

TEAS Working Group
Internet-Draft
Updates: [4090](#), [7551](#)
Intended Status: Standards Track
Expires: May 8, 2019

R. Gandhi, Ed.
Cisco Systems, Inc.
H. Shah
Ciena
J. Whittaker
Verizon
November 4, 2018

**Updates to the Fast Reroute Procedures for
Co-routed Associated Bidirectional Label Switched Paths (LSPs)
draft-ietf-teas-assoc-corouted-bidir-frr-07**

Abstract

Resource Reservation Protocol (RSVP) association signaling can be used to bind two unidirectional Label Switched Paths (LSPs) into an associated bidirectional LSP. When an associated bidirectional LSP is co-routed, the reverse LSP follows the same path as its forward LSP. This document updates the Fast Reroute (FRR) procedures defined in [RFC 4090](#) to support both single-sided and double-sided provisioned associated bidirectional LSPs. This document also updates the procedure for associating two reverse LSPs defined in [RFC 7551](#) to support co-routed bidirectional LSPs. The FRR procedures can ensure that for the co-routed LSPs, traffic flows on co-routed paths in the forward and reverse directions after a failure event.

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."

Copyright Notice

Copyright (c) 2018 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
1.1.	Assumptions and Considerations	3
2.	Conventions Used in This Document	4
2.1.	Key Word Definitions	4
2.2.	Terminology	4
2.2.1.	Forward Unidirectional LSPs	4
2.2.2.	Reverse Co-routed Unidirectional LSPs	5
3.	Problem Statement	5
3.1.	Fast Reroute Bypass Tunnel Assignment	5
3.2.	Node Protection Bypass Tunnels	6
3.3.	Bidirectional LSP Association At Mid-Points	7
4.	Signaling Procedure	8
4.1.	Associated Bidirectional LSP Fast Reroute	8
4.1.1.	Restoring Co-routing with Node Protection Bypass Tunnels	9
4.1.2.	Unidirectional Link Failures	10
4.1.3.	Revertive Behavior after Fast Reroute	10
4.1.4.	Bypass Tunnel Provisioning	10
4.1.5.	One-to-One Bypass Tunnel	11
4.2.	Bidirectional LSP Association At Mid-points	11
5.	Compatibility	11
6.	Security Considerations	12
7.	IANA Considerations	12
Appendix A.	Extended ASSOCIATION ID	12
8.	References	14
8.1.	Normative References	14
8.2.	Informative References	14
	Acknowledgments	16
	Authors' Addresses	16

1. Introduction

The Resource Reservation Protocol (RSVP) (Extended) ASSOCIATION Object is specified in [\[RFC6780\]](#) which can be used generically to associate Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) Label Switched Paths (LSPs). [\[RFC7551\]](#) defines mechanisms for binding two point-to-point unidirectional LSPs [\[RFC3209\]](#) into an associated bidirectional LSP. There are two models described in [\[RFC7551\]](#) for provisioning an associated bidirectional LSP, single-sided and double-sided. In both models, the reverse LSP of the bidirectional LSP may or may not be co-routed and follow the same path as its forward LSP.

In some packet transport networks, there are requirements where the reverse LSP of a bidirectional LSP needs to follow the same path as its forward LSP [\[RFC6373\]](#). The MPLS Transport Profile (TP) [\[RFC6370\]](#) architecture facilitates the co-routed bidirectional LSP by using the GMPLS extensions [\[RFC3473\]](#) to achieve congruent paths. However, the RSVP association signaling allows to enable co-routed bidirectional LSPs without having to deploy GMPLS extensions in the existing networks. The association signaling also allows to take advantage of the existing TE and Fast Reroute (FRR) mechanisms in the network.

[\[RFC4090\]](#) defines FRR extensions for MPLS TE LSPs and those are also applicable to the associated bidirectional LSPs. [\[RFC8271\]](#) defines FRR procedure for GMPLS signaled bidirectional LSPs, such as, coordinate bypass tunnel assignments in the forward and reverse directions of the LSP. The mechanisms defined in [\[RFC8271\]](#) are also useful for the FRR of associated bidirectional LSPs.

This document updates the FRR procedures defined in [\[RFC4090\]](#) to support both single-sided and double-sided provisioned associated bidirectional LSPs. This document also updates the procedure for associating two reverse LSPs defined in [\[RFC7551\]](#) to support co-routed bidirectional LSPs. The FRR procedures can ensure that for the co-routed LSPs, traffic flows on co-routed paths in the forward and reverse directions after fast reroute.

1.1. Assumptions and Considerations

The following assumptions and considerations apply to this document:

- o The FRR procedure for the unidirectional LSPs is defined in [\[RFC4090\]](#) and is not modified by this document.
- o The FRR procedure when using the unidirectional bypass tunnels is defined in [\[RFC4090\]](#) and is not modified by this document.

- o This document assumes that the FRR bypass tunnels used for protected associated bidirectional LSPs are also associated bidirectional.
- o The FRR bypass tunnels used for protected co-routed associated bidirectional LSPs are assumed to be co-routed associated bidirectional.
- o The FRR procedure to coordinate the bypass tunnel assignment defined in this document may be used for protected non-corouted associated bidirectional LSPs but requires that the downstream Point of Local Repair (PLR) and Merge Point (MP) pair of the forward LSP matches the upstream MP and PLR pair of the reverse LSP.
- o Unless otherwise specified in this document, the fast reroute procedures defined in [\[RFC4090\]](#) are used for associated bidirectional LSPs.

2. Conventions Used in This Document

2.1. Key Word Definitions

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.

2.2. Terminology

The reader is assumed to be familiar with the terminology defined in [\[RFC2205\]](#), [\[RFC3209\]](#), [\[RFC4090\]](#), [\[RFC7551\]](#), and [\[RFC8271\]](#).

2.2.1. Forward Unidirectional LSPs

Two reverse unidirectional point-to-point (P2P) LSPs are setup in the opposite directions between a pair of source and destination nodes to form an associated bidirectional Label Switched Path (LSP). In the case of single-sided provisioned LSP, the originating LSP with REVERSE_LSP Object [\[RFC7551\]](#) is identified as a forward unidirectional LSP. In the case of double-sided provisioned LSP, the LSP originating from the higher node address (as source) and terminating on the lower node address (as destination) is identified as a forward unidirectional LSP.

2.2.2. Reverse Co-routed Unidirectional LSPs

Two reverse unidirectional point-to-point (P2P) LSPs are setup in the opposite directions between a pair of source and destination nodes to form an associated bidirectional Label Switched Path (LSP). A reverse unidirectional LSP originates on the same node where the forward unidirectional LSP terminates, and it terminates on the same node where the forward unidirectional LSP originates. A reverse co-routed unidirectional LSP traverses along the same path as the forward direction unidirectional LSP in the opposite direction.

3. Problem Statement

As specified in [[RFC7551](#)], in the single-sided provisioning case, the RSVP TE tunnel is configured only on one endpoint node of the bidirectional LSP. An LSP for this tunnel is initiated by the originating endpoint with (Extended) ASSOCIATION Object containing Association Type set to "single-sided associated bidirectional LSP" and REVERSE_LSP Object inserted in the RSVP Path message. The remote endpoint then creates the corresponding reverse TE tunnel and signals the reverse LSP in response using the information from the REVERSE_LSP Object and other objects present in the received RSVP Path message. As specified in [[RFC7551](#)], in the double-sided provisioning case, the RSVP TE tunnel is configured on both endpoint nodes of the bidirectional LSP. Both forward and reverse LSPs are initiated independently by the two endpoints with (Extended) ASSOCIATION Object containing Association Type set to "double-sided associated bidirectional LSP". With both single-sided and double-sided provisioned bidirectional LSPs, the reverse LSP may or may not be congruent (i.e. co-routed) and follow the same path as its forward LSP.

Both single-sided and double-sided associated bidirectional LSPs require solutions to the following issues for fast reroute to ensure co-routing after a failure event.

3.1. Fast Reroute Bypass Tunnel Assignment

In order to ensure that the traffic flows on a co-routed path after a link or node failure on the protected co-routed LSP path, the mid-point Point of Local Repair (PLR) nodes need to assign matching bidirectional bypass tunnels for fast reroute. Such bypass assignment requires coordination between the forward and reverse direction PLR nodes when more than one bypass tunnels are present on a PLR node.




Figure 1: Multiple Bidirectional Bypass Tunnels

As shown in Figure 1, there are two bypass tunnels available, Bypass tunnel N (on path B-H-I-C) and Bypass tunnel S (on path B-F-G-C). The mid-point PLR nodes B and C need to coordinate bypass tunnel assignment to ensure that traffic in both directions flow through either on the Bypass tunnel N or the Bypass tunnel S, after the link B-C failure.

3.2. Node Protection Bypass Tunnels

When using a node protection bypass tunnel with a protected associated bidirectional LSP, after a link failure, the forward and reverse LSP traffic can flow on different node protection bypass tunnels in the upstream and downstream directions.

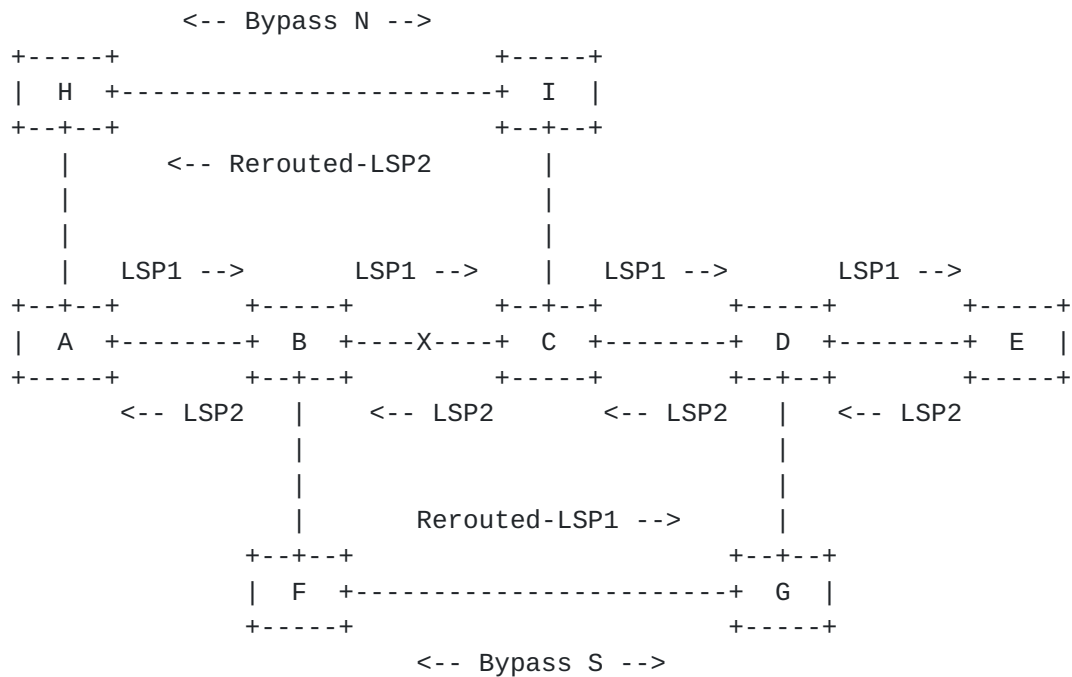


Figure 2: Node Protection Bypass Tunnels

As shown in Figure 2, after the link B-C failure, the downstream PLR node B reroutes the protected forward LSP1 traffic over the bypass tunnel S (on path B-F-G-D) to reach downstream MP node D whereas the upstream PLR node C reroutes the protected reverse LSP2 traffic over the bypass tunnel N (on path C-I-H-A) to reach the upstream MP node A. As a result, the traffic in the forward and reverse directions flows on different bypass tunnels and this can cause the co-routed associated bidirectional LSP to become non-corouted. However, unlike GMPLS LSPs, the asymmetry of paths in the forward and reverse directions does not result in RSVP soft-state timeout with the associated bidirectional LSPs.

3.3. Bidirectional LSP Association At Mid-Points

In packet transport networks, a restoration LSP is signaled after a link failure on the protected LSP path and the protected LSP may or may not be torn down [RFC8131]. In this case, multiple forward and reverse LSPs of a co-routed associated bidirectional LSP may be present at mid-point nodes with identical (Extended) ASSOCIATION Objects. This creates an ambiguity at mid-point nodes to identify the correct associated LSP pair for fast reroute bypass assignment (e.g. during the recovery phase of RSVP graceful restart procedure).

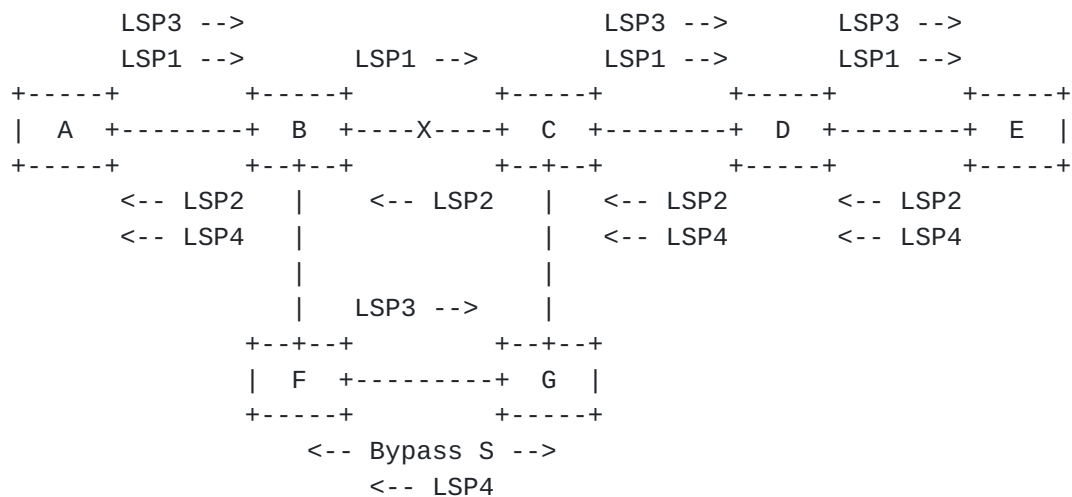


Figure 3: Restoration LSP Set-up after Link Failure

As shown in Figure 3, the protected LSPs LSP1 and LSP2 are an associated LSP pair, similarly the restoration LSPs LSP3 and LSP4 are an associated LSP pair, both pairs belong to the same associated bidirectional LSP and carry identical (Extended) ASSOCIATION Objects.

In this example, the mid-point node D may mistakenly associate LSP1 with the reverse LSP4 instead of the reverse LSP2 due to the matching (Extended) ASSOCIATION Objects. This may cause the co-routed associated bidirectional LSP to become non-corouted after fast reroute. Since the bypass assignment needs to be coordinated between the forward and reverse LSPs, this can also lead to undesired bypass tunnel assignments.

4. Signaling Procedure

4.1. Associated Bidirectional LSP Fast Reroute

For both single-sided and double-sided associated bidirectional LSPs, the fast reroute procedure specified in [RFC4090] is used. In addition, the mechanisms defined in [RFC8271] are used as following.

- o The BYPASS_ASSIGNMENT IPv4 subobject (value: 38) and IPv6 subobject (value: 39) defined in [RFC8271] are used to coordinate bypass tunnel assignment between the forward and reverse direction PLR nodes (see Figure 1). The BYPASS_ASSIGNMENT and Node-ID address [RFC4561] subobjects MUST be added by the downstream PLR node in the RECORD_ROUTE Object (RRO) of the RSVP Path message of the forward LSP to indicate the local bypass tunnel assignment using the procedure defined in [RFC8271]. The upstream node uses the bypass assignment information (namely, bypass tunnel source address, destination address and Tunnel ID) in the received RSVP

Path message of the protected forward LSP to find the associated bypass tunnel in the reverse direction. The upstream PLR node MUST NOT add the BYPASS_ASSIGNMENT subobject in the RRO of the RSVP Path message of the reverse LSP.

- o The downstream PLR node initiates the bypass tunnel assignment for the forward LSP. The upstream PLR (forward direction LSP MP) node reflects the associated bypass tunnel assignment for the reverse direction LSP. The upstream PLR node MUST NOT initiate the bypass tunnel assignment.
- o If the indicated forward bypass tunnel or the associated reverse bypass tunnel is not found, the upstream PLR SHOULD send a Notify message [RFC3473] with Error-code "FRR Bypass Assignment Error" (value: 44) and Sub-code "Bypass Tunnel Not Found" (value: 1) [RFC8271] to the downstream PLR.
- o If the bypass tunnel can not be used as described in [Section 4.5.3 in \[RFC8271\]](#), the upstream PLR SHOULD send a Notify message [RFC3473] with Error-code "FRR Bypass Assignment Error" (value: 44) and Sub-code "Bypass Assignment Cannot Be Used" (value: 0) [RFC8271] to the downstream PLR.
- o After a link or node failure, the PLR nodes in both forward and reverse directions trigger fast reroute independently using the procedures defined in [RFC4090] and send the forward and protected reverse LSP modified RSVP Path messages and traffic over the bypass tunnel. The RSVP Resv signaling of the protected forward and reverse LSPs follows the same procedure as defined in [RFC4090] and is not modified by this document.

4.1.1. Restoring Co-routing with Node Protection Bypass Tunnels

After fast reroute, the downstream MP node assumes the role of upstream PLR and reroutes the reverse LSP RSVP Path messages and traffic over the bypass tunnel on which the forward LSP RSVP Path messages and traffic are received. This procedure is defined as restoring co-routing in [RFC8271]. This procedure is used to ensure that both forward and reverse LSP signaling and traffic flow on the same bidirectional bypass tunnel after fast reroute.

As shown in Figure 2, when using a node protection bypass tunnel with protected co-routed LSPs, asymmetry of paths can occur in the forward and reverse directions after a link failure [RFC8271]. In order to restore co-routing, the downstream MP node D (acting as an upstream PLR) MUST trigger the procedure to restore co-routing and reroute the protected reverse LSP2 RSVP Path messages and traffic over the bypass tunnel S (on path D-G-F-B) to the upstream MP node B upon receiving

the protected forward LSP modified RSVP Path messages and traffic over the bypass tunnel S (on path D-G-F-B) from node B. The upstream PLR node C stops receiving the RSVP Path messages and traffic for the reverse LSP2 from node D (resulting in RSVP soft-state timeout) and it stops sending the RSVP Path messages for the reverse LSP2 over the bypass tunnel N (on path C-I-H-A) to node A.

4.1.2. Unidirectional Link Failures

The unidirectional link failures can cause co-routed associated bidirectional LSPs to become non-co-routed after fast reroute with both link protection and node protection bypass tunnels. However, the unidirectional link failures in the upstream and/or downstream directions do not result in RSVP soft-state timeout with the associated bidirectional LSPs as upstream and downstream PLRs trigger fast reroute independently. The asymmetry of forward and reverse LSP paths due to the unidirectional link failure in the downstream direction can be corrected by using the procedure to restore co-routing specified in [Section 4.1.1](#).

4.1.3. Revertive Behavior after Fast Reroute

When the revertive behavior is desired for a protected LSP after the link is restored, the procedure defined in [\[RFC4090\], Section 6.5.2](#), is followed.

- o The downstream PLR node starts sending the RSVP Path messages and traffic flow of the protected forward LSP over the restored link and stops sending them over the bypass tunnel [\[RFC4090\]](#).
- o The upstream PLR node (when the protected LSP is present) also starts sending the RSVP Path messages and traffic flow of the protected reverse LSPs over the restored link and stops sending them over the bypass tunnel [\[RFC4090\]](#).
- o In case of node protection bypass tunnels (see Figure 2), after restoring co-routing, the upstream PLR node D SHOULD start sending RSVP Path messages and traffic for the reverse LSP over the original link (C-D) when it receives the un-modified RSVP Path messages and traffic for the protected forward LSP over it and stops sending them over the bypass tunnel S (on path D-G-F-B).

4.1.4. Bypass Tunnel Provisioning

Fast reroute bidirectional bypass tunnels can be single-sided or double-sided associated tunnels. For both single-sided and double-sided associated bypass tunnels, the fast reroute assignment policies need to be configured on the downstream PLR nodes of the protected

LSPs that initiate the bypass tunnel assignments. For single-sided associated bypass tunnels, these nodes are the originating endpoints of their signaling.

4.1.5. One-to-One Bypass Tunnel

The fast reroute signaling procedure defined in this document can be used for both facility backup described in [Section 3.2 of \[RFC4090\]](#) and one-to-one backup described in [Section 3.1 of \[RFC4090\]](#). As described in [Section 5.4.2 of \[RFC8271\]](#), in one-to-one backup method, if the associated bidirectional bypass tunnel is already in-use at the upstream PLR, it SHOULD send a Notify message [\[RFC3473\]](#) with Error-code "FRR Bypass Assignment Error" (value: 44) and Sub-code "One-to-One Bypass Already in Use" (value: 2) to the downstream PLR.

4.2. Bidirectional LSP Association At Mid-points

In order to associate the LSPs unambiguously at a mid-point node (see Figure 3), the endpoint node MUST signal Extended ASSOCIATION Object and add unique Extended Association ID for each associated forward and reverse LSP pair forming the bidirectional LSP. An endpoint node MAY set the Extended Association ID to the value formatted according to the structure shown in [Appendix A](#).

- o For single-sided provisioned bidirectional LSPs [\[RFC7551\]](#), the originating endpoint signals the Extended ASSOCIATION Object with a unique Extended Association ID. The remote endpoint copies the contents of the received Extended ASSOCIATION Object including the Extended Association ID in the RSVP Path message of the reverse LSP's Extended ASSOCIATION Object.
- o For double-sided provisioned bidirectional LSPs [\[RFC7551\]](#), both endpoints need to ensure that the bidirectional LSP has a unique Extended ASSOCIATION Object for each forward and reverse LSP pair by selecting appropriate unique Extended Association IDs signaled by them. A controller can be used to provision unique Extended Association ID on both endpoints. The procedure for selecting unique Extended Association ID is outside the scope of this document.

5. Compatibility

This document updates the procedures for fast reroute for associated bidirectional LSPs defined in [\[RFC4090\]](#) and for associating bidirectional LSPs defined in [\[RFC7551\]](#). The procedures use the signaling messages defined in [\[RFC8271\]](#) and no new signaling messages are defined in this document. The procedures ensure that for the co-

LSP-ID

16-bits LSP-ID of the forward LSP [[RFC3209](#)].

Variable Length ID

Variable length ID inserted by the endpoint node of the associated bidirectional LSP [[RFC6780](#)].

An example of IPv6 Extended Association ID format is shown below:

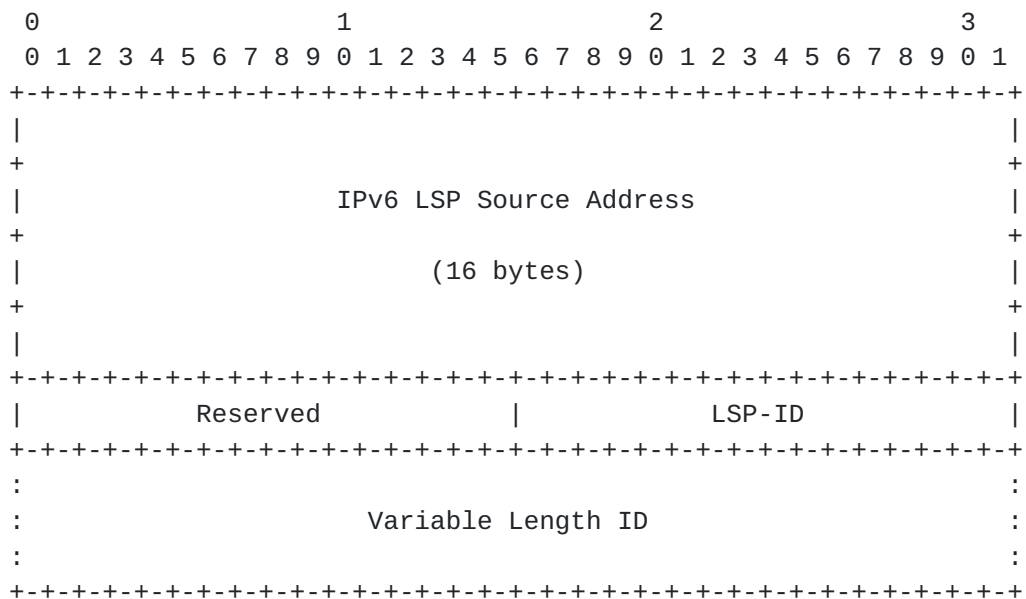


Figure 5: IPv6 Extended Association ID Format Example

LSP Source Address

IPv6 source address of the forward LSP [[RFC3209](#)].

LSP-ID

16-bits LSP-ID of the forward LSP [[RFC3209](#)].

Variable Length ID

Variable length ID inserted by the endpoint node of the associated bidirectional LSP [[RFC6780](#)].

In both IPv4 and IPv6 examples, the Reserved flags MUST be set to 0 on transmission.

8. References

8.1. 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.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.
- [RFC4561] Vasseur, J.P., Ed., Ali, Z., and S. Sivabalan, "Definition of a Record Route Object (RRO) Node-Id Sub-Object", [RFC 4561](#), June 2006.
- [RFC6780] Berger, L., Le Faucheur, F., and A. Narayanan, "RSVP Association Object Extensions", [RFC 6780](#), October 2012.
- [RFC7551] Zhang, F., Ed., Jing, R., and R. Gandhi, Ed., "RSVP-TE Extensions for Associated Bidirectional Label Switched Paths (LSPs)", [RFC 7551](#), DOI 10.17487/RFC7551, May 2015, <<https://www.rfc-editor.org/info/rfc7551>>.
- [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>>.
- [RFC8271] Taillon, M., Saad, T., Ed., Gandhi, R., Ed., Ali, Z., and M. Bhatia, "Updates to Resource Reservation Protocol for Fast Reroute of Traffic Engineering GMPLS Label Switched Paths (LSPs)", [RFC 8271](#), October 2017.

8.2. Informative 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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.

- [RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", [RFC 6370](#), September 2011.
- [RFC6373] Andersson, L., Berger, L., Fang, L., Bitar, N., and E. Gray, "MPLS Transport Profile (MPLS-TP) Control Plane Framework", [RFC 6373](#), September 2011.
- [RFC8131] Zhang, X., Zheng, H., Ed., Gandhi, R., Ed., Ali, Z., and P. Brzozowski, "RSVP-TE Signaling Procedure for End-to-End GMPLS Restoration and Resource Sharing", [RFC 8131](#), March 2017.

Acknowledgments

A special thanks to the authors of [[RFC8271](#)], this document uses the signaling mechanisms defined in that document. The authors would also like to thank Vishnu Pavan Beeram, Daniele Ceccarelli, Deborah Brungard, Adam Roach and Benjamin Kaduk for reviewing this document and providing valuable comments.

Authors' Addresses

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Himanshu Shah
Ciena

Email: hshah@ciena.com

Jeremy Whittaker
Verizon

Email: jeremy.whittaker@verizon.com

