

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: April 20, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Ori Gerstel
Cisco Systems
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG
October 21, 2013

Include Routes - Extension to
Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
draft-ali-ccamp-rsvp-te-include-route-05.txt

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."

This Internet-Draft will expire on April 20, 2014.

Copyright Notice

Copyright (c) 2013 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

Internet-Draft [draft-ali-ccamp-rsvp-te-include-route-05.txt](#)

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.

Abstract

There are scenarios that require two or more LSPs or segments of LSPs to follow same route in the network. This document specifies methods to communicate route inclusions along the loose hops during path setup using the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) protocol.

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](#) [[RFC2119](#)].

Table of Contents

Copyright Notice.....	1
1 . Introduction.....	2
2 . RSVP-TE signaling extensions.....	3
2.1 . IPv4 Point-to-Point Path ERO subobject.....	4
2.2 . IPv6 Point-to-Point Path ERO subobject.....	5
2.3 . Processing rules for Path ERO subobjects.....	7
3 . Security Considerations.....	8
4 . IANA Considerations.....	8
4.1 . New ERO subobject types.....	8
4.2 . New RSVP error sub-codes.....	9
5 . Acknowledgments.....	9
6 . References.....	9
6.1 . Normative References.....	9
6.2 . Informative References.....	10

[1](#). Introduction

There are scenarios that require two or more Label Switched Paths (LSPs) to follow same route in the network. E.g., many deployments require member LSPs of a bundle/ aggregated link (or Forwarding Adjacency (FA))) follow the same route. Possible reasons for two or more LSPs to follow the same end-to-end or

partial route include, but are not limited to:

- . Fate sharing: an application may require that two or more LSPs fail together. In the example of bundle link this would

mean that if one component goes down, the entire bundle goes down.

- . Homogeneous Attributes: it is often required that two or more LSPs have the same TE metrics like latency, latency variation, etc. In the example of a bundle/ aggregated link this would meet the requirement that all component links (FAs) of a bundle should have same latency and latency variation. As noted in [OSPF-TE-METRIC] and [ISIS-TE-METRIC], in certain networks, such as financial information networks, network performance (e.g. latency and latency variation) is becoming critical and hence having bundles with component links (FAs) with homogeneous latency and latency variation is important.

The RSVP-TE specification [RFC3209] and GMPLS extensions to RSVP-TE [RFC3473] allow abstract nodes and resources to be explicitly included in a path setup, e.g., using IPv4 prefix ERO subobject [RFC3209], IPv6 prefix ERO subobject [RFC3209] and Unnumbered Interface ID ERO subobject [RFC3477], etc. However, such inclusion may not be possible in the following scenarios:

- . Inclusion of an LSP path which, while known at the source node of the to-be-signaled LSP, where the required detail of the route is incomplete or unavailable, e.g. due to confidentiality of the path attributes. Such cases are likely to arise in Inter-domain and GMPLS overlay networks.

These use-cases require the relevant path-inclusion information to be communicated to the route expanding nodes. This document defines the necessary information, encodings, and procedures.

This document assumes that node expanding the route is normally a hop of the included LSP. However, there is a race condition in which included LSP is yet to be signaled. This draft addresses this race condition, as detailed in [Section 2.2](#).

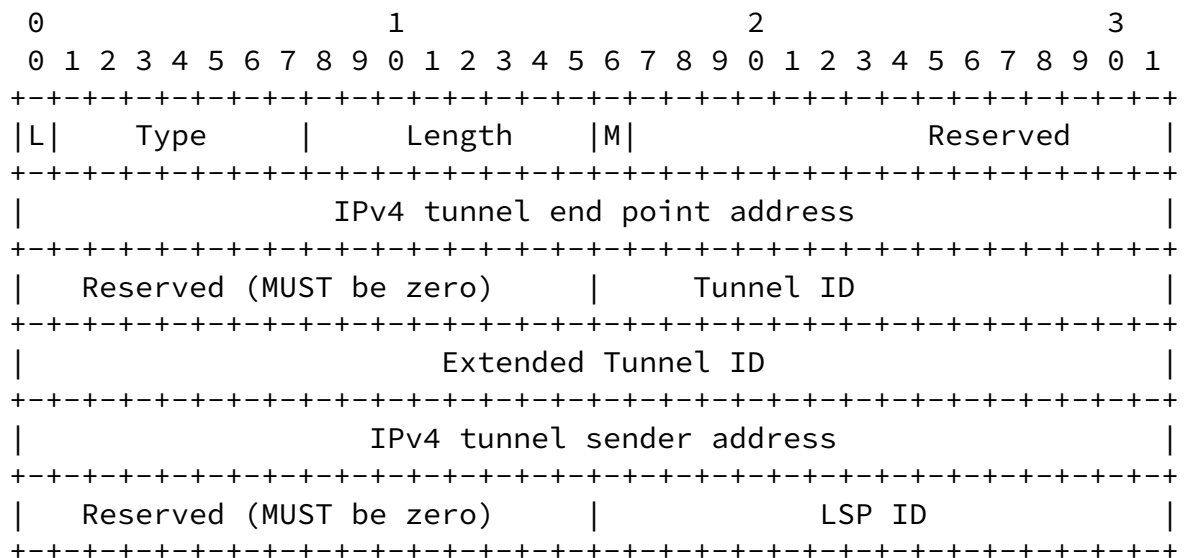
[2](#). RSVP-TE signaling extensions

New IPv4 and IPv6 Point-to-Point (P2P) Path ERO subobject types are defined in this document. These ERO subobjects are used to communicate path inclusion requirements to the ERO expanding node(s). For this purpose, the subobjects carry RSVP-TE Forwarding Equivalence Class (FEC) of the reference LSP who's Path is be used to expand the loose hop of the LSP being signaled.

Internet-Draft [draft-ali-ccamp-rsvp-te-include-route-05.txt](https://datatracker.ietf.org/doc/draft-ali-ccamp-rsvp-te-include-route-05.txt)

2.1. IPv4 Point-to-Point Path ERO subobject

The IPv4 Point-to-Point Path ERO subobject is defined as follows:



L

The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the ERO. If the bit is not set, the subobject represents a strict hop in the explicit route.

This document only defines the use of the subobject in loose hopes in the ERO, i.e., L bit MUST of set to 1.

Type

IPv4 Point-to-Point Path ERO subobject
(to be assigned by IANA; suggested value: 38).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 24.

M bit: When "mandatory inclusion" bit is set, the route of the LSP being signaled MUST follow the path specified by the Path ERO subobject. When mandatory inclusion is not set, the route

of the LSP being signaled SHOULD follow the path specified by the Path ERO subobject.

The remaining fields are used to specify RSVP-TE FEC of the reference LSP who's Path is be used to expand the route of the LSP being signaled. Specifically,

Tunnel ID

Tunnel ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

Extended Tunnel ID

Extended Tunnel ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

IPv4 tunnel sender address

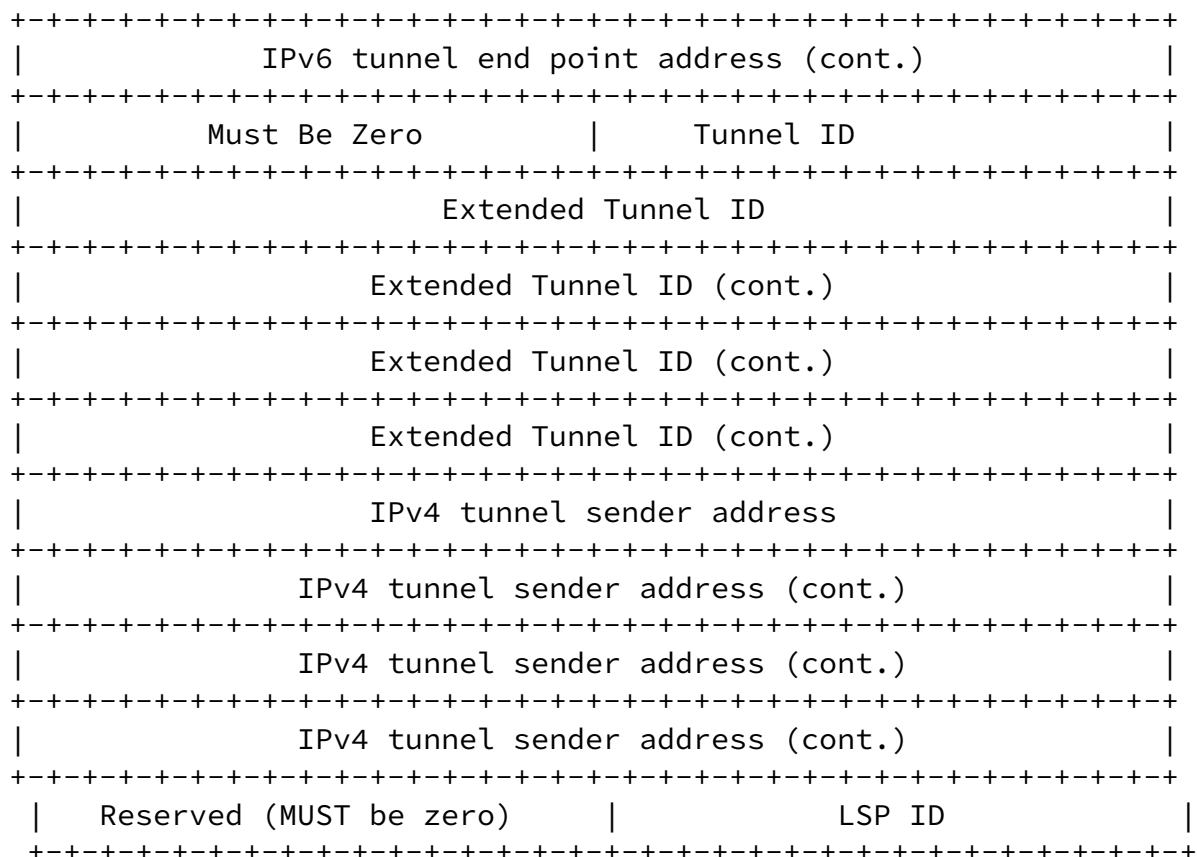
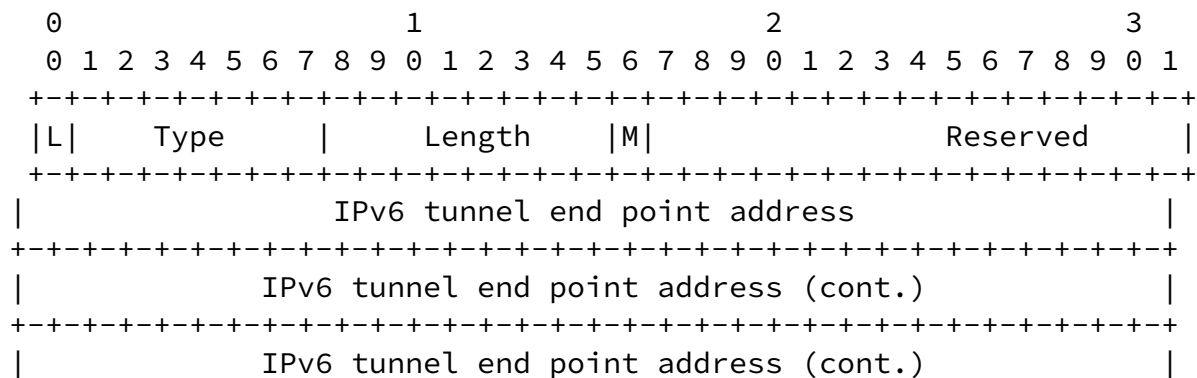
IPv4 tunnel sender address of the reference LSP who's path is be used to expand the route of the LSP being signaled.

LSP ID

LSP ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

2.2. IPv6 Point-to-Point Path ERO subobject

The IPv6 Point-to-Point Path ERO subobject is defined as follows:



The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the ERO. If the bit is not set, the subobject represents a strict hop in the explicit route.

This document only defines the use of the subobject in loose hops in the ERO, i.e., L bit MUST of set to 1.

Type

IPv6 Point-to-Point Path ERO subobject
(to be assigned by IANA; suggested value: 39).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 48.

M bit: The M bit usage is as defined for the M bit of IPv4 Point-to-Point Path ERO subobject.

The remaining fields are used to specific RSVP-TE FEC of the reference LSP who's Path is be used to expand the route of the LSP being signaled, as detailed in [Section 2.1](#).

[2.3](#). Processing rules for Path ERO subobjects

The basic processing rules of an ERO are not altered. Please refer to [\[RFC3209\]](#) for details.

If an LSR strips all local subobjects from an ERO carried in a Path message (according to the procedures in [\[RFC3209\]](#)) and finds that the next subobject is an IPv4 P2P Path ERO subobject or IPv6 P2P LSP subject, it MUST attempt to resolve the Path ERO subobject as described in the following.

If the L bit of the Path ERO subobject is not set, i.e., the subobject represents a strict hop in the explicit route, the processing node MUST respond with a PathErr message with the

error code "Routing Problem" (24) and the error value "Bad initial subobject" (4).

If the M bit is set, the processing node follows the following procedure:

- If the path taken by the LSP referenced in the Path ERO subobject is known to the processing node and the path contains the loose abstract node in the ERO hop, the processing node MUST ensure that loose hop expansion to the next abstract node follows the referenced path.
- If the path taken by the LSP referenced in the Path ERO subobject does not contain the loose abstract node in the ERO hop, the processing node MUST send a PathErr message with the error code "Routing Problem" (24) and the new error value "unknown or inconsistent LSP subobject" (value to be assigned by IANA) for the signaled LSP.
- If the path referenced in the LSP subobject is unknown to the processing node, the processing node SHOULD ignore the Path ERO subobject and SHOULD proceed with the signaling request. After sending the Resv for the signaled LSP, the processing node SHOULD return a PathErr with the error code "Notify

Error" (25) and error sub-code "TBD" (value to be assigned by IANA, suggested value: TBD) for the signaled LSP.

If the M bit is not set, the processing node follows the following procedure:

- If the path taken by the LSP referenced in the Path ERO subobject is known to the processing node and the path contains the loose abstract node in the ERO hop, the processing node SHOULD ensure that loose hop expansion to the next abstract node follows the referenced path.
- If the path taken by the LSP referenced in the Path ERO subobject is unknown to the processing node and/ or the referenced path does not contain the loose abstract node in the ERO hop, the processing node SHOULD ignore the route inclusion specified in the Path ERO subobject and SHOULD

compute a suitable path to the loose abstract node in the ERO hop and proceed with the signaling request. After sending the Resv for the signaled LSP, the processing node SHOULD return a PathErr with the error code "Notify Error" (25) and error sub-code " unknown or inconsistent LSP subobject" (value to be assigned by IANA) for the signaled LSP.

[3. Security Considerations](#)

This document does not introduce any additional security issues above those identified in [[RFC5920](#)], [[RFC2205](#)], [[RFC3209](#)], and [[RFC3473](#)] and [[RFC4874](#)].

[4. IANA Considerations](#)

[4.1. New ERO subobject types](#)

This document adds the following new subobject of the existing entry for ERO (20, EXPLICIT_ROUTE):

Value	Description
-----	-----
TBA subobject	IPv4 Point-to-Point Path ERO
TBA subobject	IPv6 Point-to-Point Path ERO

These subobject may be present in the Explicit Route Object, but not in the Route Record Object.

[4.2. New RSVP error sub-codes](#)

IANA registry: RSVP PARAMETERS
Subsection: Error Codes and Globally-Defined Error Value Sub-Codes

For Error Code "Routing Problem" (24) (see [[RFC3209](#)]) the following sub-codes are defined.

Sub-code	Value
-----	-----

Unknown or inconsistent LSP subject To be assigned by IANA.

For Error Code "Notify Error" (25) (see [[RFC3209](#)]) the following sub-codes are defined.

Sub-code	Value
-----	-----

Unknown or inconsistent LSP subject To be assigned by IANA.

[5.](#) Acknowledgments

Authors would like to thank Gabriele Maria Galimberti, Luyuan Fang and Walid Wakim for their review comments.

[6.](#) References

[6.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.

[6.2.](#) Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", [RFC 4208](#), October 2005.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiimoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", [RFC 6001](#), October 2010.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", [RFC 3477](#), January 2003.
- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", [RFC 2209](#), September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.

Authors' Addresses

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

George Swallow
Cisco Systems, Inc.
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Cisco Systems
Email: mhartley@cisco.com

Ori Gerstel
Cisco Systems
ogerstel@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Ruediger.Kunze@telekom.de