

MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: July 28, 2014

Tarek Saad
Rakesh Gandhi
Zafar Ali
Cisco Systems, Inc.
Robert H. Venator
Defense Information Systems Agency
Yuji Kamite
NTT Communications Corporation
January 24, 2014

Reoptimization of Point-to-Multipoint Traffic Engineering
Loosely Routed LSPs

draft-tsaad-mpls-p2mp-loose-path-reopt-00.txt

Abstract

This document defines signaling extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for reoptimizing loosely routed point-to-multipoint (P2MP) Traffic Engineered (TE) Label Switched Path (LSP) in an Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

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) 2014 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	3
2. Terminology	4
2.1 Conventions used in this document	4
3. Procedure for Reoptimization of a Loosely Routed P2MP-TE LSP	4
4. RSVP Signaling Extensions	5
4.1. P2MP-TE Tree Re-evaluation Request	5
4.2. Preferable P2MP-TE Tree Exists Error sub-code	5
5. Compatibility	6
6. Security Considerations	6
7. IANA Considerations	6
7.1 P2MP-TE Tree Re-evaluation Request Flag	6
7.2 Preferable P2MP-TE Tree Exists sub-code	7
8. Acknowledgments	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Author's Addresses	8

1. Introduction

This document defines RSVP signaling extensions for the reoptimization of loosely routed point-to-multipoint (P2MP) MPLS and GMPLS (Generalized Multiprotocol Label Switching) Traffic Engineered (TE) Label Switched Path (LSP).

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 at the time of its signaling by the ingress node. Such an S2L sub-LSP is signaled with no Explicit Route Object (ERO), or with an ERO that contains at least one loose hop, or with an ERO that contains an abstract node that is not a simple abstract node (that is, an abstract node that identifies more than one node).

[RFC4736] defines RSVP signaling extensions for reoptimizing loosely routed P2P TE LSP(s). Specifically, an ingress node sends a "path re-evaluation request" to a border node by setting a flag (0x20) in SESSION_ATTRIBUTES object in the Path message. A border node sends a PathErr code 25 (notify error defined in [RFC3209]) with sub-code 6 to indicate "preferable path exists" to the ingress node. The ingress node upon receiving this PathErr initiates reoptimization of the LSP.

As per [RFC4875], an ingress node may reoptimize the entire P2MP-TE LSP by resignaling all its S2L sub-LSP(s) or may reoptimize individual S2L sub-LSP(s) i.e. individual destination(s).

[RFC4736] does not define signaling extensions specific for reoptimizing entire P2MP-TE LSP tree. Mechanisms defined in [RFC4736] can be used for signaling the reoptimization of individual S2L sub-LSP(s). However, to use [RFC4736] mechanisms for reoptimizing an entire P2MP-TE LSP tree, an ingress node needs to send the query on all (typically 100s of) S2L sub-LSPs and a border node needs to notify PathErrs for all S2L sub-LSPs. In addition, a border node has to accumulate the received queries on all S2L sub-LSPs (using a wait timer) and interpret them as a reoptimization request for the P2MP-TE LSP tree. Furthermore, when the ingress node gradually receives unsolicited PathErr(s) notifications for individual S2L sub-LSP(s), it may prematurely start reoptimizing these sub-set of sub-LSPs. However, as mentioned in [RFC4875] Section 14.2, such reoptimization procedure may result in data duplication that can be avoided if the entire P2MP-TE LSP tree is reoptimized, especially if the ingress node eventually receives PathErr(s) notifications for all S2L sub-LSP(s) of the P2MP-TE LSP tree. In such cases, the ingress node may have to heuristically determine when to perform P2MP-TE LSP tree reoptimization or per S2L sub-LSP reoptimization, for example, to

wait long enough time to accumulate all PathErr(s) to be received. Such methods may produce undesired results that can be avoided by the proposed RSVP signaling extensions in this draft.

This document defines required RSVP signaling extensions to query and notify for reoptimizing loosely routed P2MP-TE LSP tree.

2. Terminology

ABR: Area Border Router.

ERO: Explicit Route Object.

TE LSP: Traffic Engineering Label Switched Path.

TE LSP ingress: head/source of the TE LSP.

TE LSP egress: tail/destination of the TE LSP.

S2L: Source-to-leaf.

Interior Gateway Protocol Area (IGP Area): OSPF Area or IS-IS level.

Inter-area TE LSP: A TE LSP whose path transits across at least two different IGP areas.

Inter-AS MPLS TE LSP: A TE LSP whose path transits across at least two different Autonomous Systems (ASes) or sub-ASes (BGP confederations).

2.1 Conventions used in this document

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

3. Procedure for Reoptimization of a Loosely Routed P2MP-TE LSP

As per [RFC4875], an ingress node may prefer to reoptimize the entire P2MP-TE LSP by resignaling all its S2L sub-LSP(s) (Section 14.1, "Make-before-Break") or reoptimize individual S2L sub-LSP(s) i.e. individual destination(s) (Section 14.2 "Sub-Group-Based Re-Optimization").

Procedures defined in [RFC4736] are used by an ingress node to reoptimize the S2L sub-LSP individually.

To reoptimize entire P2MP-TE LSP tree, in order to query border nodes to check if a preferable P2MP-TE LSP tree exists, an ingress node sends a Path message with "P2MP-TE Tree Re-evaluation Request" defined in this document.

A border node receiving the "P2MP-TE Tree Re-evaluation Request" checks for a preferable P2MP-TE LSP tree by re-evaluating loosely expanded paths for all S2L sub-LSP(s) of the P2MP-TE LSP. If a preferable P2MP-TE LSP tree is found, the border node immediately sends an RSVP PathErr to the ingress node with Error code 25 (Notify defined in [RFC3209] and Error sub-code defined in this document "Preferable P2MP-TE Tree Exists". At this point, the border node does not propagate this bit in subsequent RSVP Path messages sent downstream for the re-evaluated TE LSP. The sending of an RSVP PathErr Notify message "Preferable P2MP-TE Tree Exists" to the ingress node will notify the ingress node of the existence of a preferable P2MP-TE LSP tree. If no preferable path can be found, the recommended mode is for the border node to relay the request (by setting the "P2MP-TE Tree Re-evaluation Request" bit in the LSP_ATTRIBUTES TLV of RSVP path message sent downstream).

A border node MAY also send "Preferable P2MP-TE Tree Exists" with PathErr code 25 to the ingress node to reoptimize the entire P2MP-TE LSP tree with an unsolicited PathErr message.

4. RSVP Signaling Extensions

4.1. P2MP-TE Tree Re-evaluation Request

In order to query border nodes to check if a preferable P2MP-TE LSP tree exists, a new flag is defined in Attributes Flags TLV of the LSP_ATTRIBUTES object [RFC5420] as follows:

Bit Number (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 an S2L sub-LSP and is inserted by the ingress node.

4.2. Preferable P2MP-TE Tree Exists Error sub-code

In order to indicate to an ingress node that a preferable P2MP-TE LSP tree is available, following new sub-code for PathErr code 25 (Notify Error) [RFC3209] is defined:

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

When a preferable P2MP-TE LSP tree is found, the border node MUST send "Preferable P2MP-TE Tree Exists" PathErr to the ingress node in order to reoptimize the entire P2MP-TE LSP.

5. Compatibility

The LSP_ATTRIBUTES TLV has been defined in [RFC5420] with class numbers in the form 11bbbbbb, which ensures compatibility with non-supporting nodes. Per [RFC2205], nodes not supporting this extension will ignore the new flag defined in this document but forward it, unexamined and unmodified, in all messages resulting from this message.

6. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC3209] and [RFC4875].

7. IANA Considerations

IANA maintains a name space for RSVP-TE TE parameters "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters". From the registries in this name space "Attribute Flags" allocation of new flag is requested (sections 4.1).

IANA also maintains a name space for RSVP protocol parameters "Resource Reservation Protocol (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 4.2).

7.1 P2MP-TE Tree Re-evaluation Request Flag

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

- o P2MP-TE Tree Re-evaluation Request Flag:
 - Bit Number: To be assigned by IANA.
 - Attribute flag carried in Path message: Yes
 - Attribute flag carried in Resv message: No

7.2 Preferable P2MP-TE Tree Exists sub-code

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

- o Preferable P2MP-TE Tree Exists sub-code:
 - Sub-code for Notify PathErr code 25. To be assigned by IANA.

8. Acknowledgments

TBA.

8. References

8.1. Normative References

- [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.
- [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 A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

8.2. Informative References

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

Author's Addresses

Tarek Saad
Cisco Systems

Email: tsaad@cisco.com

Rakesh Gandhi
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