

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 17, 2014

Z. Li
T. Huang
Huawei Technologies
L. Chen
Ericsson
October 14, 2013

Alternative Constraints for Point-to-Multipoint Traffic-Engineered MPLS
Label Switched Path(LSP)
draft-li-mpls-p2mp-te-alt-path-01

Abstract

The document proposes a solution to be able to set up the alternative path for specific leaf nodes of a P2MP TE LSP. Corresponding RSVP-TE protocol extension is also defined. The solution is used to cope with the issue that in some scenarios traffic loss happens even if there exists possible path for the leaf nodes.

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 RFC 2119 [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 April 17, 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. Terminology	3
3. Problem Statement	3
4. Mechanisms	4
4.1. Path Computation in Root Node	4
4.2. Alternative Constraints Propagation	5
4.3. Resource and Label	5
5. Method of Separate Messages	5
6. Method of Single message	6
6.1. Path Message Format	6
6.2. Path Message Processing	7
6.3. Other Messages	7
7. IANA Considerations	8
8. Security Considerations	8
9. Normative References	8
Authors' Addresses	9

1. Introduction


[RFC4461] presents a set of requirements for the establishment and maintenance of Point-to-Multipoint (P2MP) Traffic-Engineered (TE) Multi-protocol Label Switching (MPLS) Label Switched Paths (LSPs). [RFC4875] defines extensions to the RSVP-TE protocol for setup of P2MP TE LSPs. P2MP TE LSPs are set up with a series of traffic engineering constraints. These constraints are applied to all S2L sub-LSPs. This may cause the issue that some S2L sub-LSPs can be set up while others cannot set up according to the constraints. There may be worse case that some S2L sub-LSPs cannot be restored after link failure according to the constraints. When P2MP TE LSPs are used for specific applications, it will cause continuous traffic loss. This document identifies the applicability issue and proposes the solution and corresponding protocol extension.

2. Terminology

This document uses terminologies defined in [RFC2205], [RFC3031], [RFC3209], [RFC3473], [RFC4090], [RFC4461] and [RFC4875].

3. Problem Statement

The P2MP TE LSP is set up with a series of traffic engineering constraints such as bandwidth, explicit path, affinity property(color), etc. These traffic engineering constraints are applied to path computation for all S2L sub-LSPs. Owing to the network provision some leaves of the P2MP LSP are not reachable according to the required constraints (it will be called primary constraints in the following text). There may be the worse case that all leaves are reachable at the beginning and they are not reachable when failure happens. In fact in the scenario these leaves can be reachable if ignore some or all of the primary constraints .



N M*****O

Figure 1. Constraints for P2MP TE LSP

An example for P2MP TE LSP setup is shown in the figure 1. A is the root node and F, N and M are leaf nodes. The link with '|' means the link with red color and the link with '*' means the link with green color. The constraint is that the link with red color should be chosen for the path. For the leaf node M, the path is A->B->E->H->K-M. When link between H and K fails, there is no path with red color can be found from A to M. This will cause the initial available traffic break until the link between H and K restores. The continuous traffic loss can cause bad user experience if the P2MP TE LSP is used for IPTV or other applications. In fact, during the course of failure, there is an alternative path from A to M (A->B->E->H->I->L->K->M) if the link with green color can be chosen.

4. Mechanisms

In order to solve the above applicability issue for P2MP TE LSP, alternative constraints can be specified for the P2MP TE LSP to calculate paths to specific leaf nodes if the path with the primary constraints is not available. The P2MP TE LSP is set up with some S2L sub-LSPs using the primary constraints while the other S2L sub-LSPs using the alternative constraints. The constraints may be used in the downstream nodes, such as ASBR node, and the alternative constraints MUST be propagated to keep the consistence through RSVP-TE protocol extensions.

4.1. Path Computation in Root Node

When alternative constraints is allowed for a specific P2MP TE LSP in the root node, the node MUST try to compute paths for all leaf nodes using the primary constraints. If paths with the primary constraints are available for all leaf nodes, the alternative constraints MUST NOT be used.

When paths with the primary constraints are not available for specific leaf nodes, the alternative constraints SHOULD be used to calculate paths for these leaf nodes. In order to get available paths, the alternative constraints should be looser than the primary constraints. The alternative constraints can be set as zero to simplify the process and the best-effort path as routing is calculated.

When calculate paths with the alternative constraints, the constraints MUST be applied to the whole S2L sub-LSP. That is, it is

prohibited that some parts of the S2L sub-LSP satisfies the primary constraints while other parts satisfies the alternative constraints. If the root node can not calculate the whole S2L sub-LSP (abstract node exists in the calculated path), the alternative constraints MUST be used in the downstream nodes path calculation.

The root node will keep trying to re-optimize to a better path to meet the primary constraints, and it is outside the scope of this document.

4.2. Alternative Constraints Propagation

When setup P2MP LSP, the primary constraint is carried according to the RSVP-TE protocol extension which is defined in [RFC4875]. If the paths to specific leaf nodes are computed using alternative constraints, the alternative constraints MUST be carried corresponding to the S2L sub-LSPs to these leaf nodes in the Path message. These alternative constraints corresponding to S2L sub-LSPs are propagated along the paths from the root node to the leaf nodes.

There are two methods for RSVP-TE protocol to propagate the alternative constraints. One is to propagate alternative constraints in separate message from primary constraints. This method can reuse current P2MP RSVP-TE Message, and does not introduce any extension. The other method is to propagate primary and alternative constraints in single RSVP Message, and need some extension on the Path Message.

When alternative constraints are received for one or more S2L sub-LSPs, they MUST be used when calculating for those S2L sub-LSPs, while the primary constraints MUST be used for other S2L sub-LSPs without alternative constraints. This will be described in detail in the section 5 and 6.

4.3. Resource and Label

When the Resv message is propagated from the leaf nodes to the root node, the transit node MUST reserve resource according to the traffic parameters specified by the required constraints. However, the common upstream node, such as A, B node in figure 1, may have different traffic parameters required if both the primary and alternative constraints exist. But no matter the parameters are same or different, all sub-LSPs in one P2MP LSP MUST share the resource and use same incoming Label on the common nodes.

5. Method of Separate Messages

Propagating alternative constraints through separate messages does not need to introduce any extension on RSVP messages based

on[RFC4875]. However, it needs to change on Path and Resv Message processing. According to [RFC4875], the constraints for all sub-LSPs that belongs to one P2MP LSP should be the same. This document introduces that sub-LSPs can have different constraints in the same P2MP LSP. In this case, a node supporting alternative sub-LSPs MUST accept such different constraints for local processing and continue to propagate them to downstream nodes. The resource reservation and Label processing are as described in Section 4.3.

Exception for the LSP attributes defined by alternative constraints, the S2L sub-LSP descriptors and Sub-Group identifier, the separate Path Message has the same objects with other Path messages for same P2MP LSP.

If a node cannot support alternative sub-LSPs, it MUST send PathErr Message back to Ingress and stop the establishment for such sub-LSPs. But other sub-LSPs with primary constraints SHOULD not be impacted.

6. Method of Single message

This method needs to extend Path Message based on [RFC4875] to carry both primary and alternative constraints in single message.

6.1. Path Message Format

```
<Path Message> ::=
    <Common Header> [ <INTEGRITY> ]
    [ [ <MESSAGE_ID_ACK> | <MESSAGE_ID_NACK> ] ... ]
    [ <MESSAGE_ID> ]
    <SESSION> <RSVP_HOP>
    <TIME_VALUES>
    [ <EXPLICIT_ROUTE> ]
    <LABEL_REQUEST>
    [ <PROTECTION> ]
    [ <LABEL_SET> ... ]
    [ <SESSION_ATTRIBUTE> ]
    [ <NOTIFY_REQUEST> ]
    [ <ADMIN_STATUS> ]
    [ <POLICY_DATA> ... ]
    <sender descriptor>
    [ <S2L sub-LSP descriptor list> ]
```

The following is the format of the S2L sub-LSP descriptor list.

```
<S2L sub-LSP descriptor list> ::= <S2L sub-LSP descriptor>
    [ <S2L sub-LSP descriptor list> ]

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

```
[ <P2MP SECONDARY_SESSION_ATTRIBUTE> ]  
[ <P2MP SECONDARY_SENDER_TSPEC> ]
```

In the Path message, S2L_SUB_LSP for specific leaf nodes can carry the alternative constraints besides the explicit route. <P2MP SECONDARY_SESSION_ATTRIBUTE> and <P2MP SECONDARY_SENDER_TSPEC> are added to specify the alternative constraints such as resource affinity, setup and holding priority and traffic parameters. The format, Class Num and C-Type of <P2MP SECONDARY_SESSION_ATTRIBUTE> and <P2MP SECONDARY_SENDER_TSPEC> are all the same as <SESSION_ATTRIBUTE> defined by [RFC3209] and <SENDER_TSPEC> defined by [RFC2210]. The downstream node can judge that the SESSION_ATTRIBUTE and SENDER_TSPEC objects are for alternative constraints of specific S2L sub-LSP when they are placed following corresponding S2L_SUB_LSP object. For convenience, we still use the names, P2MP SECONDARY_SESSION_ATTRIBUTE and P2MP SECONDARY_SENDER_TSPEC, to represent these two objects for specific sub-LSPs.

6.2. Path Message Processing

When a node receives a Path Message with P2MP SECONDARY_SESSION_ATTRIBUTE and P2MP SECONDARY_SENDER_TSPEC objects following one or more S2L_SUB_LSP objects, it can judge that such sub-LSPs are alternative sub-LSPs which have attributes identified by these two objects.

If after a branch node, the alternative sub-LSP will become alone, then the branch node will signal a new Path Message for that alternative sub-LSP in the normal way. This means, for this new path message, the content of P2MP SECONDARY_SESSION_ATTRIBUTE and P2MP SECONDARY_SENDER_TSPEC objects will be carried by the primary SESSION_ATTRIBUTE and SENDER_TSPEC like a normal P2MP Path Message, and these two new objects will not be carried any more to downstream. The SUB-Group ID for that path message will also be a new value different from the original Primary sub-LSP for the same egress.

If a transit node cannot support alternative sub-LSPs, it MUST send a PathErr Message back to ingress.

6.3. Other Messages

The format of Resv Message based on [RFC4875] does not need to be modified. But a new case for Resv Message processing is introduced that, a branch node may receive different traffic parameters in FLOWSPEC of the same P2MP LSP from different downstream nodes. It

MUST calculate the shared resource for resource reservation and carry the result as FLOWSPEC to upstream.

For other RSVP Messages based on [RFC4875], the message format and processing have no change.

7. IANA Considerations

TBD.

8. Security Considerations

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

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC4461] Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 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.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Tieying Huang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: huangtieying@huawei.com

Lei Chen
Ericsson
CDK building, No.1 Wangjing North Rd.
Beijing 100102
China

Email: charles.c.chen@ericsson.com