

Routing Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 4, 2013

S. Litkowski
Orange
April 2, 2013

Node protecting remote LFA
draft-litkowski-rtgwg-node-protect-remote-lfa-00

Abstract

This draft describes a simple extension to the remote LFA specification that computes a guaranteed node protection.

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](#)].

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 October 4, 2013.

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

Internet-Draft

Node protecting remote LFA

April 2013

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
2.	Specification	3
2.1.	Reference topology	4
2.2.	Current RLFA behavior	4
2.3.	Guaranteed node protection computation	4
2.3.1.	Computing node protecting extended P-Space	5
2.3.2.	Computing node protecting Q-Space	5
2.3.3.	Computing node protecting PQ-Spaces	5
3.	Scaling	6
4.	Security Considerations	7
5.	Acknowledgements	7
6.	IANA Considerations	7
7.	Normative References	7
	Author's Address	7

1. Introduction

The current RLFA specification described in [\[I-D.ietf-rtgwg-remote-lfa\]](#) is based on a per link computation (for optimization) that prevents determination of guaranteed node protection. This does not mean that the current solution is not able to provide node protection for a specific destination, but as computation is done from the link perspective, there is no guarantee to have node protection.




Figure 1

In the figure above, considering all metrics equal, primary path from S to D is SED. S has no LFA to reach D.

Using RLFA specification, C is a PQ from E perspective and could be used as a remote LFA in case of SE failure. In the diagram above, C is considered as a link protecting alternate (as from E perspective, it is only link protecting), even if from a topological point of view, C would be a node protecting RLFA.

There could be multiple reasons to ensure node protection in a network :

- o Presence of nodes with no high availability mechanism.
- o Avoiding transient loops in case of node crash.

This draft describes a solution to compute a guaranteed node protection using remote LFA solution.

[2.](#) Specification

Litkowski

Expires October 4, 2013

[Page 3]

Internet-Draft

Node protecting remote LFA

April 2013

[2.1.](#) Reference topology




Figure 2

[2.2.](#) Current RLFA behavior

The current remote LFA computation is based on computation of extended P-Space and Q-Space and then intersect both to determine PQ nodes that would be used for traffic tunneling. As Q-Space requires rSPF computation (one per Q-Space), it would not be scalable to compute Q-Space for each prefix. RLFA specification is proposing to compute Q Space on a per link basis, PQ will so be used by all prefixes using the link as primary nexthop.

Using figure 2 and considering protection of prefixes using E as primary nexthop on S, there would be three nodes in PQ space : P1,P2,P3. Using shortest path to PQ as tie breaker, P3 would be the best PQ.

P3 is only providing link protection to D1, D2 and D3 while P1 was providing node protection for D1 and D2, and P2 was providing node protection for D3.

In this scenario, RLFA is providing a suboptimal protection because of the approximation of using remote end of the link for computation.

[2.3.](#) Guaranteed node protection computation

To compute a guaranteed node protection using remote LFA, we just introduce a small modification in the current algorithm. Rather than computing link protecting PQs from nexthop perspective, we propose to compute node protecting PQs from nextnexthop perspective.

Our proposal is working as follows :

- o Compute node protecting PQ space for each nextnexthop.
- o Compute link protecting PQ space for each nexthop (as already done).
- o Prefixes are inheriting protection type from nexthop or nextnexthop based on defined policies. Choosing a PQ in node protecting space provides guaranteed node protection.

[2.3.1.](#) Computing node protecting extended P-Space

We define the notion of node protecting extended P-Space of a node S for a connected node E as the union of :

- o Node protecting P-Space : the set of routers reachable from S without traversing node E.
- o The set of destinations using E as primary nexthop but having a node protecting LFA.

[2.3.2.](#) Computing node protecting Q-Space

The set of routers from which a node N can be reached, by normal forwarding, without traversing the node E is termed the node protecting Q-space of N with respect to the node E. The Q-space can be obtained by computing a reverse shortest path tree (rSPT) rooted at N, with the sub-tree which traverses the failed node excised (including those which are members of an ECMP).

[2.3.3.](#) Computing node protecting PQ-Spaces

Upon termination of regular SPT, node S has knowledge of nexthops and nextnexthops to reach every destinations. For each nextnexthop on the SPT, S will compute :

- o Node protecting extended P-Space considering nexthop will fail.
- o Node protecting extended Q-Space from nextnexthop considering nexthop will fail.

Intersection of both will result in node protecting PQ-Spaces and node S will have a node protecting PQ-Space for each nextnexthop.

In figure 2, considering link SE, S has two nextnexthops (R1 and R2). S will compute node protecting PQ Space for R1 and for R2.

Node protecting extended P-Space :

- o For R1 : P1,N1
- o For R2 : P2,N2

Node protecting Q-Space :

- o For R1 : P1,D1,D2
- o For R2 : P2,D3,D2

Node protecting PQ-Space :

- o For R1 : P1

- o For R2 : P2

As a result, P1 is able to provide node protection for all destinations using R1 as nextnextthop (D1,D2). P2 is able to provide node protection for all destinations using R2 as nextnextthop (D3).

3. Scaling

Approximation introduced in remote LFA specification was done to optimize computation. Our proposal is introducing more computation but in a very acceptable degree. Computation time will be dependent on number of nextnextthops rather than nexthops.

Topology	Min	Average	Med	95th perc	Max
T1	1	20,4	19	39	107
T2	1	9,5	6	27	35
T3	2	15,7	11	41	53
T4	1	13,9	15	26,5	30
T5	1	12	8	36	72
T6	2	4,3	4	7	8
T7	1	9	6	19	22

Table 1: Number of nextnextthop

The simulation results above (from real SP topologies) show that the number of rSPF to compute is really acceptable : SPF takes today only few msec to compute even on large topologies. Moreover installing protection is not an urgent task, and it would be better to take more

time to compute a more optimal protection.

4. Security Considerations

This document does not introduce any change in security consideration compared to [[RFC5286](#)] or [[I-D.ietf-rtgwg-remote-lfa](#)].

[5.](#) Acknowledgements

[6.](#) IANA Considerations

This document has no action for IANA.

[7.](#) Normative References

[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-01](#) (work in progress), December 2012.

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

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

Author's Address

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com