

TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 22 April 2026

Z. Ali
Cisco Systems, Inc
V. Beeram
Juniper Networks
P. Schoenmaker
Meta
19 October 2025

In-Place Bandwidth Update for MPLS RSVP-TE LSPs
draft-alibee-teas-rsvp-inplace-lsp-bw-update-02

Abstract

This document describes the procedure for updating the bandwidth of an MPLS RSVP-TE Label Switched Path (LSP) tunnel in-place without employing make-before-break (MBB).

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 <https://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 22 April 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Signaling In-Place LSP Bandwidth Update	3
2.1. Procedure at the Ingress LER	5
2.2. Procedure at the Transit LSR / Egress LER	6
2.3. Backward Compatibility	6
3. Security Considerations	7
4. IANA Considerations	7
5. Acknowledgments	7
6. Contributors	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Authors' Addresses	8

1. Introduction

Sections 2.5 and 4.6.4 of [RFC3209] describe the RSVP signaling procedure for increasing the bandwidth of an MPLS TE tunnel using a make-before-break (MBB) operation. A bandwidth-update-triggered MBB operation for an MPLS RSVP-TE tunnel involves establishing a new LSP with a new LSP_ID and transferring traffic from the old LSP onto it before tearing down the old LSP. Establishing the new LSP involves programming the forwarding plane with new tunnel labels along its path, even in scenarios where both the old and new LSPs traverse the same path. Such MBB events can occur frequently in networks that deploy the 'auto-bandwidth' feature on RSVP-TE tunnels to monitor bandwidth utilization and periodically adjust tunnel bandwidth, causing a considerable amount of signaling and label programming churn in the network.

[RFC3209] does not explicitly discuss the procedure for handling a decrease in the bandwidth of an MPLS RSVP-TE tunnel, allowing an ingress LER implementation to have the option to update the LSP bandwidth "in-place" without employing MBB. The in-place LSP bandwidth update mechanism reduces signaling churn. It eliminates the need to reprogram labels at each transit Label Switch Router (LSR) along the path of the LSP and shift traffic at the ingress Label Edge Router (LER) from one LSP to another. However, the signaling procedure for handling any failures that the RSVP transit node may encounter while processing an in-place LSP bandwidth update request is unspecified. This document clarifies this procedure. It describes how an implementation can leverage the in-place LSP bandwidth update mechanism in both scenarios where the bandwidth of the TE tunnel needs to be decreased or increased.

This document does not cover the application of the in-place LSP bandwidth update procedure to anything other than point-to-point MPLS RSVP-TE tunnels.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Signaling In-Place LSP Bandwidth Update

Figure 1 illustrates an example RSVP signaling sequence for a scenario where the bandwidth of an LSP is successfully updated in-place. In the signaling sequence used for initial setup, the LSP is signaled with a bandwidth of 60 Mbps. This bandwidth value is encoded in the SENDER_TSPEC object of the PATH message sent hop-by-hop from the ingress LER to the egress LER, and upon successful admission control at each hop is encoded in the FLOWSPEC object of the RESV message sent in the reverse direction. In the in-place update signaling sequence, the same LSP instance, with no change to the LSP_ID in the SENDER_TEMPLATE object, is signaled with a bandwidth of 40 Mbps. When the ingress LER receives a RESV message with 40Mbps encoded in the FLOWSPEC object, the in-place LSP bandwidth update signaling sequence is deemed successful.

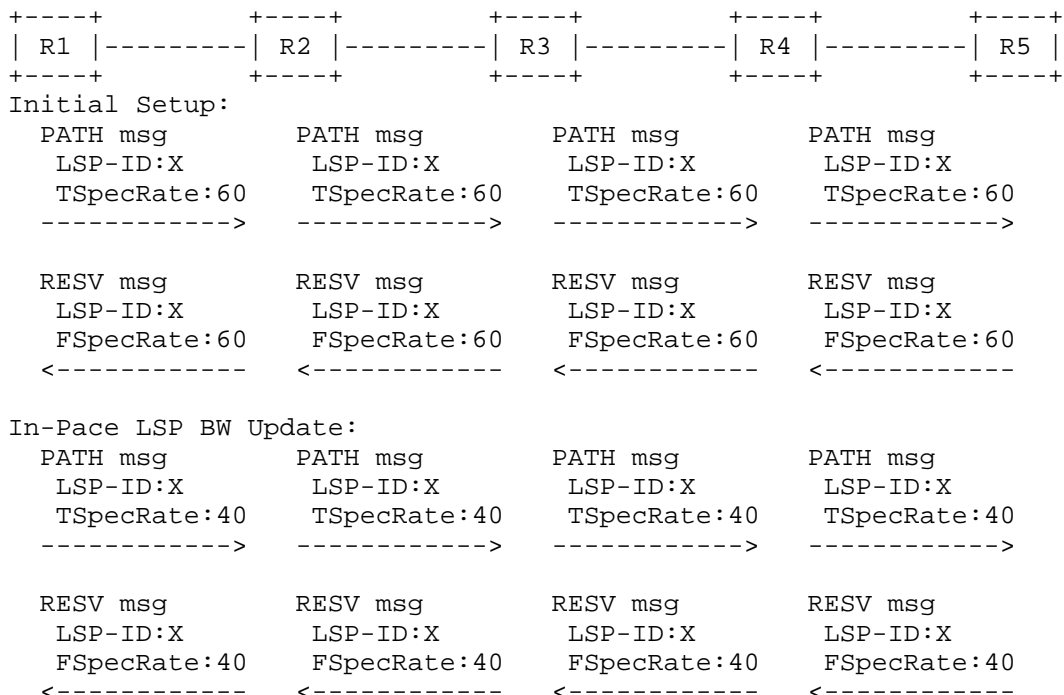


Figure 1: In-Place LSP BW Update - Success

Figure 2 illustrates an example RSVP signaling sequence for a scenario where the bandwidth of an LSP is not successfully updated in-place. In the initial setup signaling sequence, the LSP is signaled with a bandwidth of 60 Mbps. In the in-place update signaling sequence, the same LSP instance, with no change to the LSP_ID in the SENDER_TEMPLATE object, is signaled with a bandwidth of 80 Mbps. However, node R3 fails admission control and sends a PathErr with an error code of 'Admission Control Failure (1)' and error value of 'Requested bandwidth unavailable (2)'. When the ingress LER receives a PathErr message in response to an in-place LSP bandwidth update request, the in-place LSP bandwidth update signaling sequence is deemed a failure. A consequence of this failed attempt is that the bandwidth reservation in the path segment of the LSP upstream of R3 is inconsistent with the path segment downstream of R3. Another scenario in which the in-place LSP update signaling sequence is deemed a failure is when the ingress does not receive either a RESV message or a PATHErr message within a reasonable interval (bandwidth update request timed out). In such failed scenarios, the onus is on the ingress LER to reroute the path using MBB.

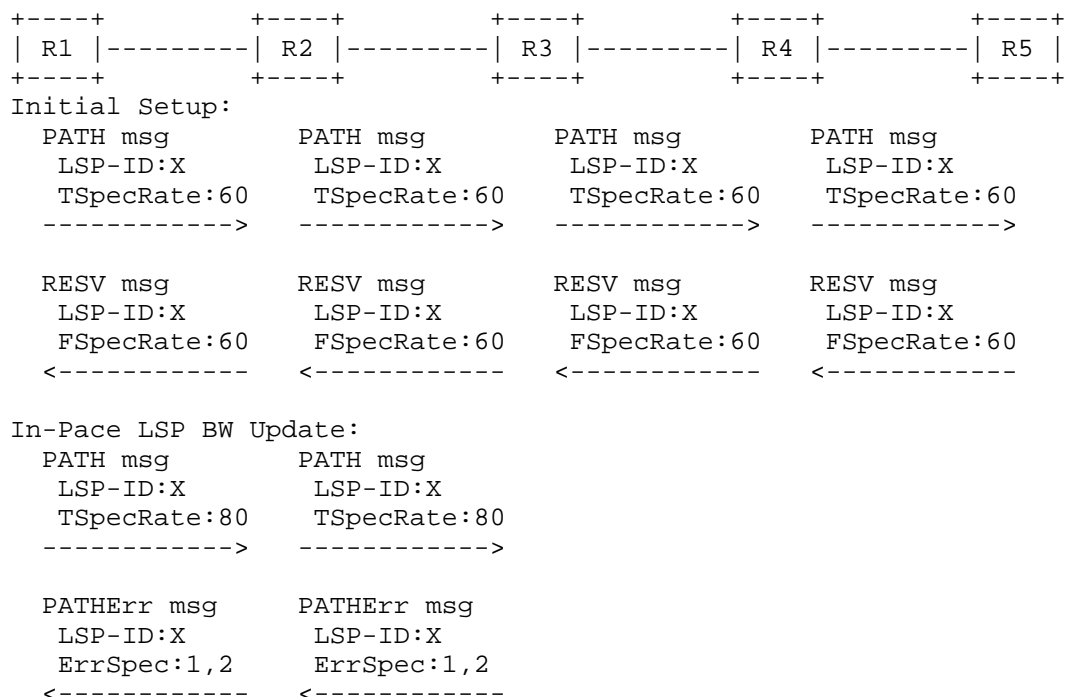


Figure 2: In-Place LSP BW Update - Failure

2.1. Procedure at the Ingress LER

An ingress LER implementation that supports in-place LSP bandwidth update MAY use local policy to determine whether to trigger the "in-place LSP bandwidth update" functionality or the "MBB" procedure to update LSP bandwidth. If the ingress LER is mandated by local policy to use in-place LSP bandwidth update, it SHOULD do the following when there is a need to update the bandwidth of the TE tunnel:

- Compute and check if the current signaled TE path can accommodate the new bandwidth:
 - * If it is determined that the current TE path cannot accommodate the new bandwidth, compute a TE path that accommodates the new bandwidth and initiate MBB signaling procedure.
 - * If it is determined that the current TE path can accommodate the new bandwidth, initiate in-place LSP bandwidth update signaling.

- ~ If in-place LSP bandwidth update signaling succeeds (a RESV message with updated FLOWSPEC is received), consider the bandwidth update procedure to be complete.
- ~ If in-place LSP bandwidth update signaling fails (non-destructive PATHErr message is received or bandwidth update request timed out), compute a TE path that accommodates the desired bandwidth and initiate MBB signaling procedure.
- ~ If a destructive PATHErr message (with Path_State_Removed flag set) or a ResvTear message is received, initiate break-before-make signaling procedure.

2.2. Procedure at the Transit LSR / Egress LER

A transit LSR / egress LER implementation that supports in-place LSP bandwidth update SHOULD perform an admission control procedure when it receives an in-place LSP bandwidth update request. If the admission control succeeds, the transit LSR / egress LER SHOULD allow the in-place LSP bandwidth signaling sequence to complete. If the admission control fails, the transit LSR / egress LER:

- SHOULD generate a non-destructive PATHErr message (without Path_State_Removed flag set) and let the ingress LER take appropriate action.
- SHOULD NOT initiate the teardown of the LSP instance.

2.3. Backward Compatibility

Since the procedure for handling an in-place LSP bandwidth update request at a transit/egress node is not specified in RFC3209, a transit/egress implementation that does not support this functionality may exhibit one of the following behaviors:

- [A] Teardown the signaling state associated with the LSP instance and generate appropriate destructive error and tear messages.
- [B] Retain the signaling state associated with the LSP instance and send an appropriate PathErr to the ingress.
- [C] Silently ignore the in-place LSP bandwidth update request, which will result in the bandwidth update request getting timed out at the ingress.

The in-place LSP bandwidth update functionality SHOULD NOT be enabled at the ingress LERs in a network with non-compliant transit/egress nodes that exhibit behavior [A]. The functionality MAY be enabled at the ingress LERs in a network with non-compliant nodes that exhibit behavior [B] or [C].

3. Security Considerations

This document does not introduce new security issues. The security considerations pertaining to the original RSVP protocol [RFC2205] and RSVP-TE [RFC3209], and those that are described in [RFC5920], remain relevant.

4. IANA Considerations

This document has no IANA actions.

5. Acknowledgments

The authors would like to thank Colby Barth, Stephane Litkowski and Robert Sawaya for their review and suggestions.

6. Contributors

The following people have contributed to this document:

Chandra Ramachandran
Juniper Networks
Email: csekar@juniper.net

Jon Parker
Cisco Systems, Inc.
Email: jdparker@cisco.com

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<https://www.rfc-editor.org/info/rfc2205>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

7.2. Informative References

- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.

Authors' Addresses

Zafar Ali
Cisco Systems, Inc
Email: zali@cisco.com

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

Peter Schoenmaker
Meta
Email: psch@meta.com