

MPLS
Internet-Draft
Expires: April 22, 2010

R. Torvi
K. Chan, Ed.
Huawei Technologies
C. Jacquenet
France Telecom
October 19, 2009

**Receiver Driven Point-To-Multi-Point Traffic Engineered Label Switched
Paths
draft-chan-torvi-jacquenet-mpls-rd-p2mp-te-00**

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 22, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

For content delivery services that rely upon the IP multicast transmission scheme, the distribution trees are receiver initiated. The delivery of such services over MPLS networking infrastructures may rely upon P2MP LSP tree structures that are sender initiated, with the root of the P2MP tree being at the LSR router directly connected to the sender. This document describes a mechanism that aims at establishing MPLS P2MP tree structures that are receiver initiated. This mechanism builds on the works of Point-to-MultiPoint Traffic Engineered Label Switched Paths (P2MP-TE LSPs). This mechanism can also be used to establish receiver driven BiDirectional P2MP TE LSPs.

Table of Contents

1.	Introduction	3
1.1.	Motivation	3
1.2.	Terminology	3
1.3.	Overview	4
2.	Signaling Protocol Extensions	6
2.1.	Path Message Extensions	6
2.2.	Resv Message Extensions	7
2.3.	PathErr Message Extensions	8
2.4.	ResvErr Message Extensions	8
2.5.	PathTear Message Extensions	9
3.	Broadcast Interfaces	9
4.	Fast Re-Route Considerations	9
5.	Re-Merging and Crossover	10
6.	Receiver-Driven Bidirectional LSPs	10
7.	Backward Compatibility	11
8.	Security Implications	11
9.	IANA Considerations	12
10.	Acknowledgements	12
11.	References	12
11.1.	Normative References	12
11.2.	Informative References	13
	Authors' Addresses	13

1. Introduction

Multiparty multimedia applications are getting greater attention in the telecom world. Such applications are QoS-demanding and can therefore benefit from the activation of MPLS traffic engineering capabilities that lead to the dynamic computation and establishment of MPLS LSPs whose characteristics comply with application-specific QoS requirements. P2MP-TE [[RFC4875](#)] defines the procedure to setup multipoint LSPs from sender to receiver. This document extends P2MP-TE to be used for receiver-driven P2MP-TE LSPs. This document complements P2MP-TE [[RFC4875](#)], allowing P2MP-TE LSPs to be established either from a sender or from a receiver, unidirectional or bidirectional.

1.1. Motivation

IP multicast distribution trees are receiver-initiated and dynamic by nature. IP multicast-enabled applications are also bandwidth savvy, especially in the area of residential IPTV services, where several hundreds of thousands of IPTV receivers need to be served with the appropriate level of quality. Current source-driven P2MP LSP establishment assumes a prior knowledge of receiver(s) location for the sake of the P2MP LSP tree structure's forwarding efficiency. But the receiver's location information is not available a priori for the root MPLS router to compute and establish the relevant P2MP tree structure.

Receiver-driven MPLS P2MP tree structure do not require sender to maintain/discover receiver information a priori, and their design should better reflect the receiver-specific QoS conditions, such as network access capabilities.

1.2. Terminology

With the receiver-driven concept, we have re-defined the following terms:

- o Sender: Sender refers to the Originator (and hence) Sender of the content/payload. As in [[RFC2205](#)].
- o Receiver: Receiver refers to the Receiver of the content/payload. As in [[RFC2205](#)].
- o Upstream: The direction of flow from content Receiver toward content Sender. As defined in [[RFC2205](#)].
- o Downstream: The direction of flow from content Sender toward content Receiver. As defined in [[RFC2205](#)].

- o Path-Sender: The sender of the RSVP PATH message, with NO correlation to direction of content/payload flow. All other control messages flow direction discussed in this document will use this as the reference.
- o Path-Receiver: The receiver of the RSVP PATH message, with NO correlation to direction of content/payload flow.
- o Path-Initiator: The Path-Sender that originated the RSVP PATH message. This being different from Path-Sender because an intermediate node can be a Path-Sender, but different from the node that created and initiated the RSVP PATH message, the Path-Initiator.
- o Path-Terminator: The Path-Receiver that does NOT propagate the Path message. This being different from Path-Receiver because an intermediate node can be a Path-Receiver.

1.3. Overview

Although receiver-driven P2MP LSPs as defined in this document use existing sender-driven syntax, there are important semantic differences that need to be defined for correct interpretation and interoperability. In the receiver-driven approach, we inverted the semantics of P2MP-TE RSVP [[RFC4875](#)] messages, while keeping the syntax unchanged.

Following are some key differences that are specific to the receiver-driven paradigm:

1. Receiver initiates RSVP PATH message towards the one or more senders. We keep same convention with respect to data flows, which are opposite to control flows.
2. S2L Destinations (the leaves) are ingress routers where user data payload traffic enter the LSP.
3. RSVP P2MP PATH messages traverse from receiver to senders.
4. RSVP P2MP RESV messages traverse from sender to receiver.
5. A node receiving a RSVP RESV message is interpreted as successful resource reservation from the upstream node.
6. A node receiving a RSVP PATH message would first allocate required resources on the interface through which the RSVP PATH message is received, before sending the RSVP PATH message upstream. So that the upstream node can send traffic soon after

successfully reserving resources on the downstream link, on which the RSVP PATH message is received.

7. Label allocation on incoming interface is done prior to sending RSVP PATH messages upstream. The syntax details are defined in [Section 2](#).

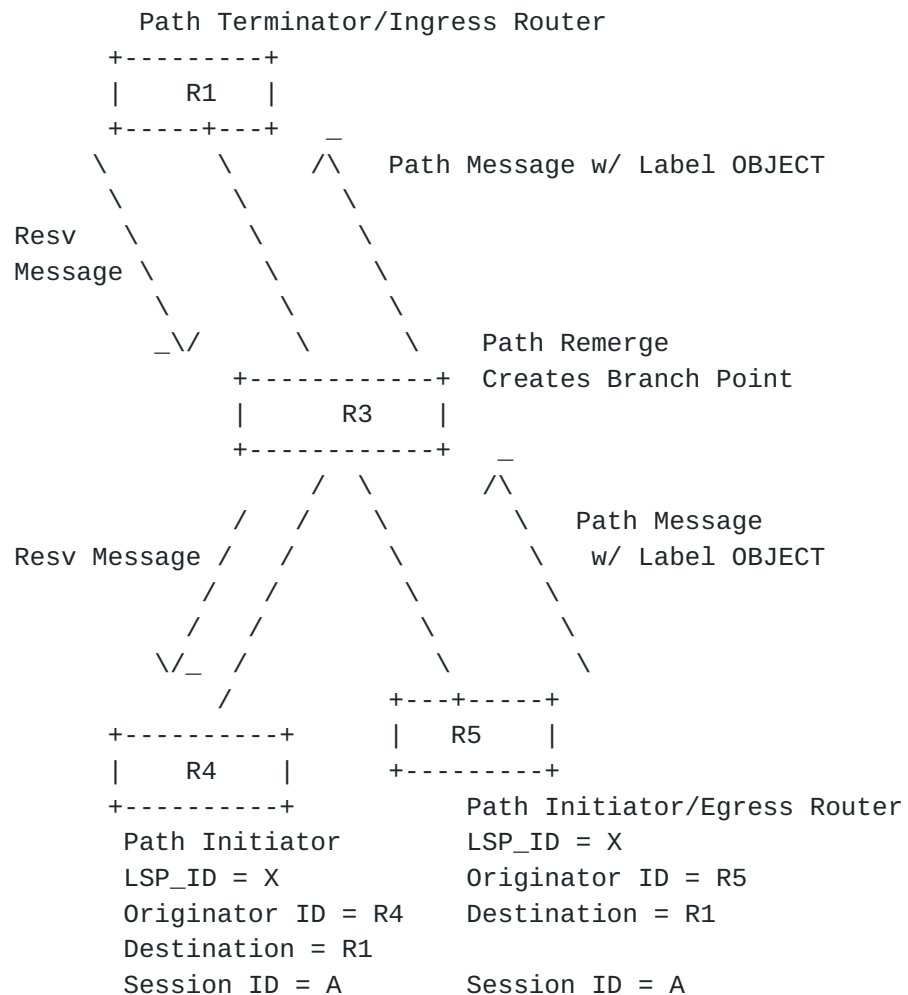


Figure 1: Receiver-Driven P2MP RSVP-TE LSP Overview

2. Signaling Protocol Extensions

Receiver-driven P2MP MPLS-TE LSP uses the RSVP-TE protocol as specified in [[RFC4875](#)], [[RFC3473](#)], and [[RFC3209](#)], but unlike what is specified in [[RFC4875](#)], receivers initiate the RSVP PATH messages toward the sender.

With receiver-driven P2MP MPLS-TE LSP, the content Receiver is the Path-Originator. The RSVP RESV messages flow in the opposite direction as compared to the RSVP PATH messages, i.e. RSVP RESV messages are generated by the content Sender or the MPLS router it is directly attached to. All other RSVP messages flow in reference to this picture.

Within this receiver-driven context, the processing of receiver-initiated P2MP RSVP-TE messages should be differentiated from the other RSVP messages. Following the method used by RSVP-TE and P2MP RSVP-TE, this document recommends the use of new SESSION C-Type as follows:

Class Name = SESSION

C-Type

XX+0	RcvrD_P2MP_LSP_TUNNEL_IPv4 C-Type
XX+1	RcvrD_P2MP_LSP_TUNNEL_IPv6 C-Type
XX+2	BiDi_P2MP_LSP_TUNNEL_IPv4 C-Type
XX+4	BiDi_P2MP_LSP_TUNNEL_IPv6 C-Type

Where XX is a number allocated by IANA.

The new SESSION C-Type MUST be used in all receiver-driven P2MP RSVP-TE messages.

The following sections describe the receiver-driven P2MP RSVP-TE extensions to the P2MP RSVP-TE protocol. When there is no difference in the protocol, usage of [[RFC4875](#)] is assumed.

2.1. Path Message Extensions

Receiver-driven P2MP MPLS-TE LSP uses the Path message to carry the LABEL object upstream towards the Sender. With receiver-driven usage of the RSVP PATH message, the LABEL_REQUEST object carried by the PATH message is no longer mandatory, it becomes optional for receiver-driven PATH messages, as indicated 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>]
```

Using [\[RFC4875\]](#) as the base specification, with the LABEL object being added to the SENDER DESCRIPTOR object:

```
<sender descriptor> ::= <SENDER_TEMPLATE> <SENDER_TSPEC>
                        [ <ADSPEC> ]
                        [ <RECORD_ROUTE> ]
                        [ <SUGGESTED_LABEL> ]
                        [ <RECOVERY_LABEL> ]
                        <LABEL>
```

With the LABEL object defined in [section 4.1 of \[RFC3209\]](#)

Please note that the receiver-driven PATH messages convey the LABEL_REQUEST as an optional object. When the receiver-driven P2MP LSP is uni-directional, the LABEL_REQUEST in the PATH message is not used. When bi-directional receiver-driven P2MP LSP is needed, the LABEL_REQUEST will operate as described in [\[RFC4875\]](#), providing the label allocation operation in the other direction.

With the receiver-driven usage, the Extended Tunnel ID of the P2MP Session object MUST NOT be set to the router ID of the Path-Initiator. This is different from [section 19.1.1 of \[RFC4875\]](#).

[2.2.](#) Resv Message Extensions

Receiver-driven P2MP RSVP-TE does not need any change in the basic RESV message illustrated in [section 6.1 of \[RFC4875\]](#), as long as the SESSION object is using one of the C-Types defined by this document.

For receiver-driven P2MP RSVP-TE, the PATH message is carrying the LABEL object, it is not necessary to have the LABEL object be carried by the RESV message anymore. Except when bi-directional P2MP RSVP-TE is needed, as indicated by the new C-Type of the SESSION object. Within the context of bi-directional P2MP tree structures, one of the directions is established as per [\[RFC3209\]](#). Thus, this document is changing the use of the LABEL object in the FF Flow Descriptor and SE Filter Spec from mandatory to optional. As indicated here:

```
<FF flow descriptor> ::= [ <FLOWSPEC> ] <FILTER_SPEC> [ <LABEL> ]  
                        [ <RECORD_ROUTE> ]  
                        [ <S2L sub-LSP flow descriptor list> ]  
  
<SE filter spec> ::=   <FILTER_SPEC> [ <LABEL> ] [ <RECORD_ROUTE> ]  
                        [ <S2L sub-LSP flow descriptor list> ]
```

2.3. PathErr Message Extensions

The receiver-driven PathErr messages have the same syntax and utilization as the PathErr message described in [\[RFC4875\]](#), with the difference in the SENDER DESCRIPTOR object carried by the PathErr message. The receiver-driven PathErr message will use the SENDER DESCRIPTOR object defined in [Section 2.1](#) of this document, the same SENDER DESCRIPTOR object carried by the Path message the PathErr message corresponds to, allowing the indication of the LABEL object in the SENDER DESCRIPTOR. With the ERROR_SPEC object being able to indicate Label errors with Error Code 24 for Routing Problem and Error Value sub-code of 9 for MPLS label allocation failure as defined in [section 7.3 of \[RFC3209\]](#).

2.4. ResvErr Message Extensions

The receiver-driven ResvErr messages have the same syntax and utilization as the ResvErr message described in [\[RFC4875\]](#). But this ResvErr message will follow the functionality defined for the receiver-driven Resv message of [Section 2.2](#) of this document. Where the FF FLOW DESCRIPTOR object and the SE FILTER SPEC object can optionally contain the LABEL object (instead of mandating the use of the LABEL object). The optional use of the LABEL object is conditioned by the nature of the P2MP tree structure, either uni-directional or bi-directional.

2.5. PathTear Message Extensions

The receiver-driven PathTear message have the same syntax and utilization as the PathTear message described in [[RFC4875](#)]. With the difference in the SENDER DESCRIPTOR object carried by the PathTear message. The receiver-driven PathTear message will use the SENDER DESCRIPTOR object defined in [Section 2.1](#) of this document, the same SENDER DESCRIPTOR object carried by the Path message the PathTear message corresponds to, allowing the indication of the LABEL object in the SENDER DESCRIPTOR.

3. Broadcast Interfaces

The receiver-driven approach interoperates with the RSVP upstream label allocation mechanism [[I-D.ietf-mpls-rsvp-upstream](#)]. Path messages originated by Path-Senders can be signaled over lower layer P2MP-LSP, using context-specific labels. Path-Senders SHOULD detect the presence of such P2MP-LSP over broadcast interfaces for each Path-Receiver. In the absence of upstream allocated P2MP-LSP, path-senders use normal procedures to establish LSP, please note that in such case user data will be replicated by upstream routers.

4. Fast Re-Route Considerations

Fast Reroute techniques are applicable in the context of receiver-driven P2MP tree structures, as stated in [[RFC4875](#)]. However, there are semantic differences with the conventions used in [[RFC4090](#)] and [[RFC4875](#)]. In a receiver-driven paradigm PLR and MP notions are inverted. For example, from a signaling point of view, PLR is the LSR that initiates a Detour S2L sub-LSP towards a MP, and an MP merges paths from primary S2L sub-LSP and detour S2L sub-LSP. However, protection switch takes place at the MP, and data merging takes place in the PLR.

Please note that this document refers to PLR and MP from the signaling point of view. With PLR sending the PATH messages toward the MP.

A Detour S2L Sub-LSP is signaled from the PLR using procedures described in [Section 2](#) of this document, with the DETOUR object conveyed in the RSVP PATH message. A MP not only merges a DETOUR S2Ls of primary S2Ls, but also associates the primary and DETOUR sub-LSPs as protected and protecting sub-LSPs respectively.

S2L Detour protection SHOULD be used in the receiver-driven context, as S2L Detour protection does not require additional extensions.

However, receiver-driven mechanism can easily be extended to P2P Bypass LSP protection.

5. Re-Merging and Crossover

[RFC4875] provides two ways to handle remerge and cross-over. However, within the context of receiver-driven P2MP, a LSR MUST allow remerge and cross-over of path messages of same LSP to form a P2MP tree, which is fundamental to the construction of a P2MP tree. Branches are formed when paths of two or more receivers merge to a common upstream path-receiver.

6. Receiver-Driven Bidirectional LSPs

In certain situations it is required to establish congruent bidirectional LSPs. A receiver-driven bidirectional P2MP tree can be established using the procedures provided in [RFC5467], along with the procedures described in [Section 2](#).

In case of IP/MPLS domains, bidirectional LSPs require an additional RSVP object in addition to the extensions mentioned in [Section 2](#).

A receiver-driven congruent bidirectional P2MP LSPs is an optional feature that could be turned on or off through configuration. Bidirectional LSPs require both upstream and downstream S2L sub-LSPs merge. Both upstream and downstream merging is done according to [RFC4875]. Upstream merging takes place when RSVP PATH messages sent by Path-Senders diverge, and downstream merging takes place when RSVP PATH messages from two or more Senders converge.




Figure 2: Receiver-Driven BiDirectional P2MP RSVP-TE LSP

7. Backward Compatibility

A receiver-driven P2MP LSP mechanism uses a unique C-Type. LSRs that do not support receiver-driven P2MP-TE LSP, send Path Error [TBD] back to the Path Initiator. Additionally, the RSVP Capability object would have the "RD-bit" set to indicate support or non-support of receiver-driven P2MP-LSP. IGP extensions to flood the additional node capabilities will be considered in the future.

8. Security Implications

A receiver **MUST** be authenticated before it is allowed to establish P2MP LSP with source, in addition to hop-by-hop security issues identified by in [RFC 3209](#) and [RFC 4206](#). How a receiver is authenticated is outside the scope of this document.

9. IANA Considerations

To be completed.

10. Acknowledgements

To be completed.

11. References

11.1. Normative References

- [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.
- [RFC4420] Farrel, A., Papadimitriou, D., Vasseur, J., and A. Ayyangar, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource Reservation Protocol-Traffic Engineering (RSVP-TE)", [RFC 4420](#), February 2006.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", [RFC 4206](#), October 2005.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), January 2003.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

11.2. Informative References

- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC5467] Berger, L., Takacs, A., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", [RFC 5467](#), March 2009.
- [I-D.ietf-mpls-rsvp-upstream]
Aggarwal, R. and J. Roux, "MPLS Upstream Label Assignment for RSVP-TE", [draft-ietf-mpls-rsvp-upstream-04](#) (work in progress), July 2009.

Authors' Addresses

Raveendra Torvi
Huawei Technologies
125 Nagog Park
Acton, MA 01720
USA

Email: traveendra@huawei.com

Kwok Ho Chan (editor)
Huawei Technologies
125 Nagog Park
Acton, MA 01720
USA

Email: khchan@huawei.com

Christian Jacquenet
France Telecom
3 avenue Francois Chateau
35000 Rennes,
France

Email: christian.jacquenet@orange-ftgroup.com

