

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2011

H. Chen
Huawei Technology, Inc.
N. So
Verison Business
July 12, 2010

Extensions to RSVP-TE for P2MP LSP Egress Local Protection
draft-chen-mpls-p2mp-egress-protection-01.txt

Abstract

This document describes extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for locally protecting egress nodes of Traffic Engineered (TE) point-to-multipoint (P2MP) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

Status of this Memo

This Internet-Draft is submitted to IETF 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 January 13, 2011.

Copyright Notice

Copyright (c) 2010 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	3
3.	Conventions Used in This Document	3
4.	Mechanism	3
4.1.	An Example of Egress Local Protection	4
4.2.	Set up of Backup P2MP sub LSP	5
4.3.	Forwarding State for Backup P2MP sub LSP	5
4.4.	Detection of Failure in Egress	6
5.	Representation of a backup P2MP Sub LSP	6
5.1.	EGRESS_BACKUP_P2MP_SUB_LSP Object	7
5.1.1.	EGRESS_BACKUP_P2MP_SUB_LSP IPv4 Object	7
5.1.2.	EGRESS_BACKUP_P2MP_SUB_LSP IPv6 Object	8
5.2.	EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE Object	8
6.	Path Message	8
6.1.	Format of Path Message	9
6.2.	Processing of Path Message	9
7.	IANA Considerations	10
8.	Acknowledgement	10
9.	References	10
9.1.	Normative References	10
9.2.	Informative References	11
	Authors' Addresses	11

1. Introduction

"Fast Reroute Extensions to RSVP-TE for LSP Tunnels" [RFC 4090](#) describes two methods for protecting P2P LSP tunnels or paths at local repair points. For a P2P LSP, the local repair points are the intermediate nodes between the ingress node and the egress node of the P2P LSP. The first method is a one-to-one protection method, where a detour backup P2P LSP for each protected P2P LSP is created at each potential point of local repair. The second method is a facility bypass backup protection method, where a bypass backup P2P LSP tunnel is created using MPLS label stacking to protect a potential failure point for a set of P2P LSP tunnels. The bypass backup tunnel can protect a set of P2P LSPs that have similar backup constraints.

"Extensions to RSVP-TE for P2MP TE LSPs" [RFC 4875](#) describes how to use the one-to-one protection method and facility bypass backup protection method to protect a link or intermediate node failure on the path of a P2MP LSP. However, there is no mention of locally protecting an egress node failure in a protected P2MP LSP.

This document defines extensions to RSVP-TE for locally protecting an egress node of a Traffic Engineered (TE) point-to-multipoint (P2MP) Label Switched Path through using a backup P2MP sub LSP.

2. Terminology

This document uses terminologies defined in [RFC 2205](#), [RFC 3031](#), [RFC 3209](#), [RFC 3473](#), [RFC 4090](#), [RFC 4461](#), and [RFC 4875](#).

3. 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 [RFC 2119](#).

4. Mechanism

This section briefly describes a solution that locally protects an egress node of a P2MP LSP through using a backup P2MP sub LSP. We first show an example, and then present different parts of the solution, which includes the creation of the backup P2MP sub LSP, the forwarding state for the backup P2MP sub LSP, and the detection of a failure in the egress node.

4.1. An Example of Egress Local Protection

The figure 1 illustrates an example of using a backup P2MP sub LSP to locally protect an egress of a P2MP LSP. The P2MP LSP to be protected is from ingress node R1 to three egress/leaf nodes: L1, L2 and L3. The P2MP LSP is represented by double lines in the figure.

La, Lb and Lc are the designated backup egress/leaf nodes for the egress/leaf nodes L1, L2 and L3 of the P2MP LSP respectively. The backup P2MP sub LSP used to protect the egress node L1 is from the previous hop node R3 of L1 to the backup egress node La. The backup P2MP sub LSP used to protect the egress node L2 is from the previous hop node R5 of L2 to the backup egress node Lb. The backup P2MP sub LSP used to protect the egress node L3 is from the previous hop node R5 of L3 to the backup egress node Lc via intermediate node Rc.

At a previous hop node such as R3 of an egress node such as L1 of the P2MP LSP, the traffic transported by the P2MP LSP is forwarded to the egress node such as L1 in a normal operation, which delivers the traffic towards its destination such as CE1. When the failure in an egress node such as L1 is detected, the previous hop node such as R3 of the egress node such as L1 forwards the traffic toward the corresponding backup egress node such as La, which delivers the traffic towards its destination such as CE1.

There may be a BFD session between an egress node such as L1 and the previous hop node such as R3 of the egress node such as L1. The previous hop node uses this BFD session to detect the failure of the egress node. When it detects the failure of the egress node, it forwards the traffic carried by the P2MP LSP into the backup P2MP sub LSP to the corresponding backup egress node. The traffic from the sub LSP is delivered from the backup egress node towards its destination.

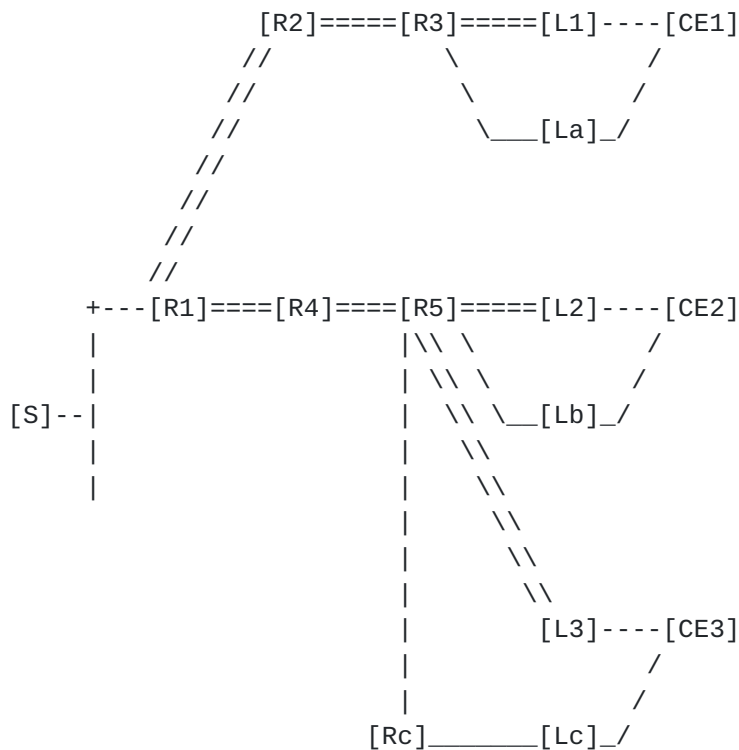


Figure 1: P2MP sub LSP for Locally Protecting Egress

4.2. Set up of Backup P2MP sub LSP

For an egress node of a P2MP LSP, a backup egress node is designated to protect the egress node. The previous-hop node of the egress node of the P2MP LSP sets up a backup P2MP sub LSP from itself to the backup egress node after receiving the information about the backup egress node.

The previous-hop node sets up the backup P2MP sub LSP, creates and maintains its state in the same way as setting up a P2MP S2L sub LSP from the signalling's point of view. It constructs and sends a RSVP-TE PATH message along the path for the backup P2MP sub LSP, receives and processes a RSVP-TE RESV message that responds to the PATH message.

4.3. Forwarding State for Backup P2MP sub LSP

The forwarding state for the backup P2MP sub LSP is different from that for a P2MP S2L sub LSP. After receiving the RSVP-TE RESV message for the backup P2MP sub LSP, the previous-hop node creates a forwarding entry with an inactive state or flag. This forwarding entry with an inactive state or flag is called an inactive forwarding entry. In a normal operation, this inactive forwarding entry is not

used to forward any data traffic. However, the forwarding entry for a P2MP S2L sub LSP is with an active state or flag, and used to forward the data traffic.

When a failure of the egress node happens, the state or flag of the forwarding entry for the backup P2MP sub LSP is set to be active. Thus, on the previous-hop node of the egress node, the data traffic will be forwarded to the backup egress node instead of to the egress node through the backup P2MP sub LSP from the P2MP LSP. From the backup egress node, the data traffic is sent towards its destination.

4.4. Detection of Failure in Egress

There are a number of failures in an egress node of a P2MP LSP. The failures in the egress that the previous hop node of the egress node should detect include two classes of failures. One class of failures is such a failure that the traffic can not be delivered to the egress node of the P2MP LSP. The death of the signalling protocol MPLS RSVP-TE in the egress node is such a failure. The death of the egress node and the failure of the link between the egress node and the previous hop node also belong to this class of failures.

Another class of failures are such failures that the egress node can not deliver the traffic from the P2MP LSP towards its destination. The failure of the link over which the traffic is delivered towards its destination such as CE1 is such a failure.

After a previous hop node detects any above failure in the egress node, it imports the traffic from the P2MP LSP into the backup P2MP sub LSP. The traffic from the backup P2MP sub LSP is delivered towards its destination at the backup egress node.

5. Representation of a backup P2MP Sub LSP

A backup P2MP sub LSP exists within the context of a P2MP LSP in a way similar to a P2MP S2L sub LSP. It is identified by the P2MP ID, Tunnel ID, and Extended Tunnel ID in the P2MP SESSION object, the tunnel sender address and LSP ID in the P2MP SENDER_TEMPLATE object, and the backup P2MP sub LSP destination address in the EGRESS_BACKUP_P2MP_SUB_LSP object. The EGRESS_BACKUP_P2MP_SUB_LSP object is defined in the section below.

An EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE Object (EB-SERO) is used to optionally specify the explicit route of a backup P2MP sub LSP that is from a previous-hop node to a backup egress node. The EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE object is defined in the following section.

The class of the EGRESS_BACKUP_P2MP_SUB_LSP IPv4 object is the same as that of the S2L_SUB_LSP IPv4 object defined in [RFC 4875](#). The C-Type of the EGRESS_BACKUP_P2MP_SUB_LSP IPv4 object is a new number 3, or may be another number assigned by Internet Assigned Numbers Authority (IANA).

This section describes extensions to the Path message defined in [RFC 4875](#). The Path message is enhanced to transport the information about a backup egress node to the previous-hop node of an egress node of a P2MP LSP through including an egress backup P2MP sub LSP descriptor list.

6.1. Format of Path Message

The format of the enhanced Path message is illustrated below.

```

<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> ]
                    [ <egress backup P2MP sub LSP descriptor list> ]

```

The format of the egress backup P2MP sub LSP descriptor list in the enhanced Path message is defined as follows.

```

<egress backup P2MP sub LSP descriptor list> ::=
    <egress backup P2MP sub LSP descriptor>
    [ <egress backup P2MP sub LSP descriptor list> ]

<egress backup P2MP sub LSP descriptor> ::=
    <EGRESS_BACKUP_P2MP_SUB_LSP>
    [ <EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE> ]

```

6.2. Processing of Path Message

The ingress node of a P2MP LSP initiates a Path message with an egress backup P2MP sub LSP descriptor list for protecting egress nodes of the P2MP LSP. In order to protect an egress node of the P2MP LSP, the ingress node MUST add an EGRESS_BACKUP_P2MP_SUB_LSP object into the Path message. The object contains the information about the backup egress node to be used to protect the failure of the egress node. An EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE object, which describes an explicit path to the backup egress node, SHOULD follow the EGRESS_BACKUP_P2MP_SUB_LSP.

After an intermediate node (a transit or branch node) receives the

Path message with an egress backup P2MP sub LSP descriptor list, for each EGRESS_BACKUP_P2MP_SUB_LSP containing a backup egress node in the list, the intermediate node of the P2MP LSP MUST put the EGRESS_BACKUP_P2MP_SUB_LSP with the directly following EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE into the Path message that is to be sent toward the direction to the previous-hop node of the egress node that is to be protected by the backup egress node if the intermediate node is not the previous-hop node of the egress node.

If the intermediate node is the previous-hop node of an egress node, when it receives the Path message with the EGRESS_BACKUP_P2MP_SUB_LSP containing the backup egress node to be assigned for protecting the egress node, the intermediate node generates a new Path message based on the information in the EGRESS_BACKUP_P2MP_SUB_LSP and the possible directly following EGRESS_BACKUP_P2MP_SECONDARY_EXPLICIT_ROUTE. The format of this new Path message is the same as that of the Path message defined in [RFC 4875](#). This new Path message is used to signal the segment of a special S2L sub-LSP of the P2MP LSP from the previous-hop node to the backup egress node. The new Path message is sent to the next-hop node along the path for the backup P2MP sub LSP.

When an egress node of the P2MP LSP receives the Path message with an egress backup P2MP sub LSP descriptor list, it SHOULD ignore the egress backup P2MP sub LSP descriptor list and generate a PathErr message.

7. IANA Considerations

TBD

8. Acknowledgement

The author would like to thank Quintin Zhao for his valuable comments on this draft.

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.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.

- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), 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.
- [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.

9.2. Informative References

- [RFC4461] Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", [RFC 4461](#), April 2006.

Authors' Addresses

Huaimo Chen
Huawei Technology, Inc.
Boston, MA
US

Email: Huaimochen@huawei.com

Ning So
Verison Business
2400 North Glenville Drive
Richardson, TX 75082
USA

Email: Ning.So@verizonbusiness.com