

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: September 04, 2012

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Ori Gerstel
Gabriele Maria Galimberti
Cisco Systems
March 05, 2012

**Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) LSP
Route Diversity using Exclude Routes**

[draft-ali-ccamp-xro-lsp-subobject-00.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 September 03, 2012.

Copyright Notice

Copyright (c) 2012 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

[RFC4874] specifies methods by which route exclusions may be communicated during RSVP-TE signaling in networks where precise explicit paths are not computed by the LSP ingress node. This document specifies signaling for additional route exclusions based on LSPs currently existing or expected to exist within the network.

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

1.	Introduction	3
1.1.	Requirements Notation	4
2.	RSVP-TE signaling extensions	5
2.1.	Terminology	5
2.2.	LSP Subobject	5
2.3.	Processing rules for the LSP subobject	8
2.4.	LSP Subobject in Explicit Exclusion Route Subobject	11
2.4.1.	Processing Rules for the EXRS with LSP subobject	12
4.	IANA Considerations	12
4.1.	New XRO subobject type	12
4.2.	New EXRS subobject type	12
4.3.	New RSVP error sub-code	12
5.	Acknowledgement	13
6.	References	13
	6.1. Normative References	13
	6.2. Informative References	13

1. Introduction

Label-Switched Path (LSP) diversity is required to ensure LSPs may be established without sharing resources, thus greatly reducing the probability of simultaneous connection failures.

LSP diversity is a well-known requirement from Service Providers. When route computation for LSPs that need to be diverse is performed at ingress node, this requirement can be met by a local decision at that node. However, there are scenarios when route computations are performed by remote nodes, in which case there is a need for relevant diversity requirements to be communicated to those nodes. These include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), e.g. inter-domain LSPs.
- . Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI) where route computation may be performed by the UNI-Network (server) node [[RFC4208](#)];

The eXclude Route Object (XRO) and Explicit Exclusion Route Subobject (EXRS) specification [[RFC4894](#)] introduces a means of specifying nodes and resources to be excluded from routes, using the XRO and/ or EXRS.

[RFC4874] facilitates the calculation of diverse routes for LSPs based on known properties of those LSPs including addresses of links and nodes traversed, and Shared Risk Link Groups (SRLGs) of traversed links. This requires that these properties of the LSP(s) from which diversity is required be known to the ingress node which initiates signaling. However, there are circumstances under which this may not be possible or desirable, including (but not limited to):

- . Exclusion of the route of a LSP which does not originate, terminate or traverse the ingress node signaling the diverse LSP, in which case the addresses and SRLGs of the LSP from which diversity is required are unknown to the ingress node.

- . Exclusion of the route of a LSP which, while known at the ingress node of the diverse LSP, has incomplete or unavailable route information, e.g. due to confidentiality of the LSP route attributes. In other words, the scenario in which the reference LSP is hosted by the ingress/ requesting node but the properties required to construct an XRO object are not known to

ingress/ requesting node. Inter-domain and GMPLS overlay networks may present such restrictions.

- . If the route of the reference LSP from which diversity is required (e.g. LSP1) is known to the ingress node, that node can use this information to construct an XRO and send it in the path message during the signaling of a diverse LSP (LSP2). However, if the route of LSP1 changes (e.g. due to re-optimization or failure in the network), the ingress node would need to change path of LSP2 to ensure that it remains diverse from LSP1. It is preferable to have this decision made by the node that calculated the path for LSP2. For example, in the case of GMPLS-UNI, it is better to have such responsibility at the server layer as opposed to at the client layer so that the diversity requirements are transparent to the client layer. Furthermore, in all networking scenarios, if the node performing the route computation/ expansion is aware of the diversity requirements of LSP1 and LSP2, it may consider joint re-optimization of the diverse LSPs.

This document addresses such scenarios and defines procedures that may be used to exclude the route taken by a particular LSP, or the route taken by all LSPs belonging to a single tunnel. Note that this diversity requirement is different from the diversity requirements of path protection where both the reference and diverse LSPs belong to the same tunnel. The diversity requirements considered in this document do not require that the LSPs in question belonging to the same tunnel or share an ingress node.

The means by which the node calculating or expanding the route of the signaled LSP discovers the route of the LSPs from which the signaled LSP requires diversity is beyond the scope of this document. However, in most cases the LSPs with route diversity requirements may transit the node expanding the route.

This document addresses only the exclusion of point-to-point tunnels; point-to-multipoint tunnels will be addressed in a future version. Similarly, at present only IPv4 addresses are considered; support for IPv6 addresses will be added in a future version.

1.1. Requirements Notation

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. RSVP-TE signaling extensions

This section describes the signaling extensions required to address the aforementioned requirements. Specifically, this document defines a new LSP subobject to be signaled in the EXCLUDE_ROUTE object (XRO) and/ or Explicit Exclusion Route Subobject (EXRS) defined in [[RFC4874](#)]. Inclusion of the LSP subobject in any other RSVP object is not defined.

2.1. Terminology

In this document, the following terminology is adapted:

LSP1/TUNNEL1: LSP1/TUNNEL1 is the LSP/tunnel from which diversity is required.

LSP2/TUNNEL2: The term LSP2/TUNNEL2 is used to refer the LSP being signaled with XRO/ EXRS containing the LSP subobject referencing LSP1/TUNNEL1.

CircuitID: The term CircuitID refers to the LSP Forwarding Equivalence Class (FEC) (LSP ID field of the FEC may be ignored depending on the context the CircuitID term is used).

CircuitIDx: The term CircuitIDx refers to CircuitID of LSPx/TUNNELx.

2.2. LSP Subobject

A new IPv4 Point-to-Point (P2P) LSP subobject is defined by this document as follows (IPv6 P2P LSP subobject will be defined in a later version of the document).


```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|   Type   |   Length   |Attribute Flags|Exclusion Flags|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               IPv4 tunnel end point address               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Must Be Zero           |   Tunnel ID   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Extended Tunnel ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               IPv4 tunnel sender address               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Must Be Zero           |   LSP ID   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

L

The L-flag is used as for the other XRO subobjects defined in [\[RFC4874\]](#).

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

Type

A new subobject type is introduced; the subobject type is to be defined by IANA (suggested value: 36).

Length

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

Attribute Flags

The Attribute Flags are used to communicate desirable attributes of the LSP being signaled (In the following, the term LSP2 is used to reference the LSP being signaled; please refer to [Section 2.1](#) for definition of LSP2). The following flags are defined. None, all or multiple attribute flags MAY be set within the same subobject.

0x01 = LSP ID to be ignored

This flag is used to indicate tunnel level exclusion. Specifically, this flag is used to indicate that the lsp-id field of the subobject is to be ignored and the exclusion applies to any LSP matching the rest of the supplied FEC. In other words, if this flag is set, the processing node MUST calculate a route based on exclusions from the routes of all known LSPs matching the tunnel-id, source, destination and extended tunnel-id specified in the subobject.

When this flag is not set, the lsp-id is not ignored and the exclusion applies only to the specified LSP (i.e., LSP level exclusion). In other words, when this flag is not set, route exclusions MUST respect the specified LSP (i.e. the lsp-id, the tunnel-id, source, destination and extended tunnel-id specified needs to be respected during exclusion).

0x02 = Destination node exception

This flag is used to indicate that the destination node may be shared even when sharing of the said node violates the exclusion flags. When this flag is not set, the exclusion flags SHOULD also be respected for the destination node.

0x04 = Processing node exception

This flag is used to indicate that the processing node may be shared even when sharing of the said node violates the exclusion flags. When this flag is not set, the exclusion flags SHOULD also be respected for the processing node.

0x08 = Penultimate node exception

This flag is used to indicate that the penultimate node may be shared even when sharing of the said node violates the exclusion flags. When this flag is not set, the exclusion flags SHOULD also be respected for the penultimate node.

Exclusion-Flags

The Exclusion-Flags are used to communicate desirable types of exclusion. The following flags are defined.

0x01 = SRLG exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be SRLG diverse from the route of the LSP or tunnel specified by the LSP subobject.

0x02 = Node exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be node diverse from the route of the LSP or tunnel specified by the LSP subobject. The node exclusion is subobject to the setting of the "Processing node exception", the "Penultimate node exception" and the "Destination node exception" Attribute Flags.

0x04 = Link exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be link diverse from the route of the LSP or tunnel specified by the LSP subobject.

The remaining fields are as defined in [[RFC3209](#)].

2.3. Processing rules for the LSP subobject

XRO processing as described in [[RFC4874](#)] is unchanged.

If the node is the destination for the LSP being signaled, it SHOULD NOT process a LSP XRO subobject.

If the L-flag is not set, the processing node follows the following procedure:

- The processing node MUST ensure that any route calculated for the signaled LSP (LSP2) respects the requested exclusion flags with respect to the route traversed by the LSP(s) referenced by the LSP subobject (LSP1/TUNNEL1), including local resources.
- If the processing node fails to find a route that meets the requested constraint, the processing node SHOULD return a PathErr with the error code "Routing Problem (24)" and error value "Route blocked by Exclude Route (67)".
- If the route of the LSP or tunnel (LSP1/TUNNEL1) referenced in the LSP subobject is unknown to the processing node, the processing node SHOULD ignore the LSP subobject in the XRO and

SHOULD proceed with the signaling request (for LSP2). However, in this case, after sending Resv for LSP2, the processing node SHOULD return a PathErr with the error code "Notify Error (25)"

and error value "Route to XRO LSP unknown (value: to be assigned by IANA, suggest value: 13)" for LSP2.

- If latter, the route of the LSP or tunnel (LSP1/TUNNEL1) referenced in the LSP subobject becomes known (e.g. when LSP1 is signaled) or the TUNNEL1 is re-optimized to a different route, such that the requested exclusion/ diversity constraints are no longer satisfied and a path that can satisfy the requested constraints exists, the node calculating or expanding the path SHOULD send a PathErr message for LSP2 with the error code "Notify Error (25)" and error value "Preferable path exists (6)". An ingress node receiving this error code/value combination MAY try to reoptimize the LSP2 to the new preferred path.
- Route computation for the LSP or tunnel (LSP1/ TUNNEL1) referenced in the LSP subobject for new setup or for re-optimization LSP SHOULD be performed to avoid situation where the requested exclusion/ diversity constraints are no longer satisfied and a path that can satisfy the requested constraints does not exist. However, if such situation arises the node that computed or expanded the route for LSP2 SHOULD send a PathErr message for LSP2 with the error code "Routing Problem (24)" and error value "Route blocked by Exclude Route (67)".

If the L-flag is set, the processing node follows the following procedure:

- The processing node SHOULD respect the requested exclusion flags with respect to the route traversed by the referenced LSP(s) (LSP1/TUNNEL1) as far as possible.
- If the processing node fails to find a route that meets the requested constraint, it SHOULD proceed with a suitable route that best meets the constraint, but after completion of signaling setup, it SHOULD return a PathErr code "Notify Error (25)" and error value "Failed to respect Exclude Route (value: to be assigned by IANA, suggest value: 14)" to the ingress node.
- If the route of the LSP or tunnel (LSP1/TUNNEL1) referenced in the LSP subobject is unknown to the processing node, the

processing node SHOULD ignore the LSP subobject in XR0 and SHOULD proceed with the signaling request (for LSP2). However, in this case, after sending Resv for LSP2, the processing node SHOULD return a PathErr with the error code "Notify Error" and error value "Route to XR0 LSP unknown" for LSP2.

- If latter, the route of the LSP or tunnel (LSP1/TUNNEL1) referenced in the LSP subobject becomes known (e.g. when LSP1 is signaled) or the TUNNEL1 is re-optimized to a different route, such that the requested exclusion/ diversity constraints are no longer satisfied and a path that can satisfy the requested constraints exists, the node calculating or expanding the path SHOULD send a PathErr message for LSP2 with the error code "Notify Error (25)" and error value "Preferable path exists". An ingress node receiving this error code/value combination MAY try to reoptimize the LSP2 to the new preferred path.
- Route computation for the LSP or tunnel (LSP1/ TUNNEL1) referenced in the LSP subobject for new setup or for re-optimization LSP SHOULD be performed to avoid situation where the requested exclusion/ diversity constraints are no longer satisfied and a path that can satisfy the requested constraints does not exist. However, if such situation arises the node that computed or expanded the route for LSP2 SHOULD send a PathErr message for LSP2 with the error code "Notify Error" and error value "Failed to respect Exclude Route".

The following rules apply equally to L = 0 and L = 1 case:

- XRO object MAY contain multiple LSP subobjects. In this case, the processing node A node receiving a Path message carrying an XRO MAY reject the message if the XRO is too large or complicated for the local implementation or the rules of local policy, as per the roles of XRO defined in [[RFC4874](#)]. In this case, the node MUST send a PathErr message with the error code "Routing Error" and error value "XRO Too Complex". An ingress node receiving this error code/value combination MAY reduce the complexity of the XRO or route around the node that rejected the XRO.
- An ingress node receiving PathErr with the error code "Notify Error" and error values "Route to XRO LSP unknown" or "Failed to respect Exclude Route" MAY take no action other than simply logging these notifications.

Note that LSP1 may be signaled with an XRO LSP subobject referencing CircuitID2 (LSP2 FEC) and LSP2 may be signaled with an XRO LSP subobject referencing CircuitID1 (LSP1 FEC). The

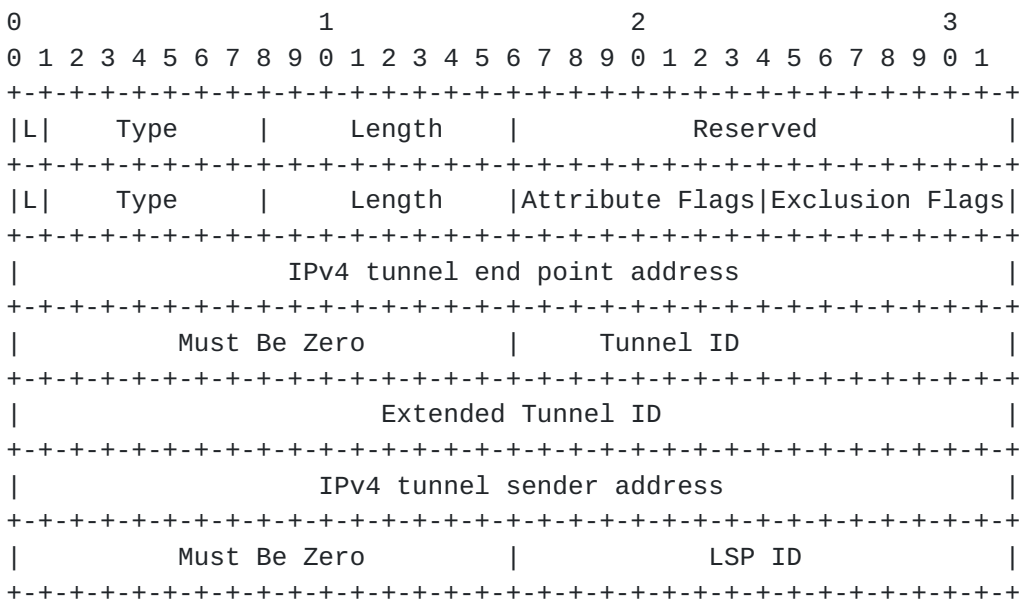
above-mentioned processing rules cover this case. In fact, if "LSP ID to be ignored" attribute flag is set when LSP1 is signaled with an XRO LSP subobject referencing CircuitID2, it is

RECOMMENDED that LSP2 is signaled with an XRO LSP subobject referencing CircuitID1.

2.4. LSP Subobject in Explicit Exclusion Route Subobject (EXRS)

[RFC4874] defines an ERO subobject called Explicit Exclusion Route Subobject (EXRS). An EXRS is used to identify abstract nodes or resources that must not or should not be used on the path between two inclusive abstract nodes or resources in the explicit route. An EXRS contains one or more subobjects of its own, called EXRS subobjects [[RFC4874](#)].

An EXRS MAY include an IPv4 Point-to-Point (P2P) LSP subobject. In this case, EXRS would look as follows:



The meaning of respective fields in EXRS header is as defined in [[RFC4874](#)]. Similarly, the meaning of respective fields in IPv4 P2P LSP subobject is as defined earlier in this document. This is with the exceptions that:

- Processing node exception applies to the node processing the ERO.

- If L bit in the ERO header is not set ($ERO.L = 0$), the IPv4 P2P LSP subobject is processed against the LSPs for which the processing node is ingress, egress or a transit node.

- Penultimate node exception applies to the penultimate node of the loose hop. This flag is only processed if EXRS.L bit is set, i.e., in the loose ERO hop case.
- Destination node exception applies to the abstract node to which the route is expanded. This flag is only processed if EXRS.L bit is set, i.e., in the loose ERO hop case.

2.4.1. Processing Rules for the EXRS with LSP subobject

Processing rules for the EXRS object are same as processing rules as described in [\[RFC4874\]](#). When the EXRS contains one or more LSP subobject(s), processing rule specified in [Section 2.3](#) applies to the node processing the ERO with EXRS subobject.

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 XRO subobject type

This document introduces a new subobject for the EXCLUDE_ROUTE object [\[RFC4874\]](#), C-Type 1.

Subobject Type	Subobject Description
-----	-----
	To be assigned by IANA (suggest value: 36) IPv4 P2P LSP subobject

4.2. New EXRS subobject type

IPv4 P2P LSP subobject is also defined as a new EXRS subobject.

4.3. New RSVP error sub-code

For Error Code = 25 "Notify Error" (see [[RFC3209](#)]) the following sub-code is defined.

Sub-code	Value
-----	-----
Route to XRO LSP unknown	To be assigned by IANA.

Suggested Value: 13.

Failed to respect Exclude Route

To be assigned by IANA.
Suggested Value: 14.

5. Acknowledgement

Authors would like to thank 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.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", [RFC 4874](#), April 2007.

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.
- [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.
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

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

Matt Hartley
Cisco Systems
Email: matt@matthartley.net

Ori Gerstel
Cisco Systems
ogerstel@cisco.com

Gabriele Maria Galimberti
Cisco Systems
ggalimbe@cisco.com