

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: August 2, 2014

Zafar Ali
George Swallow
Clarence Filisfils
Matt Hartley
Gabriele Maria Galimberti
Cisco Systems
Ori Gerstel
SDN Solutions Ltd.
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG
Julien Meuric
France Telecom Orange
February 3, 2014

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Path
Diversity using Exclude Route

draft-ietf-ccamp-lsp-diversity-03.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 August 2, 2014.

Copyright Notice

Copyright (c) 2014 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.

Abstract

RFC 4874 specifies methods by which path exclusions may be communicated during RSVP-TE signaling in networks where precise explicit paths are not computed by the LSP source node. This document specifies signaling for additional route exclusion subobjects based on Paths 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. Introduction2
- 2. RSVP-TE signaling extensions.....4
 - 2.1. Terminology.....4
 - 2.2. Path XRO Subobjects.....5
 - 2.2.1. IPv4 Point-to-Point Path subobject..... 5
 - 2.2.2. IPv6 Point-to-Point Path subobject..... 8
 - 2.3. Processing rules for the Path XRO subobjects.....9
 - 2.4. Path EXRS Subobject.....12
- 3. Security Considerations.....13
- 4. IANA Considerations.....13
 - 4.1. New XRO subobject types.....13
 - 4.2. New EXRS subobject types.....14
 - 4.3. New RSVP error sub-codes.....14
- 5. Acknowledgements.....14
- 6. References.....14
 - 6.1. Normative Reference.....14
 - 6.2. Informative References.....15

1. Introduction

Path diversity is a well-known requirement from Service Providers. Such diversity ensures Label-Switched Paths (LSPs) may be established without sharing resources, thus greatly reducing the probability of simultaneous connection failures.

When a source node has full topological knowledge and is permitted to signal an Explicit Route Object, diverse paths can be computed locally. However, there are scenarios when path computations are performed by remote nodes, thus 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 path computation may be performed by the (server layer) core node [RFC4208].

[RFC4874] introduced a means of specifying nodes and resources to be excluded from a route, using the eXclude Route Object (XRO) and Explicit Exclusion Route Subobject (EXRS).

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

- . Exclusion of a path which does not originate, terminate or traverse the source node signaling the diverse LSP, in which case the addresses and SRLGs of the path from which diversity is required are unknown to the source node.
- . Exclusion of a path which is known to the source node of the diverse LSP, however the node has incomplete or no path information, e.g. due to policy. In other words, the scenario in which the reference path is known by the source / requesting node but the properties required to construct an XRO object are not fully known. Inter-domain and GMPLS overlay networks can present such restrictions.

This document defines procedures that may be used to exclude the path taken by a particular LSP, or the paths taken by all LSPs belonging to a single tunnel. The diversity requirements considered in this document do not require that the paths in question belong to the same tunnel or share the same source or destination node.

If mutually diverse paths are desired for two LSPs belonging to different tunnels, it is recommended that they be signaled with XRO LSP subobjects referencing each other. The processing rules specified in this document cover this case.

The means by which the node calculating or expanding the route of the signaled LSP discovers the route of the path(s) from which the signaled LSP requires diversity are beyond the scope of this document.

This document addresses only the exclusion of point-to-point paths. Exclusion of point-to-multipoint paths is beyond the scope of this document.

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 adopted:

Excluded path: the path from which diversity is required.

Diverse LSP: the LSP being signaled with XRO/ EXRS containing the path subobject referencing the excluded path(s).

Processing node: the node performing a path-calculation involving exclusion specified in an XRO or EXRS.

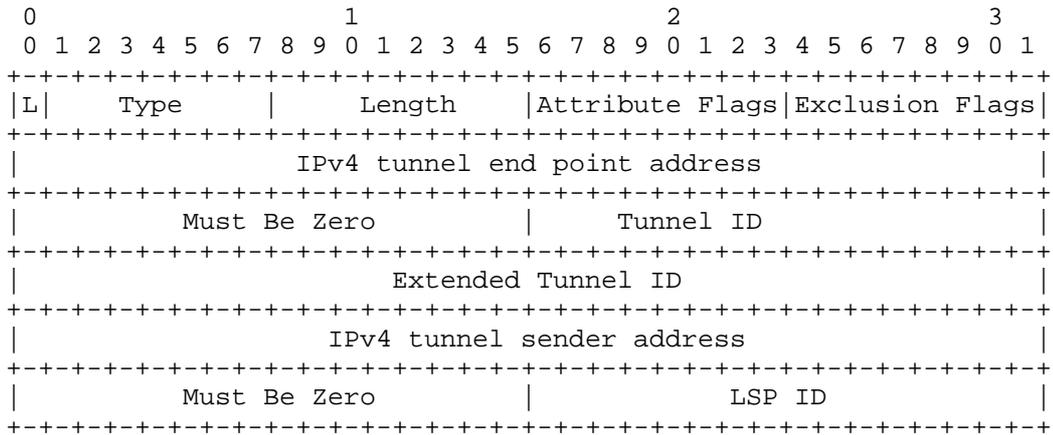
Destination node: in the context of an XRO, this is the destination of the LSP being signaled. In the context of an EXRS, the destination node is the last explicit node to which the loose hop is expanded.

Penultimate node: in the context of an XRO, this is the penultimate hop of the LSP being signaled. In the context of an EXRS, the penultimate node is the penultimate node of the loose hop undergoing expansion.

2.2. Path XRO Subobjects

New IPv4 and IPv6 Point-to-Point (P2P) Path XRO subobjects are defined by this document as follows.

2.2.1. IPv4 Point-to-Point Path subobject



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:

IPv4 Point-to-Point Path subobject (to be assigned 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. The following flags are defined. Each flag acts independently. Any combination of flags is permitted.

0x01 = LSP ID to be ignored

Indicates 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.

0x02 = Destination node exception

Indicates that exclusion does not apply to the destination node of the LSP being signaled.

0x04 = Processing node exception

Indicates that exclusion does not apply to the ERO processing node of the LSP being signaled.

0x08 = Penultimate node exception

Indicates that the penultimate node of the LSP being signaled MAY be shared with the excluded path even when this violates the exclusion flags.

Indicates that exclusion does not apply to the penultimate node of the LSP being signaled.

Exclusion Flags

The Exclusion-Flags are used to communicate the desired type(s) of exclusion. The following flags are defined.

0x01 = SRLG exclusion

Indicates that the path of the LSP being signaled is requested to be SRLG diverse from the excluded path specified by the LSP subobject.

0x02 = Node exclusion

Indicates that the path of the LSP being signaled is requested to be node diverse from the excluded path specified by the LSP subobject.

(Note: the meaning of this flag may be modified by the value of the Attribute-flags.)

0x04 = Link exclusion

Indicates that the path of the LSP being signaled is requested to be link diverse from the path specified by the LSP subobject.

The remaining fields are as defined in [RFC3209].

2.2.2. IPv6 Point-to-Point Path subobject

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	2	3	4	5	6	7	8	9										
L										Type										Length										Attribute Flags										Exclusion Flags									
IPv6 tunnel end point address																																																	
IPv6 tunnel end point address (cont.)																																																	
IPv6 tunnel end point address (cont.)																																																	
IPv6 tunnel end point address (cont.)																																																	
Must Be Zero																				Tunnel ID																													
Extended Tunnel ID																																																	
Extended Tunnel ID (cont.)																																																	
Extended Tunnel ID (cont.)																																																	
Extended Tunnel ID (cont.)																																																	
IPv4 tunnel sender address																																																	
IPv4 tunnel sender address (cont.)																																																	
IPv4 tunnel sender address (cont.)																																																	
IPv4 tunnel sender address (cont.)																																																	
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

IPv6 Point-to-Point Path subobject
(to be assigned by IANA; suggested value: 37).

Length

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

The Attribute Flags and Exclusion Flags are as defined for the IPv4 Point-to-Point LSP XRO subobject.

The remaining fields are as defined in [RFC3209].

2.3. Processing rules for the Path XRO subobjects

XRO processing as described in [RFC4874] is unchanged.

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

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

- The processing node MUST ensure that any path calculated for the signaled LSP respects the requested exclusion flags with respect to the excluded path referenced by the subobject, including local resources.
- If the processing node fails to find a path that meets the requested constraint, the processing node MUST return a PathErr with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67).
- If the excluded path 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. 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 "Route of XRO path unknown" (value to be assigned by IANA, suggested value: 13) for the signaled LSP.

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

- The processing node SHOULD respect the requested exclusion flags with respect to the excluded path to the extent possible.
- If the processing node fails to find a path that meets the requested constraint, it SHOULD proceed with signaling using a suitable path that meets the constraint as far as possible. After sending the Resv for the signaled LSP, it SHOULD return a PathErr message with error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14) to the source node.
- If the excluded path 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. After sending the Resv for signaled LSP, the processing node SHOULD return a PathErr message with the error code "Notify Error" (25) and error sub-code "Route of XRO path unknown" for the signaled LSP.

If, subsequent to the initial signaling of a diverse LSP:

- an excluded path referenced in the diverse LSP's XRO subobject becomes known to the processing node (e.g. when the excluded path is signaled), or
- A change in the excluded path becomes known to the processing node,

the processing node SHOULD re-evaluate the exclusion and diversity constraints requested by the diverse LSP to determine whether they are still satisfied.

- If the requested exclusion constraints for the diverse LSP are no longer satisfied and an alternative path for the diverse LSP that can satisfy those constraints exists, the processing node SHOULD send a PathErr message for the diverse LSP with the error code "Notify Error" (25) and a new error sub-code "compliant path exists" (value: to be assigned by IANA, suggest

value: 15). A source node receiving a PathErr message with this error code and sub-code combination MAY try to reoptimize the diverse tunnel to the new compliant path.

- If the requested exclusion constraints for the diverse LSP are no longer satisfied and no alternative path for the diverse LSP that can satisfy those constraints exists, then:
 - o If the L-flag was not set in the original exclusion, the processing node MUST send a PathErr message for the diverse LSP with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67). The PSR flag SHOULD NOT be set.
 - o If the L-flag was set in the original exclusion, the processing node SHOULD send a PathErr message for the diverse LSP with the error code error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14).

The following rules apply whether or not the L-flag is set:

- An XRO object MAY contain multiple path subobjects.
- A source node receiving a PathErr message with the error code "Notify Error" (25) and error sub-codes "Route of XRO path unknown" or "Failed to respect Exclude Route" MAY take no action.
- The attribute-flags affect the processing of the XRO subobject as follows:
 - o When the "LSP ID to be ignored" flag is set, the processing node MUST calculate a path based on exclusions from the paths of all known LSPs matching the tunnel-id, source, destination and extended tunnel-id specified in the subobject (i.e., tunnel level exclusion). 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).
 - o When the "destination node exception" flag is not set, the exclusion flags SHOULD also be respected for the destination node.

the processing rules specified in Section 2.3 apply to the node processing the ERO with the EXRS subobject.

If a loose-hop expansion results in the creation of another loose-hop in the outgoing ERO, the processing node MAY include the EXRS in the newly-created loose hop for further processing by downstream nodes.

The processing node exception for the EXRS subobject applies to the node processing the ERO.

The destination node exception for the EXRS subobject applies to the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

The penultimate node exception for the EXRS subobject applies to the node before the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

3. Security Considerations

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

4. IANA Considerations

4.1. New XRO subobject types

IANA registry: RSVP PARAMETERS
Subsection: Class Names, Class Numbers, and Class Types

This document introduces two new subobjects for the EXCLUDE_ROUTE object [RFC4874], C-Type 1.

Subobject Type	Subobject Description
To be assigned by IANA (suggested value: 36)	IPv4 P2P Path subobject
To be assigned by IANA (suggested value: 37)	IPv6 P2P Path subobject

4.2. New EXRS subobject types

The IPv4 and IPv6 P2P Path subobjects are also defined as new EXRS subobjects.

4.3. New RSVP error sub-codes

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

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

Sub-code -----	Value -----
Route of XRO path unknown	To be assigned by IANA. Suggested Value: 13.
Failed to respect Exclude Route	To be assigned by IANA. Suggested Value: 14.
Compliant path exists	To be assigned by IANA. Suggested Value: 15.

5. Acknowledgements

The 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.
- [RFC2205] Braden, R. (Ed.), Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReserVation Protocol -- Version 1 Functional Specification", RFC 2205, 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

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

Gabriele Maria Galimberti
Cisco Systems
ggalimbe@cisco.com

Ori Gerstel
SDN Solutions Ltd.
origerstel@gmail.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

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

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

Julien Meuric
France Telecom Orange
Email: julien.meuric@orange.com

George Swallow
Cisco Systems
swallow@cisco.com