

Routing Area Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: February 20, 2014

S. Litkowski  
B. Decraene  
Orange  
C. FilsFils  
K. Raza  
Cisco Systems  
August 19, 2013

Interactions between LFA and RSVP-TE  
draft-litkowski-rtgwg-lfa-rsvpte-cooperation-02

## Abstract

This document defines the behavior of a node supporting Loopfree Alternates (LFA) when the node has established RSVP TE tunnels. It first describes the decisions to be made by the LFA mechanism with respect to the use of TE tunnels as LFA candidates. Second, it discusses the use of RSVP TE tunnels as a way to complement the LFA coverage, illustrating how these technologies can benefit from each other.

## 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 February 20, 2014.

## Copyright Notice

Internet-Draft

lfa-rsvpte-cooperation

August 2013

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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	LFA FRR and MPLS-TE interactions . . . . .	<a href="#">3</a>
<a href="#">2.1.</a>	Use case : using MPLS LSP as LFA candidates . . . . .	<a href="#">3</a>
<a href="#">2.2.</a>	Specifications of interactions between LFA and TE LSP . .	<a href="#">4</a>
<a href="#">2.2.1.</a>	Having both a physical interface and a TE tunnel toward a LFA . . . . .	<a href="#">4</a>
<a href="#">2.2.2.</a>	TE ingress LSP as LFA candidate . . . . .	<a href="#">4</a>
<a href="#">2.2.3.</a>	Independence between LFA and TE FRR . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Operational considerations . . . . .	<a href="#">7</a>
<a href="#">3.1.</a>	Relevance of joint LFA FRR and RSVP-TE FRR deployments .	<a href="#">7</a>
<a href="#">3.2.</a>	Extending LFA coverage using RSVP-TE tunnels . . . . .	<a href="#">8</a>
<a href="#">3.2.1.</a>	Creating multihop tunnel to extend topology . . . . .	<a href="#">8</a>
<a href="#">3.2.2.</a>	Selecting multihop tunnels to extend topology . . . . .	<a href="#">9</a>
<a href="#">4.</a>	Security Considerations . . . . .	<a href="#">10</a>
<a href="#">5.</a>	Contributors . . . . .	<a href="#">10</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">10</a>
<a href="#">7.</a>	References . . . . .	<a href="#">10</a>
<a href="#">7.1.</a>	Normative References . . . . .	<a href="#">10</a>
<a href="#">7.2.</a>	Informative References . . . . .	<a href="#">10</a>
	Authors' Addresses . . . . .	<a href="#">11</a>

## [1.](#) Introduction

When a failure occurs in an IP network, the subsequent converge process often leads to traffic disruption. Some mechanisms are available to limit traffic disruptions by pre-computing alternate paths and locally reroute over these as soon as the failure is detected. Such techniques are commonly known as "protection

mechanisms". Currently, the protection mechanisms widely used in Service Provider networks are RSVP-TE Fast Reroute [[RFC4090](#)] and Loop Free Alternates [[RFC5286](#)]. RSVP-TE FRR permits full network coverage but with a quite high complexity in terms of operation, as well as potential scaling issues. On the other hand, LFA offer a very easy,

manageable, and scalable mechanism, but does not provide full coverage.

This document discusses how LFA and RSVP-TE should interact. It first describes how an LFA implementation should deal with existing RSVP TE tunnels established by the LFA node, as well as its behavior with respect to established IGP Shortcut tunnels [[RFC3906](#)]. Second, the document suggests the use of RSVP-TE tunnels to extend LFA coverage, and discusses the management and operational aspects of such a practice.

## [2.](#) LFA FRR and MPLS-TE interactions

This section discusses the various interactions among LFA FRR and MPLS-TE FRR. It starts with a simple example emphasizing the benefits of jointly using of LFA-FRR and MPLS-TE FRR, and then summarizes the requirements for the interactions between LFAs and MPLS-FRR.

### [2.1.](#) Use case : using MPLS LSP as LFA candidates

In some cases, typically in ring shaped parts of network topologies, links cannot be protected by LFAs. In the following topology, from the point of view of R5, LFAs are able to partially protect (49% of the destination routers) from the failure of R3, while the failure of R4 is not covered at all.




Figure 1

Many networks deploy MPLS tunnels for traffic engineering and resiliency reasons. To extend its benefit, an LFA implementation could take advantage of such existing MPLS tunnels. In the example above, if R5 has established TE tunnels bypassing R4 and R3, these could be considered as LFA candidates respectively protecting links from R5 to R4 and R3.

In the following section, we provide a detailed summary of the behavior to be applied by an LFA implementation which would consider the existence of MPLS TE tunnels to improve its applicability. The explicit configuration of such tunnels with the intent of improving LFA applicability is discussed in later sections.

## [2.2.](#) Specifications of interactions between LFA and TE LSP

Here we summarize the normative requirements for the interaction between LFA FRR and MPLS TE tunnels.

### [2.2.1.](#) Having both a physical interface and a TE tunnel toward a LFA

If a node S has both a physical interface and a TE tunnel to reach a LFA, it SHOULD use the physical interface unless :

1. The tunnel has been explicitly configured as an LFA candidate.
2. The tunnel does not pass through the link subject to LFA protection.

In other words, if a node S has an IGP/LDP forwarding entry F1 with outgoing interface i1, and S originates a TE tunnel T2 terminating on direct neighbor N2 (for example : if a TE tunnel is provisioned for link protection), T2 has an outgoing interface i2 and N2 is best LFA for F1, then an implementation MUST NOT use T2 when programming LFA repair for F1 unless T2 is configured as an LFA candidate.

### [2.2.2.](#) TE ingress LSP as LFA candidate

A TE LSP can be used as a virtual interface to reach a LFA if

1. The TE tunnel has been configured to allow its use as an LFA candidate.
2. The TE tunnel does not pass through the primary outgoing interface of D.

This would permit to extend LFA coverage as described in [[I-D.ietf-rtgwg-remote-lfa](#)], in a controlled fashion, as the tunnels used by the fast reroute mechanism are defined by configuration.

In other words, if a node S has an IGP/LDP forwarding entry F1 with outgoing interface i1 and S originates a TE tunnel T1 terminating at node Y, then an implementation SHOULD support a local policy which instructs node S to consider Y as a virtual neighbor and hence include Y as part of the LFA FRR alternate computation. In such

case, an implementation MUST not use Y as an LFA for F1 if T1's outgoing interface is i1.

### [2.2.3.](#) Independence between LFA and TE FRR

#### [2.2.3.1.](#) Tunnel head-end case

Similar requirements can be expressed for TE IGP shortcut tunnels.




Figure 2

PE to Cx metrics are 50, Cx to Cx are 1

A service provider is often providing traffic-engineered path for specific customer traffic (L3VPN, PW ...) to ensure path diversity or traffic constraints. In the diagram above, we consider a TE tunnel T2 built on a non shortest path as follows : PE1->C2->C3->PE2 and IGP shortcut is activated on PE1 to make traffic to PE2 using T2. Based on operational feedback, some implementations prevent LFA computation to run for an interface where a TE tunnel exists. In our example, if LFA is activated on N, we would not be able to have a protection for PE3 destination as a tunnel exists on the interface. This current observed behavior leads to a very limited coverage for LFA. In the other hand, it is important to keep protection mechanisms independant as much as possible to keep implementation simple. We propose the following approach :

- o If an IP prefix is reachable through a TE tunnel, LFA must not compute a protection for it.
- o If an IP prefix is reachable through a native IP path, LFA MUST compute a protection for it disregarding the presence of a tunnel or not on the primary interface.

In other words, if a node S has an IGP/LDP forwarding entry F1 with outgoing interface i1 and an IGP/LDP forwarding entry F2 with outgoing interface onto a TE tunnel T2 (due to IGP shortcut [[RFC3906](#)]) and tunnel T2 has outgoing interface i2, then an implementation MUST support enabling LFA FRR for F1 and using TE FRR for F2 as long as i1 != i2.

If i1 == i2, an implementation SHOULD allow for using LFA FRR backup for F1 and TE FRR backup for F2.

The mechanisms for using TE tunnel as an LFA candidate, and [RFC3906](#) mechanisms MUST be de-correlated -- i.e an implementation MUST support TE tunnel configuration with [RFC3906](#) only, as LFA candidate only, or both at the same time.

### 2.2.3.2. Tunnel midpoint case




Figure 3

PE to Cx metrics are 50, except PE3-C3 (60), Cx to Cx are 1

In the diagram above, we consider a TE tunnel T2 built on a non shortest path as follows : PE1->C2->C3->PE2 and IGP shortcut is activated on PE1 to make traffic to PE2 using T2. C2 is a TE tunnel midpoint router. In terms of forwarding, C2 has a MPLS TE forwarding entry for T2, as well as an IP forwarding entry to PE2. As explained in previous sections, it would be too restrictive and would limit LFA benefit on C2 if C2 would not be able to compute an LFA for the IP forwarding entry to PE2 due to the presence of a transit tunnel.

We propose the following approach for a midpoint router of a TE tunnel :

- o MPLS TE forwarding entries MUST not be protected by LFA (if an operator wants protection, TE FRR could be enabled).
- o IP forwarding entries MUST be protected by LFA disregarding the presence of a TE tunnel transiting through the primary interface of the destination.

In our example :

- o MPLS TE forwarding entry for T2 (ending on PE2) would be protected by TE-FRR (if enabled).
- o IP forwarding entry for PE2 would be protected by LFA.

In case of failure of C2-C3 :

- o traffic from PE1 to PE2 (encapsulated in T2), would be protected by TE FRR.
- o traffic from PE3 to PE2 (native IP), would be protected by LFA.

In other words, if a node S has an IGP/LDP forwarding entry F1 with outgoing interface i1 and a MPLS TE midpoint forwarding entry F2 with outgoing interface i2, then an implementation MUST support using LFA FRR for F1 and TE FRR for F2 as long as  $i1 \neq i2$ .

If  $i1 == i2$ , an implementation SHOULD allow for using LFA FRR backup for F1 and TE FRR backup for F2.

### [3.](#) Operational considerations

In this section, we first discuss the benefit of considering a joint deployment of LFA and MPLS tunnels to achieve resiliency. We then discuss one approach aiming at defining MPLS tunnels for the purpose of complementing LFA coverage.

#### [3.1.](#) Relevance of joint LFA FRR and RSVP-TE FRR deployments

This section describes the deployment scenarios where it can be beneficial to jointly use LFAs and RSVP-TE FRR.

There are many networks where RSVP-TE is already deployed. The deployment of RSVP-TE is typically for two main reasons :

- o Traffic engineering : a provider wants to route some flows on some specific paths using constraints;
- o Traffic protection using Fast-reroute ability

LFA is a feature that may bring benefits on RSVP-TE enabled networks, with no/minimal operational cost (compared to RSVP-TE FRR global roll out). These benefits include:

- o Should increase protection on network where FRR is not available



everywhere. Although it may not provide full coverage, it will increase the protection significantly.

- o May provide better protection in specific cases than RSVP-TE FRR

For IP networks that do not have any traffic protection mechanism, LFA is a very good first step to provide traffic protection even if its coverage is not 100%. Providers may want to increase protection coverage if LFA benefit is not sufficient for some destinations, in some parts of the network. The following sections discuss the use of basic RSVP-TE tunnels to extend protection coverage.

### [3.2.](#) Extending LFA coverage using RSVP-TE tunnels

We already have seen in previous sections that RSVP-TE tunnels could be established by an operator to complement LFA coverage. The method of tunnel placement depends on what type of protection (link or node) is required, as well as on the set of destinations or network parts which requires better protection than what LFA can provide.

#### [3.2.1.](#) Creating multihop tunnel to extend topology

To extend the coverage, the idea is to use a mechanism extending LFA by turning TE tunnels into LFA candidates. This mechanism is of a local significance only.

When explicitly establishing tunnels for that purpose, choices have to be made for the endpoints of such tunnels, in order to maximize coverage while preserving management simplicity. Requirements are that:

- o Endpoints must satisfy equations from [\[RFC5286\]](#), otherwise it will not be a valid LFA candidate: so when releasing traffic from tunnel, the traffic will go to the destination without flowing through the protected link or node. Depending on which equations are satisfied, node or link protection will be provided by the tunnel hop.
- o Tunnel must not flow through the link or node to be protected, explicit routing of tunnel is recommended to enforce this condition.

The approach to choose tunnel endpoints might be different here when compared to [\[I-D.ietf-rtgwg-remote-lfa\]](#) as endpoint choice is a manual one. Automatic behavior and scaling of [\[I-D.ietf-rtgwg-remote-lfa\]](#) requires:

- o Non null intersection of Extended P-Space and Q-Space
- o Computation of PQ node only for the remote end of the link

Based on this, [[I-D.ietf-rtgwg-remote-lfa](#)] may:

- o Not find a tunnel endpoint;
- o Not provide the more efficient protection : -- i.e. provides only link protection, while there is node protection possible for a specific destination

The proposed solution of manual explicitly routed tunnels is a good complement for [[I-D.ietf-rtgwg-remote-lfa](#)] and provides more flexibility:

- o Always a possibility to find a tunnel endpoint for a specific destination.
- o Possibility to provide a better protection type (link vs. node).

### [3.2.2.](#) Selecting multihop tunnels to extend topology

From a manageability point of view, computing a best Q node for each destination could lead to have one different Q node for each destination. This is not optimal in terms of number of tunnels, given that possibly one Q node may be able to serve multiple non covered destinations.

Rather than computing the best Q node per non covered destination, we would prefer to find best compromise Q nodes (best for multiple destinations). To find the best compromise between coverage increase and number of tunnels, we recommend to use a simulator performing the following computations per link:

- Step 1 : Compute for each not covered destination (routed on the link) the list of endpoints that are satisfying equations from [[RFC5286](#)] (node or link protection equations depending of required level of protection) : nodes in Q-Space
- Step 2 : Remove endpoints that are not eligible for repair (Edge nodes, low bandwidth meshed nodes, number of hops ...) : multiple attributes could be specified to exclude some nodes from Q-Space : The example of attributes include router type, metric to node, bandwidth, packet loss, RTD ...

Step 3 : Within the list of endpoints (one list per destination),  
order the endpoints by number of destination covered

Step 4 : Choose the endpoint that has the highest number of  
destination covered : some other criteria could be used to prefer  
an endpoint from another (same type of criteria that excluded some  
nodes from Q-Space)

Step 5 : Remove destinations covered by this endpoint from non  
covered list

Step 6 : If non covered list is not empty, restart from Step 1

Multiple endpoint (and so tunnels) could be necessary to have 100%  
coverage. But the idea is to find a tradeoff between number of  
tunnels configured (complexity) and number of destination covered,  
combining with traffic information would also provide a better view.

#### [4.](#) Security Considerations

TBD.

#### [5.](#) Contributors

Significant contribution was made by Pierre Francois which the  
authors would like to acknowledge.

#### [6.](#) IANA Considerations

This document has no actions for IANA.

#### [7.](#) References

##### [7.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate  
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3906] Shen, N. and H. Smit, "Calculating Interior Gateway  
Protocol (IGP) Routes Over Traffic Engineering Tunnels",  
[RFC 3906](#), October 2004.

- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.
- [RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), September 2008.

## [7.2.](#) Informative References

Litkowski, et al. Expires February 20, 2014 [Page 10]

---

Internet-Draft lfa-rsvpte-cooperation August 2013

[I-D.bryant-ipfrr-tunnels]  
Bryant, S., Filsfils, C., Previdi, S., and M. Shand, "IP Fast Reroute using tunnels", [draft-bryant-ipfrr-tunnels-03](#) (work in progress), November 2007.

[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.

### Authors' Addresses

Stephane Litkowski  
Orange

Email: [stephane.litkowski@orange.com](mailto:stephane.litkowski@orange.com)

Bruno Decraene  
Orange

Email: [bruno.decraene@orange.com](mailto:bruno.decraene@orange.com)

Clarence FilsFils  
Cisco Systems

Email: [cfilsfil@cisco.com](mailto:cfilsfil@cisco.com)

Kamran Raza  
Cisco Systems

Email: [skraza@cisco.com](mailto:skraza@cisco.com)

Litkowski, et al.

Expires February 20, 2014

[Page 11]