

Internet Engineering Task Force (IETF)
Request for Comments: 7898
Category: Experimental
ISSN: 2070-1721

D. Dhody
U. Palle
V. Kondreddy
Huawei Technologies
R. Casellas
CTTC
June 2016

Domain Subobjects
for Resource Reservation Protocol - Traffic Engineering (RSVP-TE)

Abstract

The Resource Reservation Protocol - Traffic Engineering (RSVP-TE) specification and the Generalized Multiprotocol Label Switching (GMPLS) extensions to RSVP-TE allow abstract nodes and resources to be explicitly included in a path setup. Further, Exclude Route extensions to RSVP-TE allow abstract nodes and resources to be explicitly excluded in a path setup.

This document specifies new subobjects to include or exclude Autonomous Systems (ASes), which are identified by a 4-byte AS number, and Interior Gateway Protocol (IGP) areas during path setup.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7898>.

Copyright Notice

Copyright (c) 2016 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. Scope	3
1.2. Requirements Language	4
2. Terminology	4
3. Subobjects for Domains	5
3.1. Domains	5
3.2. Explicit Route Object (ERO) Subobjects	6
3.2.1. Autonomous System	6
3.2.2. IGP Area	7
3.2.3. Mode of Operation	8
3.3. Exclude Route Object (XRO) Subobjects	9
3.3.1. Autonomous System	9
3.3.2. IGP Area	9
3.3.3. Mode of Operation	10
3.4. Explicit Exclusion Route Subobject	10
4. Interaction with Path Computation Element (PCE)	10
5. IANA Considerations	11
5.1. New Subobjects	11
6. Security Considerations	11
7. References	12
7.1. Normative References	12
7.2. Informative References	13
Appendix A. Examples	14
A.1. Inter-Area LSP Path Setup	14
A.2. Inter-AS LSP Path Setup	15
A.2.1. Example 1	15
A.2.2. Example 2	16
Acknowledgments	17
Authors' Addresses	18

1. Introduction

The RSVP-TE specification [RFC3209] and the GMPLS extensions to RSVP-TE [RFC3473] allow abstract nodes and resources to be explicitly included in a path setup using the Explicit Route Object (ERO). Further, Exclude Route extensions [RFC4874] allow abstract nodes or resources to be excluded from the whole path using the Exclude Route Object (XRO). To exclude certain abstract nodes or resources between a specific pair of abstract nodes present in an ERO, an Explicit Exclusion Route subobject (EXRS) is used.

[RFC3209] already describes the notion of abstract nodes, where an abstract node is a group of nodes whose internal topology is opaque to the ingress node of the Label Switched Path (LSP). It further defines a subobject for AS, but with a 2-byte AS number only.

This document extends the notion of abstract nodes by adding new subobjects for IGP areas and 4-byte AS numbers (as per [RFC6793]). These subobjects can be included in ERO, XRO, or EXRS.

In case of per-domain path computation [RFC5152], where the full path of an inter-domain TE LSP cannot be or is not determined at the ingress node, the signaling message could use domain identifiers. The use of these new subobjects is illustrated in Appendix A.

Further, the domain identifier could simply act as a delimiter to specify where the domain boundary starts and ends.

This is a companion document to Path Computation Element Protocol (PCEP) extensions for the domain sequence [RFC7897].

1.1. Scope

The procedures described in this document are experimental. The experiment is intended to enable research for the usage of domain subobjects for inter-domain path setup. For this purpose, this document specifies new domain subobjects as well as how they incorporate with existing subobjects.

The experiment will end two years after the RFC is published. At that point, the RFC authors will attempt to determine how widely this has been implemented and deployed.

This document does not change the procedures for handling subobjects in RSVP-TE.

The new subobjects introduced by this document will not be understood by legacy implementations. If a legacy implementation receives one of the subobjects that it does not understand in an RSVP-TE object, the legacy implementation will behave as described in [RFC3209] and [RFC4874]. Therefore, it is assumed that this experiment will be conducted only when all nodes processing the new subobject form part of the experiment.

When the result of implementation and deployment are available, this document will be updated and refined, and then it will be moved from Experimental to Standards Track.

It should be noted that there are other ways such as the use of a boundary node to identify the domain (instead of a domain identifier); the mechanism defined in this document is just another tool in the toolkit for the operator.

1.2. Requirements Language

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. Terminology

The following terminology is used in this document.

AS: Autonomous System

Domain: As per [RFC4655], any collection of network elements within a common sphere of address management or path computational responsibility. Examples of domains include IGP areas and ASes.

ERO: Explicit Route Object

EXRS: Explicit Exclusion Route subobject

IGP: Interior Gateway Protocol. Either of the two routing protocols: Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS).

IS-IS: Intermediate System to Intermediate System

OSPF: Open Shortest Path First

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Protocol

RSVP: Resource Reservation Protocol

TE LSP: Traffic Engineering Label Switched Path

XRO: Exclude Route Object

3. Subobjects for Domains

3.1. Domains

[RFC4726] and [RFC4655] define domain as a separate administrative or geographic environment within the network. A domain could be further defined as a zone of routing or computational ability. Under these definitions, a domain might be categorized as an AS or an IGP area.

As per [RFC3209], an abstract node is a group of nodes whose internal topology is opaque to the ingress node of the LSP. Using this concept of abstraction, an explicitly routed LSP can be specified as a sequence of IP prefixes or a sequence of ASes. In this document, we extend the notion to include the IGP area and 4-byte AS number.

These subobjects appear in RSVP-TE, notably in:

- o Explicit Route Object (ERO): As per [RFC3209], an explicit route is a particular path in the network topology including abstract nodes (including domains).
- o Exclude Route Object (XRO): As per [RFC4874], an Exclude Route identifies a list of abstract nodes (including domains) that should not be traversed along the path of the LSP being established.
- o Explicit Exclusion Route Subobject (EXRS): As per [RFC4874], used to specify exclusion of certain abstract nodes between a specific pair of nodes. EXRS is a subobject carried inside the ERO. These subobjects can be used to specify the domains to be excluded between two abstract nodes.

3.2. Explicit Route Object (ERO) Subobjects

As stated in [RFC3209], an explicit route is a particular path in the network topology. In addition to the ability to identify specific nodes along the path, an explicit route can identify a group of nodes (abstract nodes) to be traversed along the path.

Some subobjects are defined in [RFC3209], [RFC3473], [RFC3477], [RFC4874], and [RFC5553], but new subobjects related to domains are needed.


This document extends the support for 4-byte AS numbers and IGP areas.

Value	Description
5	4-byte AS number
6	OSPF Area ID
7	IS-IS Area ID

3.2.1. Autonomous System

[RFC3209] already defines 2-byte AS numbers.

To support 4-byte AS numbers as per [RFC6793], the following subobject is defined:



L: The L bit is an attribute of the subobject as defined in [RFC3209], i.e., it's set if the subobject represents a loose hop in the explicit route. If the bit is not set, the subobject represents a strict hop in the explicit route.

Type: 5 (indicating a 4-byte AS number).

Length: 8 (total length of the subobject in bytes).


Reserved: Zero at transmission; ignored at receipt.

AS Number: The 4-byte AS number. Note that if 2-byte AS numbers are in use, the low-order bits (16 through 31) MUST be used, and the high-order bits (0 through 15) MUST be set to zero. For the purpose of this experiment, it is advised to use a 4-byte AS number subobject as the default.

3.2.2. IGP Area

Since the length and format of Area ID is different for OSPF and IS-IS, the following two subobjects are defined:

For OSPF, the Area ID is a 32-bit number. The subobject is encoded as follows:



L: The L bit is an attribute of the subobject as defined in [RFC3209].


Type: 6 (indicating a 4-byte OSPF Area ID).

Length: 8 (total length of the subobject in bytes).

Reserved: Zero at transmission; ignored at receipt.

OSPF Area ID: The 4-byte OSPF Area ID.

For IS-IS, the Area ID is of variable length; thus, the length of the subobject is variable. The Area ID is as described in IS-IS by the ISO standard [ISO10589]. The subobject is encoded as follows:



L: The L bit is an attribute of the subobject as defined in [RFC3209].

Type: 7 (indicating the IS-IS Area ID).

Length: Variable. The length MUST be at least 8 and MUST be a multiple of 4.

Area-Len: Variable (length of the actual (non-padded) IS-IS area identifier in octets; valid values are from 1 to 13, inclusive).

Reserved: Zero at transmission; ignored at receipt.

IS-IS Area ID: The variable-length IS-IS area identifier. Padded with trailing zeroes to a 4-byte boundary.

3.2.3. Mode of Operation

The new subobjects to support 4-byte AS numbers and the IGP (OSPF / IS-IS) area could be used in the ERO to specify an abstract node (a group of nodes whose internal topology is opaque to the ingress node of the LSP).

All the rules of processing (for example, next-hop selection, L bit processing, unrecognized subobjects, etc.) are as per the [RFC3209]. Note that if a node is called upon to process subobjects defined in this document that it does not recognize, it will behave as described in [RFC3209] when an unrecognized ERO subobject is encountered. This means that this node will return a PathErr with error code "Routing Error" and error value "Bad EXPLICIT_ROUTE object" with the EXPLICIT_ROUTE object included, truncated (on the left) to the offending subobject.

3.3. Exclude Route Object (XRO) Subobjects

As stated in [RFC4874], the Exclude Route identifies a list of abstract nodes to exclude (not be traversed) along the path of the LSP being established.

Some subobjects are defined in [RFC3209], [RFC3477], [RFC4874], and [RFC6001], but new subobjects related to domains are needed.

This document extends the support for 4-byte AS numbers and IGP areas.

Value	Description
-----	-----
5	4-byte AS number
6	OSPF Area ID
7	IS-IS Area ID

3.3.1. Autonomous System

[RFC3209] and [RFC4874] already define a 2-byte AS number.

To support 4-byte AS numbers as per [RFC6793], a subobject has the same format as defined in Section 3.2.1 with the following difference:

The meaning of the L bit is as per [RFC4874], where:

0: indicates that the abstract node specified MUST be excluded.

1: indicates that the abstract node specified SHOULD be avoided.

3.3.2. IGP Area

Since the length and format of Area ID is different for OSPF and IS-IS, the following two subobjects are defined:

For OSPF, the Area ID is a 32-bit number. Subobjects for OSPF and IS-IS are of the same format as defined in Section 3.2.2 with the following difference:

The meaning of the L bit is as per [RFC4874].

3.3.3. Mode of Operation

The new subobjects to support 4-byte AS numbers and the IGP (OSPF / IS-IS) area could also be used in the XRO to specify exclusion of an abstract node (a group of nodes whose internal topology is opaque to the ingress node of the LSP).

All the rules of processing are as per [RFC4874].

Note that if a node is called upon to process a subobject defined in this document that it does not recognize, it will behave as described in [RFC4874] when an unrecognized XRO subobject is encountered, i.e., ignore it. In this case, the desired exclusion will not be carried out.

IGP area subobjects in the XRO are local to the current AS. In case of multi-AS path computation that excludes an IGP area in a different AS, an IGP area subobject should be part of EXRS in the ERO to specify the AS in which the IGP area is to be excluded. Further, policy may be applied to prune/ignore area subobjects in XRO at the AS boundary.

3.4. Explicit Exclusion Route Subobject

As per [RFC4874], the Explicit Exclusion Route is used to specify exclusion of certain abstract nodes between a specific pair of nodes or resources in the explicit route. EXRS is an ERO subobject that contains one or more subobjects of its own, called EXRS subobjects.

The EXRS subobject could carry any of the subobjects defined for XRO; thus, the new subobjects to support 4-byte AS numbers and the IGP (OSPF / IS-IS) area can also be used in the EXRS. The meanings of the fields of the new XRO subobjects are unchanged when the subobjects are included in an EXRS, except that the scope of the exclusion is limited to the single hop between the previous and subsequent elements in the ERO.

All the rules of processing are as per [RFC4874].

4. Interaction with Path Computation Element (PCE)

The domain subobjects to be used in PCEP are referred to in [RFC7897]. Note that the new domain subobjects follow the principle that subobjects used in PCEP [RFC5440] are identical to the subobjects used in RSVP-TE and thus are interchangeable between PCEP and RSVP-TE.

5. IANA Considerations

5.1. New Subobjects

IANA maintains the "Resource Reservation Protocol (RSVP) Parameters" registry at <http://www.iana.org/assignments/rsvp-parameters>. Within this registry, IANA maintains two sub-registries:

- o EXPLICIT_ROUTE subobjects (see "Sub-object type - 20 EXPLICIT_ROUTE - Type 1 Explicit Route")
- o EXCLUDE_ROUTE subobjects (see "Sub-object types of Class Types or C-Types - 232 EXCLUDE_ROUTE")

IANA has made identical additions to these registries as follows, in sync with [RFC7897]:

Value	Description	Reference
5	4-byte AS number	[RFC7897], RFC 7898
6	OSPF Area ID	[RFC7897], RFC 7898
7	IS-IS Area ID	[RFC7897], RFC 7898

Further, IANA has added a reference to this document to the new PCEP numbers that are registered by [RFC7897], as shown on <http://www.iana.org/assignments/pcep>.

6. Security Considerations

Security considerations for RSVP-TE and GMPLS signaling RSVP-TE extensions are covered in [RFC3209] and [RFC3473]. This document does not introduce any new messages or any substantive new processing, so those security considerations continue to apply. Further, general considerations for securing RSVP-TE in MPLS-TE and GMPLS networks can be found in [RFC5920]. Section 8 of [RFC5920] describes the inter-provider security considerations, which continue to apply.

The route exclusion security considerations are covered in [RFC4874] and continue to apply.

7. References

7.1. Normative References

- [ISO10589] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, November 2002.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [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, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <<http://www.rfc-editor.org/info/rfc3473>>.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, DOI 10.17487/RFC3477, January 2003, <<http://www.rfc-editor.org/info/rfc3477>>.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, DOI 10.17487/RFC4874, April 2007, <<http://www.rfc-editor.org/info/rfc4874>>.
- [RFC7897] Dhody, D., Palle, U., and R. Casellas, "Domain Subobjects for the Path Computation Element Communication Protocol (PCEP)", RFC 7897, DOI 10.17487/RFC7897, June 2016, <<http://www.rfc-editor.org/info/rfc7897>>.

7.2. Informative References


- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [RFC4726] Farrel, A., Vasseur, J., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, DOI 10.17487/RFC4726, November 2006, <<http://www.rfc-editor.org/info/rfc4726>>.
- [RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, DOI 10.17487/RFC5152, February 2008, <<http://www.rfc-editor.org/info/rfc5152>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.
- [RFC5553] Farrel, A., Ed., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, DOI 10.17487/RFC5553, May 2009, <<http://www.rfc-editor.org/info/rfc5553>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, DOI 10.17487/RFC6001, October 2010, <<http://www.rfc-editor.org/info/rfc6001>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<http://www.rfc-editor.org/info/rfc6793>>.

Appendix A. Examples

These examples are for illustration purposes only to show how the new subobjects could be encoded. They are not meant to be an exhaustive list of all possible use cases and combinations.

A.1. Inter-Area LSP Path Setup

In an inter-area LSP path setup where the ingress and the egress belong to different IGP areas within the same AS, the domain subobjects could be represented using an ordered list of IGP area subobjects in an ERO.



* All IGP areas in one AS (AS 100)

Figure 1: Domain Corresponding to IGP Area

As per Figure 1, the signaling at the ingress could be:

ERO:(A1, ABF1, area B, area C, egress)


It should be noted that there are other ways to achieve the desired signaling; the area subobject provides another tool in the toolkit and can have operational benefits when:

- o Use of PCEP-like domain sequence [RFC7897] configurations in the explicit path is such that area subobjects can be used to signal the loose path.
- o Alignment of subobjects and registries is between PCEP and RSVP-TE, thus allowing easier interworking between path computation and signaling, i.e., subobjects are able to switch between signaling and path computation (if need be).

A.2. Inter-AS LSP Path Setup

A.2.1. Example 1

In an inter-AS LSP path setup where the ingress and the egress belong to a different AS, the domain subobjects (ASes) could be used in an ERO.



* All ASes have one area (area 0)

Figure 2: Domain Corresponding to AS

As per Figure 2, the signaling at the ingress could be:

ERO:(A1, A2, AS B, AS C, egress); or

ERO:(A1, A2, AS B, area 0, AS C, area 0, egress).

Each AS has a single IGP area (area 0); the area subobject is optional.

Note that to get a domain disjoint path, the ingress could also signal the backup path with:

XRO:(AS B)

A.2.2. Example 2

As shown in Figure 3, where AS 200 is made up of multiple areas, the signaling can include both an AS and area subobject to uniquely identify a domain.




Figure 3: Domain Corresponding to AS and Area

As per Figure 3, the signaling at the ingress could be:

ERO:(X1, AS 200, area B, area C, egress).

Acknowledgments

We would like to thank Adrian Farrel, Lou Berger, George Swallow, Chirag Shah, Reeja Paul, Sandeep Boina, and Avantika for their useful comments and suggestions.

Thanks to Vishnu Pavan Beeram for shepherding this document.

Thanks to Deborah Brungard for being the responsible AD.

Thanks to Amanda Baber for the IANA review.

Thanks to Brian Carpenter for the Gen-ART review.

Thanks to Liang Xia (Frank) for the SecDir review.

Thanks to Spencer Dawkins and Barry Leiba for comments during the IESG review.

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: dhruv.ietf@gmail.com

Udayasree Palle
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: udayasree.palle@huawei.com

Venugopal Reddy Kondreddy
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: venugopalreddyk@huawei.com

Ramon Casellas
CTTC
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
Spain

Email: ramon.casellas@cttc.es

