

TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 6, 2017

T. Saad, Ed.
R. Gandhi, Ed.
Z. Ali
Cisco Systems, Inc.
R. Venator
Defense Information Systems Agency
Y. Kamite
NTT Communications Corporation
February 2, 2017

RSVP Extensions For Re-optimization of Loosely Routed
Point-to-Multipoint Traffic Engineering Label Switched Paths (LSPs)
draft-ietf-teas-p2mp-loose-path-reopt-09

Abstract

Re-optimization of a Point-to-Multipoint (P2MP) Traffic Engineered (TE) Label Switched Path (LSP) may be triggered based on the need to re-optimize an individual source-to-leaf (S2L) sub-LSP or a set of S2L sub-LSPs, both using Sub-Group-Based Re-optimization method, or the entire P2MP-TE LSP tree using the Make-Before-Break (MBB) method.

This document discusses the application of the existing mechanisms for path re-optimization of loosely routed Point-to-Point (P2P) TE LSPs to the P2MP-TE LSPs, identifies issues in doing so and defines procedures to address them. When re-optimizing a large number of S2L sub-LSPs in a tree using the Sub-Group-Based Re-optimization method, the S2L sub-LSP descriptor list may need to be semantically fragmented. This document defines the notion of a fragment identifier to help recipient nodes unambiguously reconstruct the fragmented S2L sub-LSP descriptor list.

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."

Copyright Notice

Copyright (c) 2017 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	4
2. Conventions Used in This Document	4
2.1. Key Word Definitions	4
2.2. Abbreviations	5
2.3. Terminology	5
3. Overview	5
3.1. Loosely Routed Inter-domain P2MP-TE LSP Tree	6
3.2. Existing Mechanism For Tree-Based P2MP-TE LSP Re-optimization	6
3.3. Existing Mechanism For Sub-Group-Based P2MP-TE LSP Re-optimization	7
4. Signaling Extensions For Loosely Routed P2MP-TE LSP Re-optimization	8
4.1. Tree-Based Re-optimization	8
4.2. Sub-Group-Based Re-optimization Using Fragment Identifier	9
5. Message and Object Definitions	11
5.1. P2MP-TE Tree Re-evaluation Request Flag	11
5.2. Preferable P2MP-TE Tree Exists Path Error Sub-code	11
5.3. Fragment Identifier For S2L sub-LSP Descriptor	11
6. Compatibility	12
7. IANA Considerations	13
7.1. P2MP-TE Tree Re-evaluation Request Flag	13
7.2. Preferable P2MP-TE Tree Exists Path Error Sub-code	13
7.3. Fragment Identifier For S2L sub-LSP Descriptor	14
8. Security Considerations	14
9. References	16
9.1. Normative References	16
9.2. Informative References	16
Acknowledgments	17
Author's Addresses	17

1. Introduction

This document defines Resource Reservation Protocol - Traffic Engineering (RSVP-TE) [RFC2205] [RFC3209] signaling extensions for re-optimizing loosely routed Point-to-Multipoint (P2MP) Traffic Engineered (TE) Label Switched Paths (LSPs) [RFC4875] in a Multi-Protocol Label Switching (MPLS) or Generalized MPLS (GMPLS) [RFC3473] network.

A P2MP-TE LSP is comprised of one or more source-to-leaf (S2L) sub-LSPs. A loosely routed P2MP-TE S2L sub-LSP is defined as one whose path does not contain the full explicit route identifying each node along the path to the egress node at the time of its signaling by the ingress node. Such an S2L sub-LSP is signaled with no Explicit Route Object (ERO) [RFC3209], or with an ERO that contains at least one loose next-hop, or with an ERO that contains an abstract node which identifies more than one node. This is often the case with inter-domain P2MP-TE LSPs where Path Computation Element (PCE) is not used [RFC5440].

As per [RFC4875], an ingress node may re-optimize the entire P2MP-TE LSP tree by re-signaling all its S2L sub-LSP(s) using the Make-Before-Break (MBB) method or may re-optimize individual S2L sub-LSP or a set of S2L sub-LSPs i.e. individual destination or a set of destinations, both using the Sub-Group-Based Re-optimization method.

[RFC4736] defines RSVP signaling procedure for re-optimizing the path(s) of loosely routed Point-to-Point (P2P) TE LSP(s). Those mechanisms include a method for the ingress node to trigger a new path re-evaluation request and a method for the mid-point node to notify availability of a preferred path. This document discusses the application of those mechanisms to the re-optimization of loosely routed P2MP-TE LSPs, identifies issues in doing so and defines procedures to address them.

For re-optimizing a group of S2L sub-LSPs in a tree using the Sub-Group-Based Re-optimization method, an S2L sub-LSP descriptor list can be used to signal one or more S2L sub-LSPs in an RSVP message. This RSVP message may need to be semantically fragmented when large number of S2L sub-LSPs are added to the descriptor list. This document defines the notion of a fragment identifier to help recipient nodes unambiguously reconstruct the fragmented S2L sub-LSP descriptor list.

2. Conventions Used in This Document

2.1. Key Word Definitions

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].

2.2. Abbreviations

ABR: Area Border Router.

AS: Autonomous System.

ERO: Explicit Route Object.

LSR: Label Switching Router.

S2L sub-LSP: Source-to-leaf sub Label Switched Path.

TE LSP: Traffic Engineering Label Switched Path.

TE LSP ingress: Head-end/source node of the TE LSP.

TE LSP egress: Tail-end/destination node of the TE LSP.

2.3. Terminology

The reader is assumed to be familiar with the terminology in [RFC4736] and [RFC4875].

3. Overview

[RFC4736] defines RSVP signaling extensions for re-optimizing loosely routed P2P TE LSPs as follows:

- o A mid-point LSR that expands loose next-hop(s) sends a solicited or unsolicited PathErr with the Notify error code 25 (as defined in [RFC3209]) with sub-code 6 to indicate "Preferable Path Exists" to the ingress node.
- o An ingress node triggers a path re-evaluation request at all mid-point LSR(s) that expands loose next-hop(s) by setting the "Path Re-evaluation Request" flag (0x20) in SESSION_ATTRIBUTES Object in the Path message.
- o The ingress node upon receiving this PathErr with the Notify error code either solicited or unsolicited initiates re-optimization of the LSP using the MBB method with a different LSP-ID.

The following sections discuss the issues that may arise when

applying the mechanisms defined in [RFC4736] for re-optimizing loosely routed P2MP-TE LSPs.

3.1. Loosely Routed Inter-domain P2MP-TE LSP Tree

An example of a loosely routed inter-domain P2MP-TE LSP tree is shown in Figure 1. In this example, the P2MP-TE LSP tree consists of 3 S2L sub-LSPs, to destinations (i.e. leafs) R10, R11 and R12 from the ingress node (i.e. source) R1. Nodes R2 and R5 are branch nodes and nodes ABR3, ABR4, ABR7, ABR8 and ABR9 are area border routers. For the S2L sub-LSP to destination R10, nodes ABR3, ABR7 and R10 are defined as loose next-hops. For the S2L sub-LSP to destination R11, nodes ABR3, ABR8 and R11 are defined as loose next-hops. For the S2L sub-LSP to destination R12, nodes ABR4, ABR9 and R12 are defined as loose next-hops.




Figure 1: An Example of Loosely Routed Inter-domain P2MP-TE LSP Tree

3.2. Existing Mechanism For Tree-Based P2MP-TE LSP Re-optimization

Mechanisms defined in [RFC4736] can be easily applied to trigger the re-optimization of individual or group of S2L sub-LSP(s). However, to apply these [RFC4736] mechanisms for triggering the re-optimization of a P2MP-TE LSP tree, an ingress node needs to send path re-evaluation requests on all (typically 100s of) S2L sub-LSPs and the mid-point LSR needs to send PathErrs with the Notify error code for all S2L sub-LSPs. Such mechanisms may lead to the following issues:

- o A mid-point LSR that expands loose next-hop(s) may have to

accumulate the received path re-evaluation request(s) for all S2L sub-LSPs (e.g. by using a wait timer) and interpret them as a re-optimization request for the whole P2MP-TE LSP tree. Otherwise, a mid-point LSR may prematurely notify "Preferable Path Exists" for one or a sub-set of S2L sub-LSPs.

- o Similarly, the ingress node may have to heuristically determine when to perform P2MP-TE LSP tree re-optimization and when to perform S2L sub-LSP re-optimization. For example, an implementation may choose to delay re-optimization long enough to allow all PathErr(s) to be received. Such timer-based procedures may produce undesired results.
- o The ingress node that receives (un)solicited PathErr(s) with the Notify error code for individual S2L sub-LSP(s), may prematurely start re-optimizing the sub-set of S2L sub-LSPs. However, as mentioned in [RFC4875] Section 14.2, such sub-group based re-optimization procedure may result in data duplication that can be avoided if the entire P2MP-TE LSP tree is re-optimized using the Make-Before-Break method with a different LSP-ID, especially if the ingress node eventually receives PathErrs with the Notify error code for all S2L sub-LSPs of the P2MP-TE LSP tree.

In order to address above mentioned issues and to align re-optimization of P2MP-TE LSP with P2P LSP [RFC4736], there is a need for a mechanism to trigger re-optimization of the LSP tree by re-signaling all S2L sub-LSPs with a different LSP-ID. To meet this requirement, this document defines RSVP-TE signaling extensions for the ingress node to trigger the re-evaluation of the P2MP LSP tree on every hop that has a next-hop defined as a loose or abstract hop for one or more S2L sub-LSP path, and a mid-point LSR to signal to the ingress node that a preferable LSP tree exists (compared to the current path) or that the whole P2MP-TE LSP must be re-optimized (because of maintenance required on the TE LSP path) (see Section 4.1).

3.3. Existing Mechanism For Sub-Group-Based P2MP-TE LSP Re-optimization

Applying the procedures discussed in RFC4736 in conjunction with the Sub-Group-Based Re-Optimization procedures ([RFC4875], Section 14.2), an ingress node MAY trigger path re-evaluation requests for a set of S2L sub-LSPs in a single Path message using S2L sub-LSP descriptor list. Similarly, a mid-point LSR may send a PathErr with the Notify error code 25 and sub-code 6 containing a list of S2L sub-LSPs transiting through the LSR using an S2L sub-LSP descriptor list to notify the ingress node. This method can be used for re-optimizing a sub-group of S2L sub-LSPs within an LSP tree using the same LSP-ID.

This method can alleviate the scale issue associated with sending RSVP messages for individual S2L sub-LSPs. However, this procedure can lead to the following issues when used to re-optimize the LSP tree:

- o Path message that is intended to carry the path re-evaluation request as defined in [RFC4736] with a full list of S2L sub-LSPs in S2L sub-LSPs descriptor list will be decomposed at branching LSRs, and only a subset of the S2L sub-LSPs that are routed over the same next-hop will be added in the descriptor list of the Path message propagated to downstream mid-point LSRs. Consequently, when a preferable path exists at such mid-point LSRs, the PathErr with the Notify error code can only include the sub-set of S2L sub-LSPs traversing the LSR. In this case, at the ingress node there is no way to distinguish which mode of re-optimization to invoke, i.e. sub-group based re-optimization using the same LSP-ID or tree based re-optimization using a different LSP-ID.
- o An LSR may semantically fragment a large RSVP message (when a combined message may not be large enough to fit all S2L sub-LSPs). In this case, the ingress node may receive multiple PathErrs with sub-sets of S2L sub-LSPs in each (due to either the combined Path message getting fragmented or the combined PathErr message getting fragmented) and would require additional logic to determine how to re-optimize the LSP tree (for example, waiting for some time to aggregate all possible PathErr messages before taking an action). When fragmented, RSVP messages may arrive out of order, and the receiver has no way of knowing the beginning and end of the S2L sub-LSP list.

In order to address the above mentioned issues caused by RSVP message semantic fragmentation, this document defines new fragment identifier object for the S2L sub-LSP descriptor list when combining large number of S2L sub-LSPs in an RSVP message (see Section 4.2).

4. Signaling Extensions For Loosely Routed P2MP-TE LSP Re-optimization

4.1. Tree-Based Re-optimization

To evaluate a P2MP-TE LSP tree on mid-point LSRs that expand loose next-hop(s), an ingress node MAY send a Path message with "P2MP-TE Tree Re-evaluation Request (value TBA1)" defined in this document. The ingress node selects one of the S2L sub-LSPs of the P2MP-TE LSP tree transiting a mid-point LSR to trigger the re-evaluation request. The ingress node MAY send a re-evaluation request to each border LSR on the path of the LSP tree.

A mid-point LSR that expands loose next-hop(s) for one or more S2L sub-LSP path(s) does the following upon receiving a Path message with the "P2MP-TE Tree Re-evaluation Request" flag set:

- o The mid-point LSR MUST check for a preferable P2MP-TE LSP tree by re-evaluating all S2L sub-LSP(s) that are expanded paths of the loose next-hops of the P2MP-TE LSP.
- o If a preferable P2MP-TE LSP tree is found, the mid-point LSR MUST send an RSVP PathErr with the Notify error code 25 defined in [RFC3209] and sub-code "Preferable P2MP-TE Tree Exists (value TBA2)" defined in this document to the ingress node. The mid-point LSR, in turn, SHOULD NOT propagate the "P2MP-TE Tree Re-evaluation Request" flag in the subsequent RSVP Path messages sent downstream for the re-evaluated P2MP-TE LSP.
- o If no preferable tree for P2MP-TE LSP can be found, the mid-point LSR that expands loose next-hop(s) for one or more S2L sub-LSP path(s) MUST propagate the request downstream by setting the "P2MP-TE Tree Re-evaluation Request" flag in the LSP_ATTRIBUTES Object of the RSVP Path message.

A mid-point LSR MAY send an unsolicited PathErr with the Notify error code and sub-code "Preferable P2MP-TE Tree Exists" to the ingress node to notify of a preferred P2MP-TE LSP tree when it determines it exists. In this case, the mid-point LSR that expands loose next-hop(s) for one or more S2L sub-LSP path(s) selects one of the S2L sub-LSP(s) of the P2MP-TE LSP tree to send this PathErr message to the ingress node. The mid-point LSR SHOULD consider how frequently it chooses to send such a PathErr - considering both that a PathErr may be lost on its transit to the ingress node and that the ingress node may choose not to re-optimize the LSP when such a PathErr is received.

The sending of an RSVP PathErr with the Notify error code and "Preferable P2MP-TE Tree Exists" sub-code to the ingress node notifies the ingress node of the existence of a preferable P2MP-TE LSP tree and upon receiving this PathErr, the ingress node SHOULD trigger re-optimization of the LSP using the MBB method with a different LSP-ID.

4.2. Sub-Group-Based Re-optimization Using Fragment Identifier

It might be preferable, as per [RFC4875], to re-optimize the entire P2MP-TE LSP by re-signaling all of its S2L sub-LSP(s) (Section 14.1, "Make-before-Break") or to re-optimize individual or group of S2L sub-LSP(s) i.e. individual or group of destination(s) (Section 14.2

"Sub-Group-Based Re-Optimization" in [RFC4875]), both using the same LSP-ID. For loosely routed S2L sub-LSPs, this can be achieved by using the procedures defined in [RFC4736] to re-optimize one or more S2L sub-LSP(s) of the P2MP-TE LSP.

An ingress node may trigger path re-evaluation requests using the procedures defined in [RFC4736] for a set of S2L sub-LSPs by combining multiple Path messages using an S2L sub-LSP descriptor list [RFC4875]. An S2L sub-LSP descriptor list is created using a series of S2L_SUB_LSP Objects as defined in [RFC4875]. Similarly, a mid-point LSR may send a PathErr with the Notify error code (value 25) and "Preferable Path Exists" (sub-code 6) containing a list of S2L sub-LSPs transiting through the LSR using an S2L sub-LSP descriptor list to notify the ingress node of preferable paths available.

As per [RFC4875] (Section 5.2.3, "Transit Fragmentation of Path State Information"), when a Path message is not large enough to fit all S2L sub-LSPs in the descriptor list, an LSR may semantically fragment the message. In this case, the LSR MUST add the S2L_SUB_LSP_FRAG Object defined in this document in the S2L sub-LSP descriptor to be able to rebuild the list from the received fragments that may arrive out of order.

The S2L_SUB_LSP_FRAG Object defined in this document is optional. However, a node MUST add the S2L_SUB_LSP_FRAG Object for each fragment in S2L sub-LSP descriptor when the RSVP message needs to be fragmented.

A mid-point LSR SHOULD wait to accumulate all S2L sub-LSPs before attempting to re-evaluate preferable path when a Path message for "Path Re-evaluation Request" is received with S2L_SUB_LSP_FRAG Object. If a mid-point LSR does not receive all fragments of the Path message (for example, when fragments are lost) within a configurable time interval, it SHOULD trigger re-evaluation of all S2L sub-LSPs of the P2MP-TE LSP transiting on the node. A mid-point LSR MUST receive at least one fragment of the Path message to trigger this behaviour.

An ingress node SHOULD wait to accumulate all S2L sub-LSPs before attempting to trigger re-optimization when a PathErr with Notify error code and "Preferable Path Exists" sub-code is received with a S2L_SUB_LSP_FRAG Object. If an ingress node does not receive all fragments of the PathErr message (for example, when fragments are lost) within a configurable time interval, it SHOULD trigger re-optimization of all S2L sub-LSPs of the P2MP-TE LSP transiting on the mid-point node that had sent the PathErr message. An ingress node MUST receive at least one fragment of the PathErr message to trigger this behaviour.

The S2L_SUB_LSP_FRAG Object defined in this document has a wider applicability in addition to the P2MP-TE LSP re-optimization. It can also be used (in Path and Resv messages) to setup a new P2MP-TE LSP, send other PathErr messages as well as Path Tear and Resv Tear messages for a set of S2L sub-LSPs. This is outside the scope of this document.

5. Message and Object Definitions

5.1. P2MP-TE Tree Re-evaluation Request Flag

In order to trigger a tree re-evaluation request, a new flag is defined in Attributes Flags TLV of the LSP_ATTRIBUTES Object [RFC5420] as follows:

Bit Number (TBA1, to be assigned by IANA): P2MP-TE Tree
Re-evaluation Request flag

The "P2MP-TE Tree Re-evaluation Request" flag is meaningful in a Path message of a P2MP-TE S2L sub-LSP and is inserted by the ingress node using the message format defined in [RFC6510].

5.2. Preferable P2MP-TE Tree Exists Path Error Sub-code

In order to indicate to an ingress node that a preferable P2MP-TE LSP tree exists, the following new sub-code for PathErr with Notify error code 25 [RFC3209] is defined:

Sub-code (TBA2, to be assigned by IANA): Preferable P2MP-TE Tree
Exists sub-code

When a preferable path for P2MP-TE LSP tree exists, the mid-point LSR sends a solicited or unsolicited "Preferable P2MP-TE Tree Exists" sub-code with PathErr with Notify error code 25 to the ingress node of the P2MP-TE LSP.

5.3. Fragment Identifier For S2L sub-LSP Descriptor

The S2L_SUB_LSP Object [RFC4875] identifies a particular S2L sub-LSP belonging to the P2MP-TE LSP. An S2L sub-LSP descriptor list is created using a series of S2L_SUB_LSP Objects as defined in [RFC4875]. The RSVP message may need to be semantically fragmented [RFC4875] due to large number of S2L sub-LSPs added in the descriptor list, and such fragments may be received out of order. To be able to rebuild the fragmented S2L sub-LSP descriptor list correctly, the

following Object is defined to identify the fragments.

S2L_SUB_LSP_FRAG: Class-Num TBA3 by IANA

Length (8 bytes)	Class-Num TBA3	C-Type 1
Fragment ID	Fragments Tot	Fragment Num

Fragment ID: 16-bit integer in the range of 1 to 65535.

This value is incremented for each new RSVP message that needs to be semantically fragmented. The fragment ID is reset to 1 when it reaches the maximum value of 65535. The scope of the fragment ID is limited to the RSVP message type (e.g. Path) carrying the fragment. In other words, fragment IDs do not have any correlation between different RSVP message types (e.g. Path and PathErr). The receiver does not check to ensure if the consecutive new RSVP messages (e.g. Path messages) are received with fragment IDs incremented by 1.

Fragments Total: 8-bit integer in the range of 1 to 255.

This value indicates the number of fragments sent for the given RSVP message. This value MUST be the same in all fragmented RSVP messages with a common Fragment ID.

Fragment Number: 8-bit integer in the range of 1 to 255.

This value indicates the position of this fragment in the given RSVP message.

The format of an S2L sub-LSP descriptor message is as follows:

```
<S2L sub-LSP descriptor> ::=
    [ <S2L_SUB_LSP_FRAG> ]
    <S2L_SUB_LSP>
    [ <P2MP SECONDARY_EXPLICIT_ROUTE> ]
```

The S2L_SUB_LSP_FRAG Object is added before adding the S2L_SUB_LSP Object in the semantically fragmented RSVP message.

6. Compatibility

The LSP_ATTRIBUTES Object has been defined in [RFC5420] and its message formats in [RFC6510] with class numbers in the form 1lbbbbbb, which ensures compatibility with non-supporting nodes. Per [RFC2205], nodes not supporting this extension will ignore the new flag defined for this Object in this document but forward it without modification.

The S2L_SUB_LSP_FRAG Object has been defined with class numbers in the form 1lbbbbbb, which ensures compatibility with non-supporting nodes. Per [RFC2205], nodes not supporting this Object will ignore the Object but forward it without modification.

7. IANA Considerations

IANA is requested to administer assignment of new values for namespace defined in this document and summarized in this section.

7.1. P2MP-TE Tree Re-evaluation Request Flag

IANA maintains a name space for RSVP-TE TE parameters "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters" (see <http://www.iana.org/assignments/rsvp-te-parameters>). From the registries in this name space "Attribute Flags", allocation of new flag is requested (Section 5.1).

The following new flag is defined for the Attributes Flags TLV in the LSP_ATTRIBUTES Object [RFC5420]. The numeric value is to be assigned by IANA.

o P2MP-TE Tree Re-evaluation Request Flag:

Bit No	Attribute Flag Name	Carried in Path	Carried in Resv	Carried in RRO or ERO	Reference
TBA1 by IANA	P2MP-TE Tree Re-evaluation	Yes	No	No	This document

7.2. Preferable P2MP-TE Tree Exists Path Error Sub-code

IANA maintains a name space for RSVP protocol parameters "Resource Reservation Protocol (RSVP) Parameters" (see <http://www.iana.org/assignments/rsvp-parameters>). From the sub-registry "Sub-Codes - 25 Notify Error" in registry "Error Codes

and Globally-Defined Error Value Sub-Codes", allocation of a new error code is requested (Section 5.2).

As defined in [RFC3209], the Error Code 25 in the ERROR SPEC Object corresponds to PathErr with Notify error. This document adds a new sub-code for this PathErr as follows:

- o Preferable P2MP-TE Tree Exists sub-code:

Sub-code value	Sub-code Description	PathErr Code	PathErr Name	Reference
TBA2 by IANA	Preferable P2MP-TE Tree Exists	25	Notify Error	This document

7.3. Fragment Identifier For S2L sub-LSP Descriptor

IANA maintains a name space for RSVP protocol parameters "Resource Reservation Protocol (RSVP) Parameters" (see <http://www.iana.org/assignments/rsvp-parameters>). From the registry "Class Names, Class Numbers, and Class Types", allocation of new Class-Num is requested (Section 5.3).

- o S2L_SUB_LSP_FRAG Object:

Class-Num value	Description	Reference
TBA3 by IANA	S2L_SUB_LSP_FRAG	This document

8. Security Considerations

This document defines RSVP-TE signaling extensions to allow an ingress node of a P2MP-TE LSP to request the re-evaluation of the LSP tree downstream of a node, and for a mid-point LSR to notify the ingress node of the existence of a preferable tree by sending a PathErr. As per [RFC4736], in the case of a P2MP-TE LSP S2L sub-LSP spanning multiple domains, it may be desirable for a mid-point LSR to modify the RSVP PathErr message defined in this document to preserve confidentiality across domains.

This document also defines fragment identifier for the S2L sub-LSP descriptor when combining large number of S2L sub-LSPs in an RSVP

message and the message needs to be semantically fragmented. The introduction of the fragment identifier, by itself, introduces no additional information to signaling. For a general discussion on MPLS and GMPLS related security issues, see the MPLS/GMPLS security framework [RFC5920].

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC4736] Vasseur, JP., Ikejiri, Y. and Zhang, R, "Reoptimization of Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Loosely Routed Label Switched Path (LSP)", RFC 4736, November 2006.
- [RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and Ayyangar, A., "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

9.2. Informative References

- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [RFC6510] Berger, L. and G. Swallow, "Resource Reservation Protocol (RSVP) Message Formats for Label Switched Path (LSP) Attributes Objects", RFC 6510, February 2012.

Acknowledgments

The authors would like to thank Loa Andersson, Sriganesh Kini, Curtis Villamizar, Dimitri Papadimitriou, Nobo Akiya, Vishnu Pavan Beeram and Joel M. Halpern for reviewing this document and providing many useful comments and suggestions. The authors would also like to thank Ling Zeng with Cisco Systems for implementing mechanisms defined in this document. A special thanks to Adrian Farrel for his thorough review of this document.

Author's Addresses

Tarek Saad (editor)
Cisco Systems

EMail: tsaad@cisco.com

Rakesh Gandhi (editor)
Cisco Systems

EMail: rgandhi@cisco.com

Zafar Ali
Cisco Systems

EMail: zali@cisco.com

Robert H. Venator
Defense Information Systems Agency

EMail: robert.h.venator.civ@mail.mil

Yuji Kamite
NTT Communications Corporation

EMail: y.kamite@ntt.com