PCE Working Group
Internet-Draft
Intended status: Experimental
Expires: November 1, 2015

D. Dhody
U. Palle
Huawei Technologies
R. Casellas
CTTC

April 30, 2015

Standard Representation of Domain-Sequence
draft-ietf-pce-pcep-domain-sequence-08

Abstract

The ability to compute shortest constrained Traffic Engineering Label Switched Paths (TE LSPs) in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks across multiple domains has been identified as a key requirement. In this context, a 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). This document specifies a standard representation and encoding of a Domain-Sequence, which is defined as an ordered sequence of domains traversed to reach the destination domain to be used by Path Computation Elements (PCEs) to compute inter-domain constrained shortest paths across a predetermined sequence of domains . This document also defines new subobjects to be used to encode domain identifiers.

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 November 1, 2015.

Copyright Notice

Copyright (c) 2015 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

A Path Computation Element (PCE) may be used to compute end-to-end paths across multi-domain environments using a per-domain path computation technique [RFC5152]. The backward recursive path computation (BRPC) mechanism [RFC5441] also defines a PCE-based path computation procedure to compute inter-domain constrained path for (G)MPLS TE LSPs. However, both per-domain and BRPC techniques assume that the sequence of domains to be crossed from source to destination is known, either fixed by the network operator or obtained by other means. Also for inter-domain point-to-multi-point (P2MP) tree computation, [RFC7334] assumes the domain-tree is known in priori.

The list of domains (Domain-Sequence) in point-to-point (P2P) or a domain tree in point-to-multipoint (P2MP) is usually a constraint in inter-domain path computation procedure.

The Domain-Sequence (the set of domains traversed to reach the destination domain) is either administratively predetermined or discovered by some means like H-PCE.

[RFC5440] defines the Include Route Object (IRO) and the Explicit Route Object (ERO). [RFC5521] defines the Exclude Route Object (XRO) and the Explicit Exclusion Route Subobject (EXRS). The use of Autonomous System (AS) (albeit with a 2-Byte AS number) as an abstract node representing a domain is defined in [RFC3209], this document specifies new subobjects to include or exclude domains including IGP area or an Autonomous Systems (4-Byte as per [RFC6793]).

Further, the domain identifier may simply act as delimiter to specify where the domain boundary starts and ends in some cases.

This document further illustrates how the new subobjects defined in this document, along with the existing subobjects, are incorporated in IRO, XRO and ERO to represent a Domain-Sequence.

This is a companion document to Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) extensions for the domain identifiers [DOMAIN-SUBOBJ].

1.1. Scope

The procedures described in this document are experimental. The experiment is intended to enable research for the usage of Domain-Sequence at the PCEs for inter-domain paths. For this purpose this document specify new domain subobjects as well as how they incorporate with existing subobjects to represent a Domain-Sequence.

This document does not change the procedures for handling existing subobjects in PCEP.

The new subobjects introduced by this document will not be understood by a legacy implementation. If one of the subobjects is received in a PCEP object that does not understand it, it will behave as described in Section 3.4.3. Therefore, it is assumed that this experiment will be conducted only when both the PCE and the PCC form part of the experiment. It is possible that a PCC or PCE can operate with peers some of which form part of the experiment and some that do not. In this case, since no capabilities exchange is used to identify which nodes can use these extensions, manual configuration should be used to determine which peerings form part of the experiment.

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

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.

ABR: OSPF Area Border Router. Routers used to connect two IGP areas.

AS: Autonomous System.

ASBR: Autonomous System Boundary Router.

BN: Boundary Node, Can be an ABR or ASBR.

BRPC: Backward Recursive Path Computation

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) area and Autonomous System (AS).

Domain-Sequence: An ordered sequence of domains traversed to reach the destination domain.

ERO: Explicit Route Object

H-PCE: Hierarchical PCE

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

IRO: Include Route Object

IS-IS: Intermediate System to Intermediate System.

OSPF: Open Shortest Path First.

PCC: Path Computation Client: any client application requesting a path computation to be performed by a Path Computation Element.

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.

P2MP: Point-to-Multipoint

P2P: Point-to-Point

RSVP: Resource Reservation Protocol

TE LSP: Traffic Engineering Label Switched Path.

XRO: Exclude Route Object

3. Detail Description

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. Each AS can be made of several IGP areas. In order to encode a Domain-Sequence, it is required to uniquely identify a domain in the Domain-Sequence. A domain can be uniquely identified by area-id or AS number or both.

3.2. Domain-Sequence

A Domain-Sequence is an ordered sequence of domains traversed to reach the destination domain.

A Domain-Sequence can be applied as a constraint and carried in a path computation request to PCE(s). A Domain-Sequence can also be the result of a path computation. For example, in the case of Hierarchical PCE (H-PCE) [RFC6805], Parent PCE could send the Domain-Sequence as a result in a path computation reply.

In a P2P path, the domains listed appear in the order that they are crossed. In a P2MP path, the domain tree is represented as a list of Domain-Sequences.

A Domain-Sequence enables a PCE to select the next domain and the PCE serving that domain to forward the path computation request based on the domain information.

Domain-Sequence can include Boundary Nodes (ABR or ASBR) or Border links (Inter-AS-links) to be traversed as an additional constraint.

Thus a Domain-Sequence can be made up of one or more of -

- o AS Number
- o Area ID
- o Boundary Node ID
- o Inter-AS-Link Address

These are encoded in the new subobjects defined in this document as well as the existing subobjects to represent a Domain-Sequence.

Consequently, a Domain-Sequence can be used:

1. by a PCE in order to discover or select the next PCE in a collaborative path computation, such as in BRPC [RFC5441];

2. by the Parent PCE to return the Domain-Sequence when unknown; this can then be an input to the BRPC procedure [RFC6805];

3. by a Path Computation Client (PCC) or a PCE, to constraint the domains used in inter-domain path computation, explicitly specifying which domains to be expanded or excluded;

4. by a PCE in the per-domain path computation model [RFC5152] to identify the next domain;

3.3. Domain-Sequence Representation

Domain-Sequence appears in PCEP messages, notably in -

- Include Route Object (IRO): As per [RFC5440], IRO can be used to specify a set of network elements to be traversed to reach the destination, which includes subobjects used to specify the Domain-Sequence.

- Exclude Route Object (XRO): As per [RFC5521], XRO can be used to specify certain abstract nodes, to be excluded from whole path, which includes subobjects used to specify the Domain-Sequence.

- Explicit Exclusion Route Subobject (EXRS): As per [RFC5521], EXRS can be used to specify exclusion of certain abstract nodes (including domains) between a specific pair of nodes. EXRS are a subobject inside the IRO.

- Explicit Route Object (ERO): As per [RFC5440], ERO can be used to specify a computed path in the network. For example, in the case of H-PCE [RFC6805], a Parent PCE can send the Domain-Sequence as a result, in a path computation reply using ERO.

3.4. Include Route Object (IRO)

As per [RFC5440], IRO (Include Route Object) can be used to specify that the computed path needs to traverse a set of specified network elements or abstract nodes.

3.4.1. Subobjects

Some subobjects are defined in [RFC3209], [RFC3473], [RFC3477] and [RFC4874], but new subobjects related to Domain-Sequence are needed.

This document extends the support for 4-Byte AS numbers and IGP Areas.

```
             Type   Subobject
              TBD1  Autonomous system number (4 Byte)
              TBD2  OSPF Area id
              TBD3  ISIS Area id
```

3.4.1.1. Autonomous system

[RFC3209] already defines 2 byte AS number.

To support 4 byte AS number as per [RFC6793] following subobject is defined:

```
 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    |          Reserved             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          AS-ID (4 bytes)                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L: The L bit is an attribute of the subobject as defined in [RFC3209] and usage in IRO subobject updated in [IRO-UPDATE].

Type: (TBD1 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) MUST be used and the high order bits (0 through 15) MUST be set to zero.

3.4.1.2. IGP Area

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

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

```
 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    |         Reserved              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    OSPF Area Id (4 bytes)                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L: The L bit is an attribute of the subobject as defined in [RFC3209] and usage in IRO subobject updated in [IRO-UPDATE].

Type: (TBD2 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 [ISO10589]. The subobject is encoded as follows:

```
 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    |  Area-Len     |  Reserved     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//                        IS-IS Area ID                        //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L: The L bit is an attribute of the subobject as defined in [RFC3209] and usage in IRO subobject updated in [IRO-UPDATE].

Type: (TBD3 by IANA) indicating 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 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.

3.4.2. Update in IRO specification

[RFC5440] describes IRO as an optional object used to specify network elements to be traversed by the computed path. It further state that the L bit of such subobject has no meaning within an IRO. It also did not mention if IRO is an ordered or un-ordered list of subobjects.

An update to IRO specification [IRO-UPDATE] makes IRO as an ordered list, as well as support for loose bit (L-bit) is added.

The use of IRO for Domain-Sequence, assumes the updated specification for IRO, as per [IRO-UPDATE].

3.4.3. IRO for Domain-Sequence

Some subobjects for the IRO are defined in [RFC3209], [RFC3477], and [RFC4874]; further some new subobjects related to Domain-Sequence are also added in this document as mentioned in Section 3.4.

The subobject type for IPv4, IPv6, and unnumbered Interface ID can be used to specify Boundary Nodes (ABR/ASBR) and Inter-AS-Links. The subobject type for the AS Number (2 or 4 Byte) and the IGP Area are used to specify the domain identifiers in the Domain-Sequence.

The IRO can incorporate the new domain subobjects with the existing subobjects in a sequence of traversal.

Thus an IRO, comprising of subobjects that represents a Domain-Sequence, define the domains involved in an inter-domain path computation, typically involving two or more collaborative PCEs.

A Domain-Sequence can have varying degrees of granularity. It is possible to have a Domain-Sequence composed of, uniquely, AS identifiers. It is also possible to list the involved IGP areas for a given AS.

In any case, the mapping between domains and responsible PCEs is not defined in this document. It is assumed that a PCE that needs to obtain a "next PCE" from a Domain-Sequence is able to do so (e.g. via administrative configuration, or discovery).

A PCC builds an IRO to encode the Domain-Sequence, so that the cooperating PCEs could compute an inter-domain shortest constrained path across the specified sequence of domains.

For each inclusion, the PCC clears the L-bit to indicate that the PCE is required to include the domain, or sets the L-bit to indicate that the PCC simply desires that the domain be included in the Domain-Sequence.

If a PCE receives an IRO in a Path Computation request (PCReq) message that contains subobjects defined in this document, that it does not recognize, it will respond according to the rules for a malformed object as per [RFC5440]. The PCE MAY also include the IRO in the PCErr to indicate in which case the IRO SHOULD be terminated immediately after the unrecognized subobject.

PCE MUST act according to the requirements expressed in the subobject. That is, if the L-bit is clear, the PCE(s) MUST produce a path that follows the Domain-Sequence in order identified by the subobjects in the path. If the L-bit is set, the PCE(s) SHOULD produce a path along the Domain-Sequence unless it is not possible to construct a path complying with the other constraints expressed in the request.

A successful path computation reported in a path computation reply message (PCRep) MUST include an ERO to specify the path that has been computed as specified in [RFC5440] following the sequence of domains.

In a PCRep, PCE MAY also supply IRO (with Domain-Sequence information) with the NO-PATH object indicating that the set of elements (domains) of the request's IRO prevented the PCEs from finding a path.

Following processing rules apply for Domain-Sequence in IRO -

o The Area subobject is optional.

o The AS subobject is optional.

o If an Area subobject is present then it changes the Area for all subsequent subobjects that do not change the area themselves. Subobjects that may change the Area are:

   * IP addresses that are present in another Area (via IPv4/IPv6 subobject)

   * Area ID (via OSPF/ISIS area subobjects)

   * AS number of another AS (via AS subobjects)

- o If an Area subobject is not preceded by an AS subobject then the receiver MUST act as though there is no change in AS from the previous subobject.

- o If an AS subobject is present then it changes the AS for all subsequent subobjects that do not change the AS themselves. Subobjects that may change the AS are:

  * IP addresses that are present in another AS (via IPv4/IPv6 subobject)

  * Unnumbered interfaces that are present in another AS (via Unnumbered Interface ID subobject)

  * AS number (via AS subobjects)

- o AS and Area subobjects may be interspersed with other subobjects without change to the previously specified processing of those subobjects in the IRO.

3.5. Exclude Route Object (XRO)

The Exclude Route Object (XRO) [RFC5521] is an optional object used to specify exclusion of certain abstract nodes or resources from the whole path.

3.5.1. Subobjects

Some subobjects to be used in XRO as defined in [RFC3209], [RFC3477], [RFC4874], and [RFC5520], but new subobjects related to Domain-Sequence are needed.

This document extends the support for 4-Byte AS numbers and IGP Areas.

Type	Subobject
TBD1	Autonomous system number (4 Byte)
TBD2	OSPF Area id
TBD3	ISIS Area id

3.5.1.1. Autonomous system

The new subobjects to support 4 byte AS and IGP (OSPF / ISIS) Area MAY also be used in the XRO to specify exclusion of certain domains in the path computation procedure.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|X|    Type     |     Length    |         Reserved              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          AS-ID (4 bytes)                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The X-bit indicates whether the exclusion is mandatory or desired.

0: indicates that the AS specified MUST be excluded from the path computed by the PCE(s).

1: indicates that the AS specified SHOULD be avoided from the inter-domain path computed by the PCE(s), but MAY be included subject to PCE policy and the absence of a viable path that meets the other constraints.

All other fields are consistent with the definition in Section 3.4.

3.5.1.2. IGP Area

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

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

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|X|    Type     |     Length    |         Reserved              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    OSPF Area Id (4 bytes)                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The X-bit indicates whether the exclusion is mandatory or desired.

0: indicates that the OSFF Area specified MUST be excluded from the path computed by the PCE(s).

1: indicates that the OSFF Area specified SHOULD be avoided from the inter-domain path computed by the PCE(s), but MAY be included subject to PCE policy and the absence of a viable path that meets the other constraints.

All other fields are consistent with the definition in Section 3.4.

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 [ISO10589]. The subobject is encoded as follows:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|X|    Type     |     Length    |  Area-Len     |  Reserved     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//                        IS-IS Area ID                        //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The X-bit indicates whether the exclusion is mandatory or desired.

0: indicates that the ISIS Area specified MUST be excluded from the path computed by the PCE(s).

1: indicates that the ISIS Area specified SHOULD be avoided from the inter-domain path computed by the PCE(s), but MAY be included subject to PCE policy and the absence of a viable path that meets the other constraints.

All other fields are consistent with the definition in Section 3.4.

All the processing rules are as per [RFC5521].

Note that, if a PCE receives an XRO in a PCReq message that contains subobjects defined in this document, that it does not recognize, it will respond according to the rules for a malformed object as per [RFC5440]. The PCE MAY also include the XRO in the PCErr to indicate in which case the XRO SHOULD be terminated immediately after the unrecognized subobject.

3.6. Explicit Exclusion Route Subobject (EXRS)

Explicit Exclusion Route Subobject (EXRS) [RFC5521] is used to specify exclusion of certain abstract nodes between a specific pair of nodes.

The EXRS subobject can carry any of the subobjects defined for inclusion in the XRO, thus the new subobjects to support 4 byte AS and IGP (OSPF / ISIS) 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 scope of the exclusion is limited to the single hop between the previous and subsequent elements in the IRO.

All the processing rules are as per [RFC5521].

Note that, if a PCE that supports the EXRS in an IRO, parses an IRO, and encounters an EXRS that contains subobjects defined in this document, that it does not recognize, it will act according to the setting of the X-bit in the subobject as per [RFC5521].

3.7. Explicit Route Object (ERO)

The Explicit Route Object (ERO) [RFC5440] is used to specify a computed path in the network. PCEP ERO subobject types correspond to RSVP-TE ERO subobject types as defined in [RFC3209], [RFC3473], [RFC3477], [RFC4873], [RFC4874], and [RFC5520]. The subobjects related to Domain-Sequence are further defined in [DOMAIN-SUBOBJ].

The new subobjects to support 4 byte AS and IGP (OSPF / ISIS) Area can 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 LSP). Using this concept of abstraction, an explicitly routed LSP can be specified as a sequence of domains.

In case of Hierarchical PCE [RFC6805], a Parent PCE can be requested to find the Domain-Sequence. Refer example in Section 4.6. The ERO in reply from parent PCE can then be used in Per-Domain path computation or BRPC.

If a PCC receives an ERO in a Path Computation response (PCRep) message that contains subobject defined in this document, that it does not recognize, it will respond according to the rules for a malformed object as per [RFC5440]. The PCC MAY also include the ERO in the PCErr to indicate in which case the ERO SHOULD be terminated immediately after the unrecognized subobject.

4. Other Considerations

The examples in this section are for illustration purposes only; to highlight how the new subobjects could be encoded. They are not meant to be an exhaustive list of all possible usecases and combinations.

4.1. Inter-Area Path Computation

In an inter-area path computation where the ingress and the egress nodes belong to different IGP areas within the same AS, the Domain-Sequence could be represented using a ordered list of Area subobjects.

```
    -----------------                              -----------------
   |                 |                            |                 |
   |          +--+   |                            |     +--+        |
   | +--+     |  |   |                            |     |  |        |
   | |  |     +--+   |                            |     +--+   +--+ |
   | +--+            |                            |            |  | |
   |                 |                            |            +--+ |
   |        +--+     |                            |                 |
   |        |  |     |                            |     +--+        |
   |        +--+     |                            |     |  |        |
   |                 | -------------------------- |     +--+        |
   |                +--+                       +--+                 |
   |                |  |         +--+          |  |                 |
   |Area 2          +--+         |  |          +--+  Area 4          |
    ----------------- |          +--+            | -----------------
                      |                          |
                      |                +--+      |
                      |    +--+        |  |      |
                      |    |  |        +--+      |
                      |    +--+                  |
                      |                          |
                      |                          |
                      |                          |
                      |                          |
                      |           +--+           |
                      |           |  |           |
                      |           +--+           |
    ----------------- |                          | ------------------
   |                 +--+                      +--+                  |
   |                 |  |                      |  |                  |
   |                 +--+    Area 0            +--+                  |
   |                 | -------------------------- |     +--+         |
   |          +--+   |                            |     |  |         |
   |          |  |   |                            |     +--+         |
   | +--+     +--+   |                            |                  |
   | |  |            |                            |            +--+  |
   | +--+            |                            |            |  |  |
   |                 |                            |            +--+  |
   |       +--+      |                            |                  |
   |       |  |      |                            |     +--+         |
   |       +--+      |                            |     |  |         |
   |                 |                            |     +--+         |
   |                 |                            |                  |
   | Area 1          |                            |  Area 5          |
    -----------------                              ------------------
```

Figure 1: Inter-Area Path Computation

AS Number is 100.

This could be represented in the IRO as:

```
+---------+ +---------+ +---------+
|IRO      | |Sub      | |Sub      |
|Object   | |Object   | |Object   |
|Header   | |Area 0   | |Area 4   |
|         | |         | |         |
|         | |         | |         |
+---------+ +---------+ +---------+

or

+---------+ +---------+ +---------+ +---------+
|IRO      | |Sub      | |Sub      | |Sub      |
|Object   | |Object   | |Object   | |Object   |
|Header   | |Area 2   | |Area 0   | |Area 4   |
|         | |         | |         | |         |
|         | |         | |         | |         |
+---------+ +---------+ +---------+ +---------+

or

+---------+ +---------+ +---------+ +---------+ +---------+
|IRO      | |Sub      | |Sub      | |Sub      | |Sub      |
|Object   | |Object AS| |Object   | |Object   | |Object   |
|Header   | |100      | |Area 2   | |Area 0   | |Area 4   |
|         | |         | |         | |         | |         |
|         | |         | |         | |         | |         |
+---------+ +---------+ +---------+ +---------+ +---------+
```

The Domain-Sequence can further include encompassing AS information in the AS subobject.

4.2. Inter-AS Path Computation

In inter-AS path computation, where ingress and egress belong to different AS, the Domain-Sequence could be represented using an ordered list of AS subobjects. The Domain-Sequence can further include decomposed area information in the Area subobject.

4.2.1. Example 1

As shown in Figure 2, where AS has a single area, AS subobject in the domain-sequence can uniquely identify the next domain and PCE.

```
              AS A                AS E                AS C
         <------------->     <---------->      <------------->

                  A4----------E1---E2---E3---------C4
                 /           /                       \
               /           /                           \
             /           /       AS B                    \
           /           /      <---------->                 \
     Ingress------A1---A2------B1---B2---B3------C1---C2------Egress
           \                                    /          /
             \                                /          /
               \                            /          /
                 \                        /          /
                  A3----------D1---D2---D3---------C3

                              <---------->
                                  AS D

     * All AS have one area (area 0)
```

Figure 2: Inter-AS Path Computation

This could be represented in the IRO as:

```
+-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    |
|Object | |Object | |Object |
|Header | |AS B   | |AS C   |
|       | |       | |       |
+-------+ +-------+ +-------+
```

or

```
+-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object |
|Header | |AS A   | |AS B   | |AS C   |
|       | |       | |       | |       |
+-------+ +-------+ +-------+ +-------+
```

or

```
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object | |Object | |Object | |Object |
|Header | |AS A   | |Area 0 | |AS B   | |Area 0 | |AS C   | |Area 0 |
|       | |       | |       | |       | |       | |       | |       |
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+ +-------+
```

Note that to get a domain disjoint path, the ingress could also request the backup path with -

```
+-------+ +-------+
|XRO    | |Sub    |
|Object | |Object |
|Header | |AS B   |
|       | |       |
+-------+ +-------+
```

As described in Section 3.4.3, domain subobject in IRO changes the domain information associated with the next set of subobjects; till you encounter a subobject that changes the domain too. Consider the following IRO:

```
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object | |Object | |Object |
|Header | |AS B   | |IP     | |IP     | |AS C   | |IP     |
|       | |       | |B1     | |B3     | |       | |C1     |
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+
```

On processing subobject "AS B", it changes the AS of the subsequent subobjects till we encounter another subobject "AS C" which changes the AS for its subsequent subobjects.

Consider another IRO:

```
+-------+ +-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object | |Object |
|Header | |AS D   | |IP     | |IP     | |IP     |
|       | |       | |D1     | |D3     | |C3     |
+-------+ +-------+ +-------+ +-------+ +-------+
```

Here as well, on processing "AS D", it changes the AS of the subsequent subobjects till you encounter another subobject "C3" which belong in another AS and changes the AS for its subsequent subobjects.

Further description for the Boundary Node and Inter-AS-Link can be found in Section 4.3.

4.2.2. Example 2

In Figure 3, AS 200 is made up of multiple areas.

```
               |
               |  +-------------+                +----------------+
               |  |Area 2       |                |Area 4          |
               |  |         +--+|                |          +--+  |
               |  |         |  ||                |          | B|  |
               |  |  +--+   +--+|                |   +--+   +--+  |
               |  |  |  |       |                |   |  |         |
               |  |  +--+       |                |   +--+         |
               |  |        +--+ |                |          +--+  |
               |  |        |  | |                |          |  |  |
               |  |        +--+ |                |   +--+   +--+  |
               |  |  +--+       |+--------------+|   |  |         |
               |  |  |  |       +--+          +--+   +--+         |
+-------------+|  |  +--+       |  |          |  |                |
|             ||  |             +--+          +--+                |
|         +--+||  +-------------+|              |+----------------+
|         |  |||                 |     +--+     |
|         +--+||                 |     |  |     |
|    +--+     ||                 |     +--+     |
|    |  |  +---+                +--+            |
|    +--+  |   |----------------|  |            |
|          +---+   Inter-AS     +--+   +--+     |
|+--+         ||    Links        |     |  |     |
```

```
||A |      +---+                +--+   +--+     |
|+--+      |   |----------------|  |            |
|          +---+                +--+   +--+     |
|    +--+     ||  +------------+ |     |  |     |+----------------+
|    |  |     ||  |Area 3      +--+    +--+   +--+ Area 5         |
|    +--+     ||  |            |  |           |  |                |
|             ||  |            +--+           +--+                |
|         +--+||  |       +--+ | |  Area 0      ||   +--+         |
|         |  |||  |       |  | | +--------------+|   |  |         |
|         +--+||  |       +--+ |                 |   +--+         |
|             ||  |            |                 |          +--+  |
|Area 0       ||  |   +--+     |                 |   +--+   |  |  |
+-------------+|  |   |  |     |                 |   |  |   +--+  |
               |  |   +--+  +--+                 |   +--+         |
               |  |         |  |                 |                |
               |  |         +--+                 |          +--+  |
               |  |   +--+     |                 |          | C|  |
               |  |   |  |     |                 |          +--+  |
               |  |   +--+     |                 |                |
               |  |            |                 |                |
               |  +------------+                 +----------------+
               |
               |
    AS 100     |  AS 200
               |
```

Figure 3: Inter-AS Path Computation

The Domain-Sequence for the LSP (A-B) can be carried in the IRO as shown below:

```
+-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object |
|Header | |AS 200 | |Area 0 | |Area 4 |
|       | |       | |       | |       |
+-------+ +-------+ +-------+ +-------+
```

or

```
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+
|IRO    | |Sub    | |Sub    | |Sub    | |Sub    | |Sub    |
|Object | |Object | |Object | |Object | |Object | |Object |
|Header | |AS 100 | |Area 0 | |AS 200 | |Area 0 | |Area 4 |
|       | |       | |       | |       | |       | |       |
+-------+ +-------+ +-------+ +-------+ +-------+ +-------+
```

The Domain-Sequence for the LSP (A-C) can be carried in the IRO as shown below:

```
+--------+ +--------+ +--------+ +--------+
|IRO     | |Sub     | |Sub     | |Sub     |
|Object  | |Object  | |Object  | |Object  |
|Header  | |AS 200  | |Area 0  | |Area 5  |
|        | |        | |        | |        |
+--------+ +--------+ +--------+ +--------+
```

or

```
+--------+ +--------+ +--------+ +--------+ +--------+ +--------+
|IRO     | |Sub     | |Sub     | |Sub     | |Sub     | |Sub     |
|Object  | |Object  | |Object  | |Object  | |Object  | |Object  |
|Header  | |AS 100  | |Area 0  | |AS 200  | |Area 0  | |Area 5  |
|        | |        | |        | |        | |        | |        |
+--------+ +--------+ +--------+ +--------+ +--------+ +--------+
```

4.3. Boundary Node and Inter-AS-Link

A PCC or PCE can include additional constraints covering which Boundary Nodes (ABR or ASBR) or Border links (Inter-AS-link) to be traversed while defining a Domain-Sequence. In which case the Boundary Node or Link can be encoded as a part of the Domain-Sequence.

Boundary Nodes (ABR / ASBR) can be encoded using the IPv4 or IPv6 prefix subobjects usually the loopback address of 32 and 128 prefix length respectively. An Inter-AS link can be encoded using the IPv4 or IPv6 prefix subobjects or unnumbered interface subobjects.

For Figure 1, an ABR (say 203.0.113.1) to be traversed can be specified in IRO as:

```
+----------+ +----------+ +----------++----------+ +----------+
|IRO       | |Sub       | |Sub       ||Sub       | |Sub       |
|Object    | |Object    | |Object    ||Object    | |Object    |
|Header    | |Area 2    | |IPv4      ||Area 0    | |Area 4    |
|          | |          | |203.0.    ||          | |          |
|          | |          | |112.1     ||          | |          |
+----------+ +----------+ +----------++----------+ +----------+
```

For Figure 3, an inter-AS-link (say 198.51.100.1 - 198.51.100.2) to be traversed can be specified as:

```
          +---------+  +---------+ +---------+ +---------+
          |IRO      |  |Sub      | |Sub      | |Sub      |
          |Object   |  |Object AS| |Object   | |Object AS|
          |Header   |  |100      | |IPv4     | |200      |
          |         |  |         | |198.51.  | |         |
          |         |  |         | |100.2    | |         |
          +---------+  +---------+ +---------+ +---------+
```

4.4. PCE Serving multiple Domains

A single PCE can be responsible for multiple domains; for example PCE function deployed on an ABR could be responsible for multiple areas. A PCE which can support adjacent domains can internally handle those domains in the Domain-Sequence without any impact on the other domains in the Domain-Sequence.

4.5. P2MP

[RFC7334] describes an experimental inter-domain P2MP path computation mechanism where the path domain tree is described as a series of Domain-Sequences, an example is shown in the below figure:

```
   D1-D3-D6, D1-D3-D5 and D1-D2-D4.
               D1
              /  \
             D2   D3
            /    /  \
           D4   D5   D6
```

The domain sequence handling described in this document could be applied to P2MP path domain tree.

4.6. Hierarchical PCE

In case of H-PCE [RFC6805], the parent PCE can be requested to determine the Domain-Sequence and return it in the path computation reply, using the ERO. . For the example in section 4.6 of [RFC6805], the Domain-Sequence can possibly appear as:

```
+---------+ +---------+ +---------+ +---------+
|ERO      | |Sub      | |Sub      | |Sub      |
|Object   | |Object   | |Object   | |Object   |
|Header   | |Domain 1 | |Domain 2 | |Domain 3 |
|         | |         | |         | |         |
|         | |         | |         | |         |
+---------+ +---------+ +---------+ +---------+
```

or

```
+---------+ +---------+ +---------+
|ERO      | |Sub      | |Sub      |
|Object   | |Object   | |Object   |
|Header   | |BN 21    | |Domain 3 |
|         | |         | |         |
|         | |         | |         |
+---------+ +---------+ +---------+
```

4.7. Relationship to PCE Sequence

Instead of a Domain-Sequence, a sequence of PCEs MAY be enforced by policy on the PCC, and this constraint can be carried in the PCReq message (as defined in [RFC5886]).

Note that PCE-Sequence can be used along with Domain-Sequence in which case PCE-Sequence MUST have higher precedence in selecting the next PCE in the inter-domain path computation procedures.

4.8. Relationship to RSVP-TE

[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 LSP. It further defines a subobject for AS but with a 2-Byte AS Number.

[DOMAIN-SUBOBJ] extends the notion of abstract nodes by adding new subobjects for IGP Areas and 4-byte AS numbers. These subobjects can be included in Explicit Route Object (ERO), Exclude Route object (XRO) or Explicit Exclusion Route Subobject (EXRS) in RSVP-TE.

In any case subobject type defined in RSVP-TE are identical to the subobject type defined in the related documents in PCEP.

5. IANA Considerations

5.1. New Subobjects

IANA maintains the "Path Computation Element Protocol (PCEP) Numbers" at http://www.iana.org/assignments/pcep/pcep.xhtml. Within this registry IANA maintains two sub-registries:

- "IRO Subobjects": http://www.iana.org/assignments/pcep/pcep.xhtml#iro-subobject
- "XRO Subobjects": http://www.iana.org/assignments/pcep/pcep.xhtml#xro-subobject

Upon approval of this document, IANA is requested to make identical additions to these registries as follows:

Subobject	Type	Reference
TBD1	4 byte AS number	[This I.D.]
TBD2	OSPF Area ID	[This I.D.]
TBD3	IS-IS Area ID	[This I.D.]

6. Security Considerations

This document specifies a standard representation of Domain-Sequence and new subobjects, which could be used in inter-domain PCE scenarios as explained in other RFC and drafts. The new subobjects and Domain-Sequence mechanisms defined in this document allow finer and more specific control of the path computed by a cooperating PCE(s). Such control increases the risk if a PCEP message is intercepted, modified, or spoofed because it allows the attacker to exert control over the path that the PCE will compute or to make the path computation impossible. Therefore, the security techniques described in [RFC5440] are considered more important.

Note, however, that the Domain-Sequence mechanisms also provide the operator with the ability to route around vulnerable parts of the network and may be used to increase overall network security.

7. Manageability Considerations

7.1. Control of Function and Policy

The exact behaviour with regards to desired inclusion and exclusion of domains MUST be available for examination by an operator and MAY be configurable. Manual configurations is needed to identify which PCEP peers understand the new domain subobjects defined in this document.

7.2. Information and Data Models

A MIB module for management of the PCEP is being specified in a separate document [RFC7420]. That MIB module allows examination of individual PCEP messages, in particular requests, responses and errors. The MIB module MUST be extended to include the ability to view the Domain-Sequence extensions defined in this document.

7.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

7.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440].

7.5. Requirements On Other Protocols

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, a signaling message may use the domain identifiers. The Subobjects defined in this document SHOULD be supported by RSVP-TE. [DOMAIN-SUBOBJ] extends the notion of abstract nodes by adding new subobjects for IGP Areas and 4-byte AS numbers.

Apart from this, mechanisms defined in this document do not imply any requirements on other protocols in addition to those already listed in [RFC5440].

7.6. Impact On Network Operations

The mechanisms described in this document can provide the operator with the ability to exert finer and more specific control of the path computation by inclusion or exclusion of domain subobjects. There may be some scaling benefit when a single domain subobject may substitute for many subobjects and can reduce the overall message size and processing.

Backward compatibility issues associated with the new subobjects arise when a PCE does not recognize them, in which case PCE responds according to the rules for a malformed object as per [RFC5440]. For successful operations the PCEs in the network would need to be upgraded.

8. Acknowledgments

Authors would like to especially thank Adrian Farrel for his detailed reviews as well as providing text to be included in the document.

Further, we would like to thank Pradeep Shastry, Suresh Babu, Quintin Zhao, Fatai Zhang, Daniel King, Oscar Gonzalez, Chen Huaimo, Venugopal Reddy, Reeja Paul, Sandeep Boina, Avantika and Sergio Belotti 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.

[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.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5441] Vasseur, JP., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, April 2009.

[RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", RFC 5521, April 2009.

[RFC6805] King, D. and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, November 2012.

[ISO10589]
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, 1992.

[IRO-UPDATE]
Dhody, D., "Update to Include Route Object (IRO) specification in Path Computation Element communication Protocol (PCEP. (draft-dhody-pce-iro-update-02)", December 2014.

[DOMAIN-SUBOBJ]
Dhody, D., Palle, U., Kondreddy, V., and R. Casellas, "Domain Subobjects for Resource ReserVation Protocol - Traffic Engineering (RSVP-TE). (draft-ietf-teas-rsvp-te-domain-subobjects-00)", April 2015.

9.2. Informative References

[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.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.

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

[RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.

[RFC5886] Vasseur, JP., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", RFC 5886, June 2010.

[RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, December 2012.

[RFC7334] Zhao, Q., Dhody, D., King, D., Ali, Z., and R. Casellas, "PCE-Based Computation Procedure to Compute Shortest Constrained Point-to-Multipoint (P2MP) Inter-Domain Traffic Engineering Label Switched Paths", RFC 7334, August 2014.

[RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", RFC 7420, December 2014.

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India

EMail: dhruv.ietf@gmail.com

Udayasree Palle
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India

EMail: udayasree.palle@huawei.com

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

EMail: ramon.casellas@cttc.es