

March 2007

MPLS Upstream Label Assignment for RSVP-TE

[draft-ietf-mpls-rsvp-upstream-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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>.

Abstract

This document describes procedures for distributing upstream-assigned labels for Resource Reservation Protocol - Traffic Engineering (RSVP-TE). It also describes how these procedures can be used for avoiding branch LSR traffic replication on a LAN for RSVP-TE point-to-multipoint (P2MP)LSPs.

Table of Contents

1	Specification of requirements	2
2	Introduction	2
3	RSVP-TE Upstream Label Assignment Capability	3
4	Distributing Upstream-Assigned Labels in RSVP-TE	4
4.1	Procedures	4
5	RSVP-TE Tunnel Identifier Exchange	5
6	RSVP-TE Point-to-Multipoint LSPs on a LAN	6
7	Acknowledgements	7
8	References	7
8.1	Normative References	7
8.2	Informative References	8
9	Author Information	8
10	Intellectual Property Statement	8
11	Full Copyright Statement	9

[1. Specification of requirements](#)

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. Introduction](#)

This document describes procedures for distributing upstream-assigned labels [[MPLS-UPSTREAM](#)] for Resource Reservation Protocol with Traffic Engineering (RSVP-TE). These procedures follow the architecture for MPLS Upstream Label Assignment described in [[MPLS-UPSTREAM](#)].

This document describes extensions to RSVP-TE that a LSR can use to advertise to its neighboring LSRs whether the LSR supports upstream label assignment.

This document also describes extensions to RSVP-TE to distribute upstream-assigned labels.

The usage of MPLS upstream label assignment using RSVP-TE for

avoiding branch LSR [RSVP-P2MP] traffic replication on a LAN for RSVP-TE P2MP TE LSPs [[RSVP-TE-P2MP](#)] is also described.

3. RSVP-TE Upstream Label Assignment Capability

According to [[MPLS-UPSTREAM](#)], upstream-assigned label bindings MUST NOT be used unless it is known that a downstream LSR supports them. This implies that there MUST be a mechanism to enable a LSR to advertise to its RSVP-TE neighbor LSR(s) its support of upstream-assigned labels.

[RSVP-RESTART] defines a CAPABILITY object to be carried within Hello messages, and used to indicate the set of capabilities supported by a node. Currently one flag is defined, the R flag indicating the support for RecoveryPath Srefresh. This document defines a new flag, the U flag, to signal a LSR's support of upstream label assignment to its RSVP-TE neighbors.

The format of a Capability object is:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               | Class-Num(TBA) | C-Type  (1) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reserved                               |U|R|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Recovery Path Srefresh Capable (R): 1 bit, defined in [[RSVP-RESTART](#)].

Upstream Label Assignment Capable (U): 1 bit When set this means that the LSR is capable of both distributing upstream-assigned label bindings and receiving upstream-assigned label bindings

Reserved bits MUST be set to zero on transmission and MUST be ignored on receipt.

The usage of RSVP-TE Hello messages for exchanging upstream label assignment capability implies that a LSR MAY exchange RSVP-TE Hellos with a neighbor before sending/receiving any other RSVP-TE messages to/from that neighbor.

4. Distributing Upstream-Assigned Labels in RSVP-TE

An optional RSVP-TE object, the UPSTREAM_ASSIGNED_LABEL object is introduced to signal an upstream-assigned label. The Class-Num for this object comes from the 0bbbbbbb space and is to be determined by IANA.

UPSTREAM_ASSIGNED_LABEL C-Num = TBD

[illegible]

The label can be encoded as in [RFC3209] when the C-Type is 1 or as a Generalized Label [RFC3473] when the C-Type is 2 or 3.

4.1. Procedures

A RSVP-TE LSR that assigns Upstream-Assigned Labels, distributes them to the downstream LSRs by including them in RSVP-TE Path messages.

A RSVP-TE LSR MUST NOT distribute the UPSTREAM_ASSIGNED_LABEL Object to a downstream LSR if the downstream LSR had not previously advertised the CAPABILITY object with the U bit set in its RSVP-TE Hello messages.

If a downstream RSVP-TE LSR receives a Path message that carries an UPSTREAM_ASSIGNED_LABEL Object and the LSR does not support the object C-Num/C-Type it will return an "Unknown Object C-Num/C-Type" error. If the LSR does support the object, but is unable to process the upstream-assigned label as described in [[MPLS-UPSTREAM](#)] it SHOULD send a PathErr with the error code "Routing problem" and the error value "MPLS Upstream Assigned Label Processing Failure". If the LSR successfully processes the Path message and the upstream-assigned label it MUST send a Resv message upstream as per [[RFC3209](#)] but it MUST NOT include the LABEL object with a downstream assigned label in the Resv Message. This is because as described in [[MPLS-UPSTREAM](#)] two LSRs Ru and Rd for a LSP that is bound to FEC F, MUST use either downstream-assigned label distribution or upstream-assigned label distribution, for FEC F, but NOT both, for packets that are to be transmitted on the LSP from Ru to Rd.

5. RSVP-TE Tunnel Identifier Exchange

As described in [[MPLS-UPSTREAM](#)] an upstream LSR Ru MAY transmit a MPLS packet, the top label of which (L) is upstream-assigned, to a downstream LSR Rd, by encapsulating it in an IP or MPLS tunnel. In this case the fact that L is upstream-assigned is determined by Rd by the tunnel on which the packet is received. There must be a mechanism for Ru to inform Rd that a particular tunnel from Ru to Rd will be used by Ru for transmitting MPLS packets with upstream-assigned MPLS labels.

When RSVP-TE is used for upstream label assignment, the IF_ID RSVP_HOP object is used for signaling the Tunnel Identifier. If Ru uses an IP or MPLS tunnel to transmit MPLS packets with upstream assigned labels to Rd, Ru MUST include the IF_ID RSVP_HOP object [[RFC3473](#)] in Path messages along with the UPSTREAM_ASSIGNED_LABEL Object.

Three new TLVs are introduced in the IF_ID RSVP_HOP object [[RFC3471](#)] to support RSVP-TE P2MP LSPs, IP Multicast Tunnels and context labels. The TLV value acts as the tunnel identifier.

1. RSVP-TE P2MP LSP TLV. Type = TBD. Value of the TLV is the RSVP-TE P2MP Session Object and optionally the P2MP Sender Template Object [[RSVP-TE-P2MP](#)]. The TLV value identifies the RSVP-TE P2MP LSP. This mechanism extends RSVP-TE P2P Hierarchy [[LSP-HIER](#)] to RSVP-TE P2MP Hierarchy. It allows Ru to tunnel an "inner" P2MP LSP, the label for which is upstream assigned, over an "outer" P2MP LSP that has leaves <Rd1...Rdn>. The P2MP LSP IF_ID TLV allows Ru to signal to <Rd1...Rdn> the binding of the inner P2MP LSP to the outer P2MP LSP. The control plane signaling between Ru and <Rd1...Rdn> for the inner P2MP LSP uses directed RSVP-TE signaling messages as in [[LSP-HIER](#)].

2. IP Multicast Tunnel TLV. Type = TBD. In this case the TLV value is a <Source Address, Multicast Group Address> tuple. Source Address is the IP address of the root of the tunnel i.e. Ru, and Multicast Group Address is the Multicast Group Address used by the tunnel.

3. MPLS Context Label TLV. Type = TBD. In this case the TLV value is a <Source Address, MPLS Context Label> tuple. The Source Address belongs to Ru and the MPLS Context Label is an upstream assigned label, assigned by Ru. This allows Ru to tunnel an "inner" RSVP-TE P2MP LSP, the label of which is upstream assigned, over an "outer" MPLS LSP, where the outer LSP has the following property:

- + The label pushed by Ru for the outer MPLS LSP is an upstream assigned context label, assigned by Ru. When <Rd1...Rdn> perform a MPLS label lookup on this label a combination of this label and the incoming interface MUST be sufficient for <Rd1...Rdn> to uniquely determine Ru's context specific label space to lookup the next label on the stack in. <Rd1...Rdn> MUST receive the data sent by Ru with the context specific label assigned by Ru being the top label on the label stack.

Currently the usage of the context label TLV is limited only to RSVP-TE P2MP LSPs on a LAN as specified in the next section. The context label TLV MUST NOT be used for any other purposes.

6. RSVP-TE Point-to-Multipoint LSPs on a LAN

This section describes one application of upstream label assignment using RSVP-TE. Further applications are to be described in separate documents.

[RSVP-TE-P2MP] describes how to setup RSVP-TE P2MP LSPs. On a LAN the solution described in [[RSVP-TE-P2MP](#)] relies on "ingress replication". A LSR on a LAN (say I1), that is a branch LSR for a P2MP LSP, (say Ru) sends a separate copy of a packet that it receives on the P2MP LSP to each of the downstream LSRs on the LAN (say <Rd1...Rdn> that are adjacent to it in the P2MP LSP.

In order to increase efficiency of bandwidth utilization, it is desirable for Ru to send a single copy of the packet for the P2MP LSP on the LAN, when there are multiple downstream routers on the LAN that are adjacent in that P2MP LSP. This requires that each of <Rd1...Rdn> must be able to associate the label L, used by Ru to transmit packets for the P2MP LSP on the LAN, with that P2MP LSP. It is possible to achieve this using RSVP-TE upstream-assigned labels with the following procedures. Assume that Ru and <Rd1...Rdn> support upstream label assignment.

Ru sends a Path message for the P2MP LSP to each of <Rd1...Rdn> that is adjacent on the P2MP LSP, with the same UPSTREAM_ASSIGNED_LABEL object. This object carries an upstream assigned label, L. This message also carries a MPLS Context Label TLV, as described in the previous section, with the value of the MPLS label set to a value assigned by Ru on interface I1 as specified in [[MPLS-UPSTREAM](#)]. <Rd1...Rdn> "reserve" the upstream assigned label in the separate Upstream Neighbor Label Space that they maintain for Ru [[MPLS-UPSTREAM](#)].

Ru can then transmit a single packet for the P2MP LSP to <Rd1...Rdn>

with a top label L using procedures defined in [[MPLS-UPSTREAM](#)] and [[MPLS-MCAST-ENCAPS](#)]. The MPLS packet transmitted by Ru contains as the top label the context label assigned by Ru on the LAN interface, I1. The bottom label is the upstream label assigned by Ru to the RSVP-TE P2MP LSP. The top label is looked up in the context of the LAN interface, I1, [[MPLS-UPSTREAM](#)] by a downstream LSR on the LAN. This lookup enables the downstream LSR to determine the context specific label space to lookup the inner label in.

If a subset of <Rd1...Rdn> do not support upstream label assignment these procedures can still be used between Ru and the remaining subset of <Rd1...Rdn>. Ingress replication and downstream label assignment will continue to be used for LSRs that do not support upstream label assignment.

[7. Acknowledgements](#)

Thanks to Yakov Rekhter for his contribution. Thanks to Ina Minei and Thomas Morin for their comments.

[8. References](#)

[8.1. Normative References](#)

[RFC3031] "MPLS Architecture", E. Rosen, A. Viswanathan, R. Callon, [RFC 3031](#).

[MPLS-UPSTREAM] R. Aggarwal, Y. Rekhter, E. Rosen, "MPLS Upstream Label Assignment and Context Specific Label Space", [draft-ietf-mpls-upstream-label-02.txt](#)

[MPLS-MCAST-ENCAPS] T. Eckert, E. Rosen, R. Aggarwal, Y. Rekhter, [draft-ietf-mpls-codepoint-02.txt](#)

[RFC3209] Awduche et. al. "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.

[RFC2119] "Key words for use in RFCs to Indicate Requirement Levels.", Bradner, March 1997

[RFC3473] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.

[RFC3471] Berger, L. Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#) January 2003.

[RSVP-RESTART] A. Satyanarayana et. al., "Extensions to GMPLS RSVP Graceful Restart", [draft-ietf-ccamp-rsvp-restart-ext-02.txt](#)

8.2. Informative References

[MVPN] E. Rosen, R. Aggarwal [Editors], "Multicast in BGP/MPLS VPNs", [draft-ietf-l3vpn-2547bis-mcast-02.txt](#)

[RSVP-TE-P2MP] R. Aggarwal, D. Papadimitriou, S. Yasukawa [Editors], "Extensions to RSVP-TE for Point to Multipoint TE LSPs", [draft-ietf-mpls-rsvp-te-p2mp-07.txt](#)

9. Author Information

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
Email: rahul@juniper.net

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
France
E-mail: jeanlouis.leroux@orange-ftgroup.com

10. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an

attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

11. Full Copyright Statement

Copyright (C) The IETF Trust (2007). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

