above metric costs except $D_{opt}(N_{pq}, Dst)$, can be obtained with forward and reverse SPF's with N_p (the primary nexthop) as the root, run as part of the regular LFA and Remote-LFA implementation. The $D_{opt}(N_{pq}, Dst)$ metric can only be determined by the additional forward SPF run with N_{pq} (PQ-node) as the root. With reference to the topology in Figure 1, Table 3 below shows how the above condition can be used to determine node-protection with a PQ-node.

Destination (Dst)	Primary-NH (N_p)	PQ-node (N_{pq})	$D_{opt}(N_{pq}, Dst)$	$D_{opt}(N_p, N_{pq})$	$D_{opt}(N_p, Dst)$	Condition Met
D	E	C	1	1	1	Yes
E	E	C	2	2	0	No
F	E	C	3	2	1	No
G	E	C	2	2	1	Yes

Table 3: Using Node-protecting condition for Remote-LFA

As seen in the above example above, C does not meet the node-protecting inequality for destination E, and F. And so, once again, while C is a node-protecting Remote-LFA nexthop for D and G, it is not so for E and F.

The procedure described in this document helps no more than to determine whether a given Remote-LFA alternate provides node-protection for a given destination or not. It does not find out any new Remote-LFA alternate nexthops, outside the ones already computed by standard Remote-LFA procedure. However, in case of availability of more than one PQ-node (Remote-LFA alternates) for a destination, and node-protection is required for the given primary nexthop, this procedure will eliminate the PQ-nodes that do not provide node-protection and choose only the ones that does.

3. Manageability of Remote-LFA Alternate Paths

3.1. The Problem

With the regular Remote-LFA functionality the computing router may compute more than one PQ-node as usable Remote-LFA alternate nexthops. Additionally a alternate selection policy may be configured to enable the network operator to choose one of them as the most appropriate Remote-LFA alternate. For such policy-based alternate selection to run, all the relevant path characteristics for

each the alternate paths (one through each of the PQ-nodes), needs to be collected. The Remote-LFA alternate path through a given PQ-node to a given destination comprises of two path segments as follows.

1. Path segment from the computing router to the PQ-node (Remote-LFA alternate nexthop), and
2. Path segment from the PQ-node to the destination being protected.

The first Path segment can be calculated from the regular forward SPF done as part of standard and remote LFA computations. However without the mechanism proposed in this document, there is no way to determine the path characteristics for the second path segment. In the absence of the path characteristics for the second path segment, two Remote-LFA alternate path may be equally preferred based on the first path segments characteristics only, although the second path segment attributes may be different.

3.2. The Solution

The additional forward SPF computation being proposed in this document shall also collect links, nodes and path characteristics along the second path segment. This shall enable collection of complete path characteristics for a given Remote-LFA alternate path to a given destination. The complete alternate path characteristics shall then facilitate more accurate alternate path selection while running the alternate selection policy.

4. Prior Solutions

A recent Node Protecting Remote-LFA [I-D.litkowski-rtgwg-node-protect-remote-lfa] draft proposes a solution for providing node-protection with Remote-LFA. It requires the computing router to additionally run reverse SPFs rooted at the nextnexthop routers (i.e. all the 2-hop neighborhood) as well. A simple study of standard IGP network topologies in real-life deployments shall reveal that, the increase in the number of required SPF computations is exponential, and can be a substantial computation overhead.

4.1. Advantage of this Proposal

Following are the advantages of the mechanism proposed in this document.

- o The recent Remote-LFA node-protection document [I-D.litkowski-rtgwg-node-protect-remote-lfa] proposes an extra reverse SPF computation for each nextnexthop of the computing

router. The mechanism in this document proposes an extra forward SPF for each of the PQ-nodes. Considering some of the standard IGP network topologies in real-life service-provider deployments, the number of nextnexthops will be substantially higher than the number of PQ-nodes discovered in those topologies. Hence the number of additional SPFs required in the proposed mechanism in this document will be considerably less compared to the procedures outlined in [I-D.litkowski-rtgwg-node-protect-remote-lfa], and imply less computational overhead.

- o Also the extra reverse SPF proposed per nextnexthop in Remote-LFA node-protection [I-D.litkowski-rtgwg-node-protect-remote-lfa] specification does not provide a means to collect the path characteristics for the alternate path segment from the PQ-node to the destination. The additional forward SPF for each PQ-node, as proposed in this document facilitates the same.

5. Acknowledgements

6. IANA Considerations

N/A. - No protocol changes are proposed in this document.

7. Security Considerations

This document does not introduce any change in any of the protocol specifications. It simply proposes to run an extra SPF rooted on each PQ-node discovered in the whole network.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[I-D.ietf-rtgwg-lfa-manageability]

Litkowski, S., Decraene, B., Filsfils, C., and K. Raza, "Operational management of Loop Free Alternates", draft-ietf-rtgwg-lfa-manageability-00 (work in progress), May 2013.

[I-D.ietf-rtgwg-remote-lfa]

Bryant, S., Filsfils, C., Previdi, S., Shand, M., and S. Ning, "Remote LFA FRR", draft-ietf-rtgwg-remote-lfa-02 (work in progress), May 2013.

[I-D.litkowski-rtgwg-node-protect-remote-lfa]

Litkowski, S., "Node protecting remote LFA", draft-litkowski-rtgwg-node-protect-remote-lfa-00 (work in progress), April 2013.

[RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, September 2008.

Authors' Addresses

Pushpasis Sarkar (editor)
Juniper Networks, Inc.
Electra, Exora Business Park
Bangalore, KA 560103
India

Email: psarkar@juniper.net

Hannes Gredler
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

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

Email: shraddha@juniper.net

Harish Raghuvver
Juniper Networks, Inc.
Electra, Exora Business Park
Bangalore, KA 560103
India

Email: hraghuveer@juniper.net