

CCAMP Working Group
Internet-Draft
Intended status: Experimental
Expires: January 10, 2014

D. Dhody
U. Palle
V. Kondreddy
Huawei Technologies
R. Casellas
CTTC
July 9, 2013

Domain Subobjects for Resource ReserVation Protocol - Traffic
Engineering (RSVP-TE)
draft-dhody-ccamp-rsvp-te-domain-subobjects-02

Abstract

The Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) specification [[RFC3209](#)] and the Generalized Multiprotocol Label Switching (GMPLS) extensions to RSVP-TE [[RFC3473](#)] allow abstract nodes and resources to be explicitly included in a path setup. Further Exclude Routes extensions [[RFC4874](#)] allow abstract nodes and resources to be explicitly excluded in a path setup.

This document specifies new subobjects to include or exclude domains during path setup where domain is a collection of network elements within a common sphere of address management or path computational responsibility (such as an Interior Gateway Protocol (IGP) area or an Autonomous System (AS)). Note that the use of AS as an abstract node representing domain is already defined in [[RFC3209](#)] and [[RFC4874](#)], albeit with a 2-Byte AS number.

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 January 10, 2014.

Copyright Notice

Internet-Draft

DOMAIN SUBOBJ

July 2013

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 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.	Requirements Language	3
2.	Terminology	3
3.	Subobjects for Domains	4
3.1.	Domains	4
3.2.	Explicit Route Object (ERO)'s Subobjects	5
3.2.1.	Autonomous system	6
3.2.2.	IGP Area	6
3.2.3.	Mode of Operation	7
3.3.	Exclude Route Object (XRO)'s Subobjects	8
3.3.1.	Autonomous system	8
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)	11
5.	Examples	11
5.1.	Inter-Area LSP Path Setup	11
5.2.	Inter-AS LSP Path Setup	12
5.2.1.	Example 1	12
5.2.2.	Example 2	13
6.	IANA Considerations	15
6.1.	New Subobjects	15
7.	Security Considerations	15
8.	Acknowledgments	15
9.	References	15
9.1.	Normative References	15
9.2.	Informative References	16

Internet-Draft

DOMAIN SUBOBJ

July 2013

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 Routes 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, a 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 [[RFC4893](#)]). These subobjects MAY be included in Explicit Route Object (ERO), Exclude Route object (XRO) or Explicit Exclusion Route Subobject (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, and signaling message may use domain identifiers. The use of these new subobjects is illustrated in [Section 5](#).

1.1. 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 Interior Gateway Protocol (IGP) areas and Autonomous Systems (ASs).

Dhody, et al.

Expires January 10, 2014

[Page 3]

Internet-Draft

DOMAIN SUBOBJ

July 2013

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 may 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 Autonomous Systems. In this document we extend the notion to include IGP area and 4-Byte AS number.

The sub-objects MAY 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 (domains).

- o Exclude Route Object (XRO): As per [RFC4874], an exclude route identifies a list of abstract nodes (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 are a subobject carried inside the ERO. These subobjects are used to specify the domains that must be excluded between two abstract nodes.

3.2. Explicit Route Object (ERO)'s 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) that must be traversed along the path.

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

The following subobject types are used in ERO.

Type	Subobject
1	IPv4 prefix
2	IPv6 prefix
3	Label
4	Unnumbered Interface ID
32	Autonomous system number (2 Byte)
33	Explicit Exclusion (EXRS)
34	SRLG
64	IPv4 Path Key
65	IPv6 Path Key

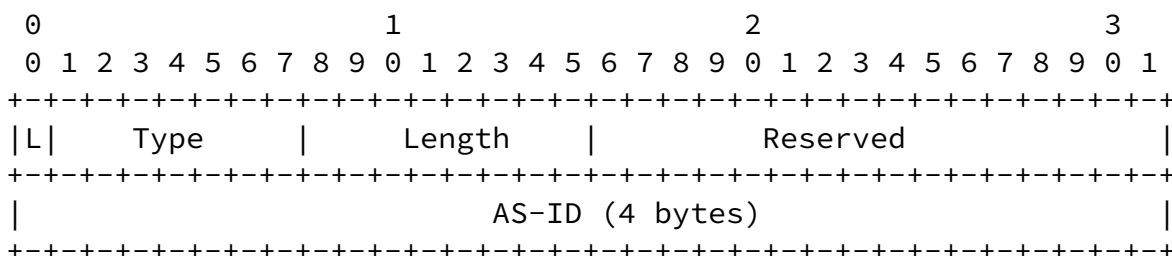
This document extends the above list to support 4-Byte AS numbers and IGP Areas.

Type	Subobject
TBD	Autonomous system number (4 Byte)
TBD	OSPF Area id
TBD	ISIS Area id

[3.2.1.](#) Autonomous system

[RFC3209] already defines 2-Byte AS number.

To support 4-Byte AS numbers as per [[RFC4893](#)], the following subobject is defined:



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

Type: (TBA by IANA) indicating a 4-Byte AS Number.

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

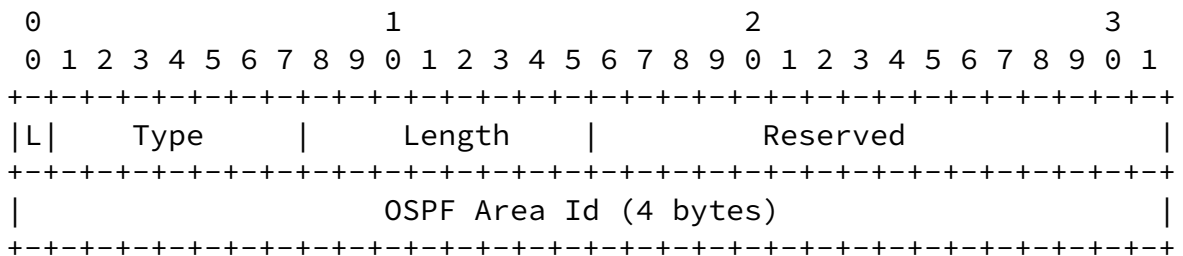
Reserved: Zero at transmission, ignored at receipt.

AS-ID: The 4-Byte AS Number. Note that if 2-Byte AS numbers are in use, the low order bits (16 through 31) should be used and the high order bits (0 through 15) should be set to zero.

[3.2.2.](#) IGP Area

Since the length and format of Area-id is different for OSPF and ISIS, 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 [\[RFC3209\]](#).

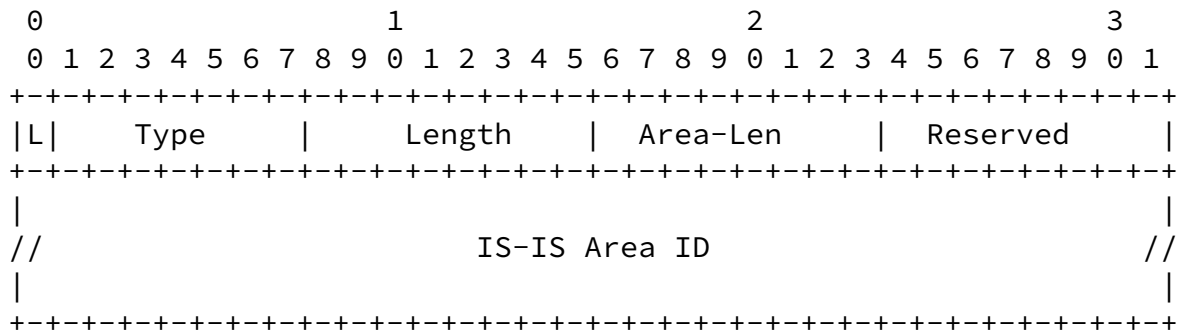
Type: (TBA by IANA) indicating a 4-Byte OSPF Area ID.

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 and thus the length of the subobject is variable. The Area-id is as described in IS-IS by ISO standard [ISO 10589]. The subobject is encoded as follows:



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

Type: (TBA by IANA) indicating IS-IS Area ID.

Length: Variable. As per [\[RFC3209\]](#), the total length of the subobject in bytes, including the L, Type and Length fields. The Length MUST be at least 4, 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 2 to 11 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 four-byte boundary.

[3.2.3.](#) Mode of Operation

The new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY also 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

All the rules of processing (for example Next Hop Selection, L bit processing, unrecognized subobjects etc) are as per the [[RFC3209](#)].

[3.3.](#) Exclude Route Object (XRO)'s Subobjects

As stated in [[RFC4874](#)], the exclude route identifies a list of abstract nodes that should 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.

The following subobject types are used in XRO.

Type	Subobject
1	IPv4 prefix
2	IPv6 prefix
3	Label
4	Unnumbered Interface ID
32	Autonomous system number (2 Byte)
34	SRLG

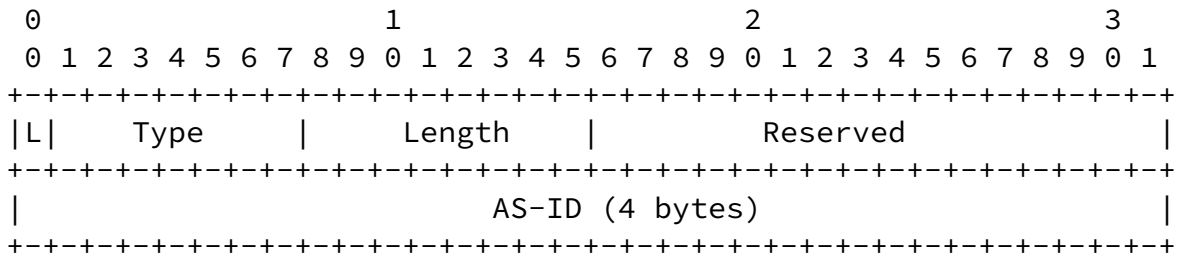
This document extends the above list to support 4-Byte AS numbers and IGP Areas.

Type	Subobject
TBD	Autonomous system number (4 Byte)
TBD	OSPF Area id
TBD	ISIS 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 [[RFC4893](#)], the following subobject is defined:



The meaning of the L bit, similar to [\[RFC4874\]](#), is as follows:

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


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

The meaning of all the other elements (Type, Length, Reserved and 4-Byte AS Id) is same as explained above in [Section 3.2.1](#).

[3.3.2](#). IGP Area

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

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



The meaning of the L bit, similar to [\[RFC4874\]](#), is as follows:

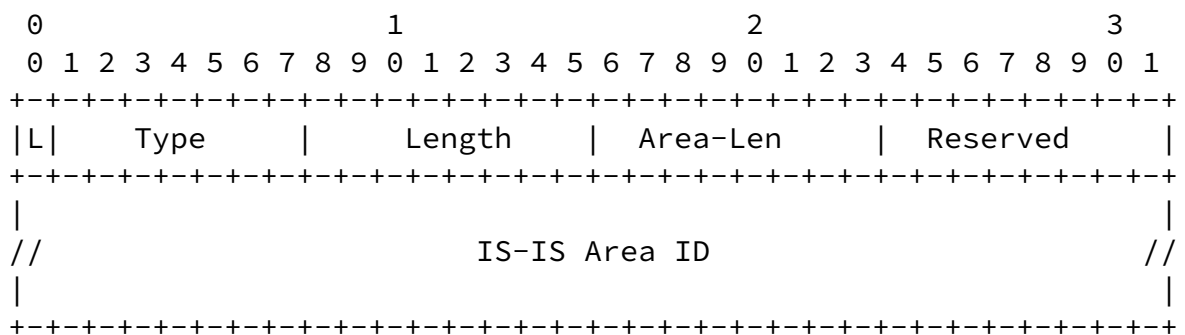
0: indicates that the abstract node (OSPF Area)) specified MUST be excluded.

1: indicates that the abstract node (OSPF Area) specified SHOULD be avoided.

The meaning of all the other elements (Type, Length, Reserved and OSPF Area Id) is same as explained above in [Section 3.2.2](#).

For IS-IS, the area-id is of variable length and thus the length of the subobject is variable. The Area-id is as described in IS-IS by

ISO standard [ISO 10589]. The subobject is encoded as follows:



The meaning of the L bit, similar to [RFC4874], is as follows:

- 0: indicates that the abstract node (IS-IS Area) specified MUST be excluded.
- 1: indicates that the abstract node (IS-IS Area) specified SHOULD be avoided.

The meaning of all the other elements (Type, Length, Area-Len, Reserved and IS-IS Area Id) is same as explained above in [Section 3.2.2](#).

3.3.3. Mode of Operation

The new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY 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 the [RFC4874].

3.4. Explicit Exclusion Route Subobject

As per [RFC4874], the Explicit Exclusion Route defines 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.

EXRS is an ERO subobject that contains one or more subobjects of its own, called EXRS subobjects.

The EXRS subobject may carry any of the subobjects defined for XRO, thus the new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY 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 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 the [\[RFC4874\]](#).

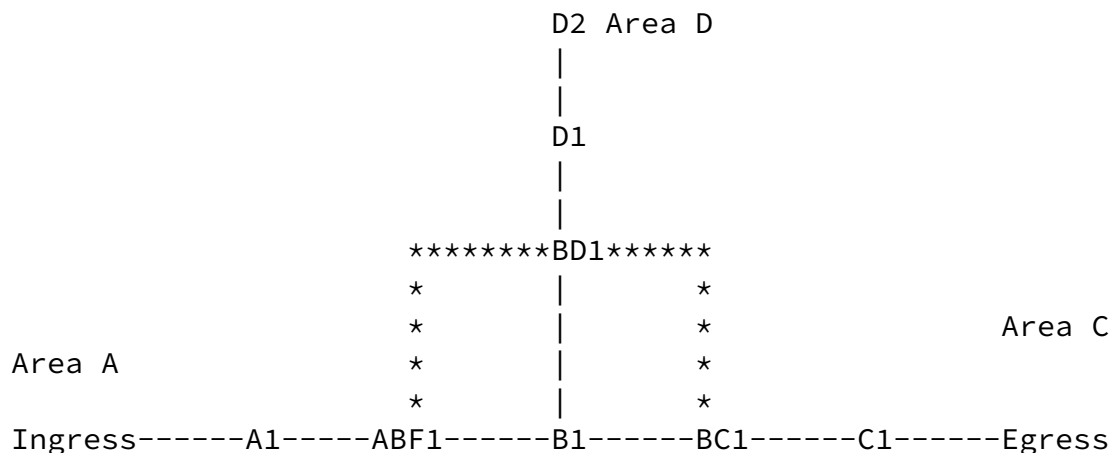
[4.](#) Interaction with Path Computation Element (PCE)


The domain subobjects to be used in Path Computation Element Protocol (PCEP) are referred to in [\[PCE-DOMAIN\]](#). Note that the new domain subobjects follow the principle that subobjects used in PCEP [\[RFC5440\]](#) are identical to the subobjects used in RSVP-TE.

[5.](#) Examples

[5.1.](#) 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 MAY be represented using an ordered list of IGP area subobjects in an ERO. The AS number MAY be skipped, as area information is enough to uniquely identify a domain.





* All IGP Area in one AS (AS 100)

Figure 1: Domain Corresponding to IGP Area

As per Figure 1, the signaling at Ingress MAY be -

ERO:(A1, ABF1, Area B, Area C, Egress); or


ERO:(A1, ABF1, AS 100, Area B, AS 100, Area C, Egress).


The AS subobject is optional and it MAY be skipped. An RSVP-TE implementation should be able to understand both notations and there is no change in the processing rules as mentioned in [\[RFC3209\]](#).

5.2. Inter-AS LSP Path Setup

5.2.1. Example 1

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





* All AS have one area (area 0)

Figure 2: Domain Corresponding to AS

As per Figure 2, the signaling at Ingress MAY 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), Area subobject is optional and it MAY be skipped as AS is enough to uniquely identify a domain. An RSVP-TE implementation should be able to understand both notations and there is no change in the processing rules as mentioned in [\[RFC3209\]](#).

Note that to get a domain disjoint path, the ingress may also signal the backup path with -

XRO:(AS B)

[5.2.2.](#) Example 2

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

[6.1.](#) New Subobjects

IANA registry: RSVP PARAMETERS

Subsection: Class Names, Class Numbers, and Class Types

IANA is requested to add further subobjects to the existing entry for:

20	EXPLICIT_ROUTE
232	EXCLUDE_ROUTE

Subobject	Type	Reference
TBA	4-Byte AS number	[This I.D.]
TBA	OSPF Area ID	[This I.D.]
TBA	IS-IS Area ID	[This I.D.]

[7.](#) Security Considerations

Security considerations for MPLS-TE and GMPLS signaling are covered in [[RFC3209](#)] and [[RFC3473](#)]. This document does not introduce any new messages or any substantive new processing, and so those security considerations continue to apply.

The route exclusion security consideration are covered in [[RFC4874](#)] and continue to apply.

[8.](#) Acknowledgments

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

[9.](#) References

[9.1.](#) Normative References

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

9.2. 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.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource Reservation Protocol - Traffic Engineering (RSVP-TE)", [RFC 3477](#), January 2003.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC4726] Farrel, A., Vasseur, J., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", [RFC 4726](#), November 2006.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource Reservation Protocol-Traffic Engineering (RSVP-TE)", [RFC 4874](#), April 2007.
- [RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", [RFC 4893](#), May 2007.
- [RFC5152] Vasseur, JP., Ayyangar, A., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", [RFC 5152](#), February 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC5553] Farrel, A., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", [RFC 5553](#), May 2009.
- [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](#), October 2010.

Internet-Draft

DOMAIN SUBOBJ

July 2013

[PCE-DOMAIN] Dhody, D., Palle, U., and R. Casellas, "Standard Representation Of Domain Sequence. ([draft-ietf-pce-pcep-domain-sequence](#))", July 2013.

[ISO 10589] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002.

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruv.dhody@huawei.com

Udayasree Palle
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: udayasree.palle@huawei.com

Venugopal Reddy Kondreddy
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: venugopalreddyk@huawei.com

Ramon Casellas

CTTC
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
SPAIN

E-Mail: ramon.casellas@cttc.es

Dhody, et al.

Expires January 10, 2014

[Page 17]