

Path Computation Element Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2019

O. Dugeon
J. Meuric
Orange Labs
Y. Lee
Huawei Technologies
D. Ceccarelli
Ericsson
March 04, 2019

PCEP Extension for Stateful Inter-Domain Tunnels
draft-dugeon-pce-stateful-interdomain-02

Abstract

This document proposes to combine a Backward Recursive or Hierarchical method in Stateful PCE with PCInitiate message to setup independent paths per domain, and combine these different paths together in order to operate them as end-to-end inter-domain paths without the need of signaling session between AS border routers. A new Stitching Label is defined, new Path Setup Types and a new Association Type are considered for that purpose.

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 RFC 2119 [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 5, 2019.

Copyright Notice

Copyright (c) 2019 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.	General assumptions	5
1.2.	Terminology	6
2.	Stitching Label	7
2.1.	Definition	8
2.2.	Inter-domain LSP-TYPE	8
3.	Backward Recursive PCInitiate procedure	9
3.1.	Mode of operation	9
3.2.	Example	12
3.3.	Inter-domain LSP setup procedure completion failure	13
4.	Hierarchical PCInitiate procedure	14
4.1.	Mode of operation	14
4.2.	Inter-domain LSP setup procedure completion failure	16
4.3.	Example for Stateful H-PCE Stching procedure	17
5.	Inter-domain LSP Management	21
5.1.	Identification of inter-domain tunnels	21
5.2.	Inter-domain association group	21
5.3.	Inter-domain LSP management	22
5.4.	Modification of inter-domain LSP	23
5.5.	Removal of inter-domain LSP	23
6.	Applicability	24
6.1.	RSVP-TE	24
6.2.	Segment Routing	24
6.3.	Mixing technology	25
7.	IANA Considerations	26
7.1.	Path Setup Type values	26
7.2.	Association Type value	27
7.3.	PCEP Error values	27
8.	Security Considerations	27
9.	Acknowledgements	27
10.	Disclaimer	28

11. References	28
11.1. Normative References	28
11.2. Informative References	29
Authors' Addresses	30

1. Introduction

The Path Computation Element (PCE) working group (WG) has produced a set of RFCs to standardize the behavior of the Path Computation Element as a tool to help MPLS-TE, GMPLS LSP tunnels and Segment Routing paths placement. This also includes the ability to compute inter-domain LSPs or Segment Routing paths following a distributed or hierarchical approach. To complement the original stateless mode, a stateful mode has been added. In particular, the new PCInitiate message allows a PCE to directly ask a PCC to setup an MPLS-TE, GMPLS LSP tunnel or a Segment Routing path. However, once computed, the inter-domain LSPs or Segment Routing path are hard to setup in the underlying network. Especially, in operational network, RSVP-TE signaling is not enabled between AS border routers. But, such RSVP-TE signaling is mandatory to setup contiguous LSP tunnels or to stitch or nest independent LSP tunnels to form the end-to-end inter-domain paths.

Looking to the different RFCs that describe the PCE architecture and in particular PCE based architecture [RFC4655], PCE protocol [RFC5440], BRPC [RFC5441] and H-PCE [RFC6805], the Path Computation Element (PCE) is able to compute inter-domain paths in complement to intra-domain computation. Such inter-domain paths could then serve as the Explicit Route Object input for the RSVP-TE signaling to setup the tunnels within the underlying network. Three kinds of inter-domain paths could be established:

- o Contiguous tunnel ([RFC3209] and [RFC3473]): The RSVP-TE signaling crosses the boundary between two domains, e.g. between two AS Border Routers (ASBRs) as if they were two routers of the same domain. This kind of tunnel is not recommended mostly for security and scalability purpose. In addition, the initiating domain imposes huge constraints on subsequent domains, because they undergo the tunnel request without being able to control it.
- o Stitching tunnel ([RFC5150]): Each domain establishes in its own network the corresponding part of the inter-domain path independently. Then, a second end-to-end RSVP-TE Path message is sent by the initiating domain to stitch the different tunnel parts to form the inter-domain path. In fact, this second RSVP-TE Path message is used by border nodes to exchange the label that must be used by the previous domain to send the traffic in order that the

MPLS packets follow the next LSP tunnel in the following domain. These labels are conveyed in the RSVP-TE Resv message.

- o Nesting tunnel ([RFC4206]): This is similar to the stitching mode but, this time, with the possibility to setup tunnel hierarchy. For example, an LSP tunnel between two edge domains crossing a transit domain could be carried over a tunnel of a higher level in the transit domain. Again, a second end-to-end RSVP-TE Path message is sent from the source to the destination. Labels that must be used by the ASBRs of transit domains to identify flows to be nested are carried by the RSVP-TE Resv message.

In all case, RSVP-TE signaling must be exchanged between the different domains. However, from an operational point of view, looking to different networks under the responsibility of different administrative entities, only BGP sessions are setup and configured between ASBRs. Technologically speaking, this is possible and many RFCs describe how to use RSVP-TE for inter-domain. But, due to security, scalability, management and contract constraints, RSVP-TE is not exposed at the network boundary. To circumvent some of the security issues, RSVP-TE can be carried inside an IPsec tunnel between ASBRs, but, this does not eliminate the scalability aspect nor the constraints imposed by setting up inter-domain paths.

The purpose of this memo is to take the benefit of PCE Stateful [RFC8231] and PCE Initiated [RFC8281] modes to stitch or nest inter-domain paths directly using PCEP between domains' PCEs instead of using RSVP-TE signaling at the inter-domain border nodes, while keeping each operator free to independently setup their respective part of the inter-domain paths. PCInitiate message is used in a Backward Recursive way like the PCReq message in BRPC [RFC5441], to recursively setup the end-to-end tunnel. PCRep message is used to automatically stitch or nest the different local LSP tunnels. And, PCRep in conjunction of PCUpd messages are used to maintain, modify and remove inter-domain paths. This method is also applicable to Segment Routing to build inter-domain segment paths.

H-PCE [RFC6805] describes a Hierarchical PCE architecture which can be used for computing end-to-end paths for inter-domain MPLS Traffic Engineering (TE) and GMPLS Label Switched Paths (LSPs). Within this architecture, the Parent PCE (P-PCE) is used to compute a multi-domain path based on the domain connectivity information. A Child PCE (C-PCE) may be responsible for a single domain or multiple domains, it is used to compute the intra-domain path based on its domain topology information.

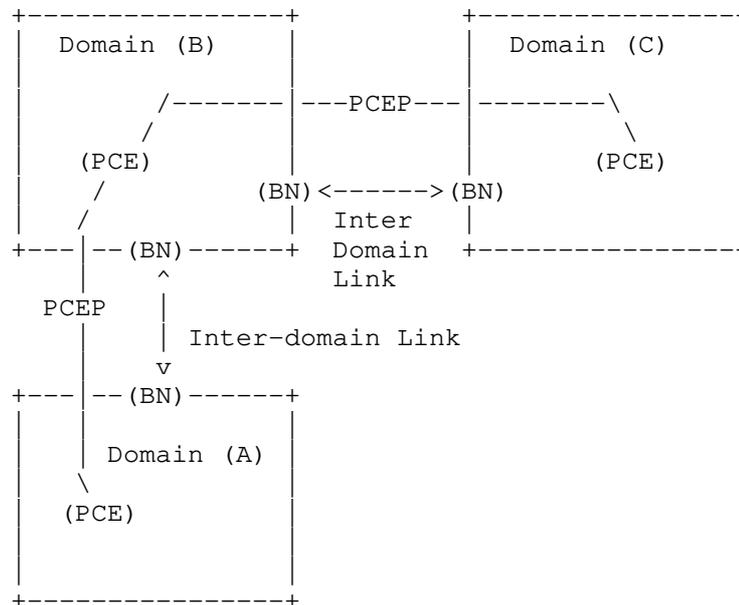
Stateful H-PCE [I-D.ietf-pce-stateful-hpce] presents general considerations for stateful PCE(s) in hierarchical PCE architecture.

In particular, the behavior changes and additions to the existing stateful PCE mechanisms (including PCE-initiated LSP setup and active PCE usage) in the context of networks using the H-PCE architecture. Section 3.3.1 [I-D.ietf-pce-stateful-hpce] describes the per domain stitched LSP mode, where the individual per domain LSP are stitched together. PCInitiate message is also used to stitch the end-to-end tunnel. See section 4 for details.

1.1. General assumptions

In the rest of this document, we used the same references as per BRPC [RFC5441] and make the following set of assumptions (see figure below):

- o Domain refers to an IGP area or an Autonomous System (AS).
- o Inter-domain path is used to refer to a path that cross two or more different domains as defined previously,
- o At least, one PCE is deployed in each domain. These PCE are all stateful active capable and could request to enforce LSP tunnels in their respective domain by means of PCInitiate messages.
- o LSRs, including border nodes, are PCC enable and support stateful active mode. PCEP sessions is established between these routers and their domains' PCE.
- o Each PCE establishes a PCEP session with its respective neighbor domain's PCE. The way a PCE discover its neighboring PCE is out of scope of this draft. These information could be fulfill administratively or automatically discovered through, for example per draft 'BGP Extensions for Path Computation Element (PCE) Discovery' [I-D.dong-pce-discovery-proto-bgp].
- o PCEs are able to compute and end-o-end path as per BRPC procedure [RFC5441] or as per H-PCE procedure (stateless [RFC6805] or stateful [I-D.ietf-pce-stateful-hpce]).
- o Tunnels refer to LSPs setup by mean of RSVP-TE or Segment Path in a Segment Routing network.



Example of the representation of 3 domains with 3 PCEs

1.2. Terminology

ABR: Area Border Routers. Routers used to connect two IGP areas (areas in OSPF or levels in IS-IS).

AS: Autonomous System

ASBR: Autonomous System Border Router. Router used to connect together ASes of the same or different service providers via one or more inter-AS links.

Border Node (BN): a boundary node is either an ABR in the context of inter-area Traffic Engineering or an ASBR in the context of inter-AS Traffic Engineering.

BN-en(i): Entry BN of domain(i) connecting domain(i-1) to domain(i) along a determined sequence of domains. Multiple entry BN-en(i) could be used to connect domain(i-1) to domain(i).

BN-ex(i): Exit BN of domain(i) connecting domain(i) to domain(i+1) along a determined sequence of domains. Multiple exit BN-ex(i) could be used to connect domain(i) to domain(i+1).

Domains: Autonomous System (AS) or IGP Area. An Autonomous System is composed by one or more IGP area.

ERO(i): The Explicit Route Object scoped to domain(i)

IGP-TE: Interior Gateway Protocol with Traffic Engineering support. Both OSPF-TE and IS-IS-TE are identified in this category.

Inter-domain path: A path that crosses two or more domains through a pair of Border Node (BN-ex, BN-en).

LK(i): A Link that connect BN-ex(i-1) to BN-en(i). Note that BN-ex(i-1) could be connected to BN-en(i) by more than one link. LK(i) serves to identify which of the multiple links will be used for the inter-domain LSP setup.

Local LSP tunnel: A LSP tunnel that do not cross a domain. It is setup between entry BN-en to output BN-ex, any source to output BN-ex or entry BN-en to any destination of the same domain. This LSP could be enforce by means of RSVP-TE signaling or Segment Routing labels stack.

Local LSP tunnel(i): A local LSP tunnel of domain(i)

PLSP-ID(i): A PLSP-ID that identify the local tunnel part of an inter-domain tunnel in the domain(i).

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.

PCE(i) is a PCE with the scope of domain(i).

PST: Path Setup Type

R(i, j): The router j of domain i

Stitching Label (SL): A dedicated label that is used to stitch two RSVP-TE tunnels or two Segment Routing paths.

SL(i): A Stitching Label that link domain(i-1) to domain(i).

2. Stitching Label

This section introduce the concept of Stitching Label that allows stitching and nesting of local LSP tunnels in order to form inter-domain path that cross several different domains.

2.1. Definition

The operation of stitch or nest a local LSP tunnel(i) to a local LSP tunnel(i+1) in order to form an inter-domain path simply consist in defining the label that the output BN-ex(i) will use to send its traffic to the entry BN-en(i+1). Indeed, the entry BN-en(i+1) needs to identify the incoming traffic i.e. IP packets, in order to know if this traffic must follow the local LSP tunnel(i+1) or not. Forwarding Equivalent Class (FEC) could be used for that purpose. But, when stitching or nesting tunnels, the FEC is reduce to the incoming label that the entry BN-en(i+1) as chosen for the local LSP tunnel(i+1).

In this memo, we introduce the named of 'Stitching Label (SL)' to designate this label. Such label is usually exchange between output BN-ex(i) and entry BN-en(i+1) with the RSVP-TE signaling. But, as we want to avoid to use RSVP-TE signaling due to operational constraints, this Stitching Label will be convey by PCEP protocol. In fact, the Explicit Route Object (ERO) and the Record Route Object (RRO) are defined in order to transport this Stitching Label in the RSVP-TE signaling. As PCEP protocol used RSVP-TE Objects, and in particular the ERO and RRO, it is able to convey the Stitching Label without any modification of the PCEP protocol nor the PCE or RSVP-TE Objects.

As per RFC4003 [RFC4003], the Stitching Label will be convey as a companion of an IP address. In our case, this is one of the IP address of the link LK(i) which connects BN-ex(i) to BN-en(i+1) and carries the traffic from the domain(i) to domain(i+1). It is left to implementation to select which of the two IP address of the link LK(i) is used.

2.2. Inter-domain LSP-TYPE

However, even if PCEP could convey the Stitching Label, a PCC is not aware that a PCE requests or provides such label. For that purpose, this memo propose to use the PST as defined in [RFC8408] with new values (See IANA section of this memo) defined as follow:

- o TBD1: Inter-Domain Traffic engineering end-to-end path is setup using Backward Recursive or Hierarchical method. This new PST value MUST be set in a PCInitiate messages sends by a PCE(i) to its neighbor PCE(i+1) in the Backward Recursive method or by the Parent PCE to the Child PCE(i) to initiate a new inter-domain path. In turn, neighbor PCE(i+1) or Child PCE(i) MUST return a Stitching Label SL with the IP address of the associated link in the RRO of the PCRpt message to PCE(i) or Parent PCE.

- o TBD2: Inter-Domain Traffic engineering local path is setup using RSVP-TE. This new PST value MUST be set in the PCInitiate message sends by a PCE(i) requesting to a PCC of domain(i) to initiate a new local LSP tunnel(i) which is part of an inter-domain path. This PST value MUST be used by the PCE(i) only after receiving a PCInitiate message with an PST equal to TBD1 from a neighbor PCE(i+1) in the Backward Recursive method or Parent PCE in the Hierarchical method. In turn, the PCC of domain(i) MUST return a Stitching Label SL with the IP address of associated link in the RRO of the PCRpt message.
- o TBD3: Inter-Domain Traffic engineering local path is setup using Segment Routing. This new PST value MUST be set in the PCInitiate message sends by a PCE(i) requesting to a PCC of domain(i) to initiate a new Segment Routing path which is part of an inter-domain Segment Routing path. This PST value MUST be used by the PCE(i) only after receiving a PCInitiate message with an PST equal to TBD1 from a neighbor PCE(i+1). In turn, the PCC MUST return a Stitching Label SL with the IP address of the associated link in the RRO of the PCRpt message.

3. Backward Recursive PCInitiate procedure

This section describes how to setup inter-domain paths than cross several different domains by using a Backward Recursive method which is compatible to inter-domain path computation by means of the BRPC procedure as describe in RFC5441 [RFC5441].

3.1. Mode of operation

This section describes how PCInitiate and PCRpt messages are combined between PCE in order to setup inter-domain paths between a source domain(1) to a destination domain(n). S and D are respectively the source and destination of the inter-domain path. Domain(1) and domain(n) are different and connected through 0 or more intermediate domains denoted domain(i) with $i = (2, n-1)$. Domains are directly connected when $n = 2$.

First, the PCE(1) runs standard BRPC algorithm as per RFC5441 [RFC5441] with its neighbor PCEs in order to compute the inter-domain path from S to D, where S and D are respectively a node in the domain(1) and domain(n). Path Key confidentiality as per RFC5520 [RFC5520] SHOULD be used to obfuscate the detailed ERO(i) of the different domains(i). The resulting ERO is of the form {S, PKS(1), BN-ex(1), ..., BN-en(i), PKS(i), BN-ex(i), ..., BN-en(n), PKS(n), D} when Path Key is used and of the form {S, R(1,1), ..., R(1,k), BN-ex(1), ..., BN-en(i), R(i,1), ..., R(i,l), BN-ex(i), ..., BN-en(n), R(n,1), ..., R(n,m), D} otherwise. As subsequent domains are not

aware about the final computed ERO in case of multiple VSPTs, the final ERO selected by the PCE(1) MUST be sent in the PCInitiate message to indicate to the subsequent PCEs which solution has been finally chosen. PCE(1) MUST ensure that this ERO is self comprehensive by subsequent PCEs. Indeed, when a PCE(i) receives the ERO, it MUST be able to verify that it is in the scope of this ERO and to determine the PCE(i+1). When Path Key is used, PCEs MUST encode the Path Key with a reachable IP address in order for previous PCEs in the AS chain to join them. When Path Key is not used, the PCEs MUST be able to retrieve IP address of the next PCE from the ERO.

The complete procedure with Path Key follow the different steps described below:

Steps 1: Initialization

Once ERO(S, D) is computed, PCE(1) sends a PCInitiate message to PCE(2) containing an ERO equal to {S, PKS(2), ..., PKS(i), ..., PKS(n), D}, PST = TBD1 and End-Points Object = (S, D). The ERO corresponds to the one PCE(1) has received from PCE(2) during the BRPC process in which only Path Key are kept. In case of multiple EROs, i.e. VSPT, PCE(1) has chosen one of them and used the selected one for the PCInitiate message. PKS(i) could be replaced by the full ERO description if Path Key is not used by PCE(i).

When PCE(i) receives a PCInitiate message from domain(i-1) with PST = TBD1 and ERO = {PKS(i), PKS(i+1), ..., PKS(n), D}, it sends a PCInitiate message to PCE(i+1) with a popped ERO and records its received PKS(i) part. All PCE(i)s generate the appropriate PCInitiate message to PCE(i+1) up to PCE(n), i.e. to the destination domain(n).

Steps 2: Actions taken at the destination domain(n) by PCE(n)

When PCInitiate message propagation reach the destination domain(n), PCE(n) retrieves the ERO from the PKS(n) if necessary and sends to BN-en(n) a PCInitiate message with the ERO(n) = {BN-en(n), ..., D}, PST = TBD2 and End-Points Object = {BN(n), D} in order to inform the PCC BN-en(n) that this local LSP tunnel(n) is part of an inter-domain path. When the PCC BN-en(n) received the PCInitiate message from its PCE(n), it setup the local LSP tunnels from entry BN-en(n) to D by means of RSVP-TE signaling with the given ERO(n). Once the tunnel setup, it chooses a free label for the Stitching Label SL(n) and add a new entry in its MPLS L(F)IB with this SL(n) label. Then, it sends a PCRpt message to its PCE(n) with an RRO equal to {[LK(n), SL(n)], RRO(n)} and PLSP-ID(n). Once PCE(n) receives the PCRpt from the PCC BN-en(n) with the RRO, PLSP-ID and PST = TBD2, it sends to the

PCE(n-1) a PCRpt containing the RRO equal to {[LK(n), SL(n)]} and PLSP-ID(n). PCE(n) MAY add {PKS(n), D} in the RRO.

Steps i: Actions performed by all intermediate domains(i), for i = 2 to n-1

1. When the PCE(i) receives a PCRpt message from domain(i+1) with PST = TBD1, RRO = {[LK(i+1), SL(i+1)]} and PLSP-ID(i+1), it retrieves the ERO(i) from the PKS(i), recorded in step 1, and sends to the PCC BN-en(i) a PCInitiate message with ERO = {ERO(i), [LK(i+1), SL(i+1)]}, PST = TBD2 and End-Points Object = {BN-en(i), BN-ex(i)} in order to inform the PCC BN-en(i) that this local LSP tunnel(i) is part of an inter-domain path.
2. When the PCC BN-en(i) received the PCInitiate message from its PCE(i), it setup the local LSP tunnels from BN-en(i) to BN-ex(i) by means of RSVP-TE signaling with the given ERO(i).
3. Egress Control mechanism, as per RFC4003 section 2.1 [RFC4003], is used to instruct the egress node of domain(i), i.e. BN-ex(i), to forward packets belonging to this tunnel with the Stitching Label. Both Stitching Label and IP address of outgoing interface are carried in the ERO = {..., [LK(i+1), SL(i+1)]} as the last SubObject in conformance to [RFC4003]. So that, BN-ex(i) installs in its MPLS L(F)IB the SWAP instruction to label SL(i+1) with forward to LK(i+1).
4. Once the tunnel setup, PCC BN-en(i) chooses a free label for the Stitching Label SL(i) and adds a new entry in its MPLS L(F)IB with this SL(i) label. Then, it sends a PCRpt message to its PCE(i) with an RRO equal to {[LK(i), SL(i)], RRO(i)} and PLSP-ID(i).
5. Once PCE(i) receives the PCRpt from the PCC BN-en(i) with the RRO and PST = TBD2, it sends to the PCE(i-1) a PCRpt message containing the RRO equal to {[LK(i), SL(i)]} and the PLSP-ID(i). PCE(i) MAY add {PKS(i), ..., PKS(n)} in the RRO.

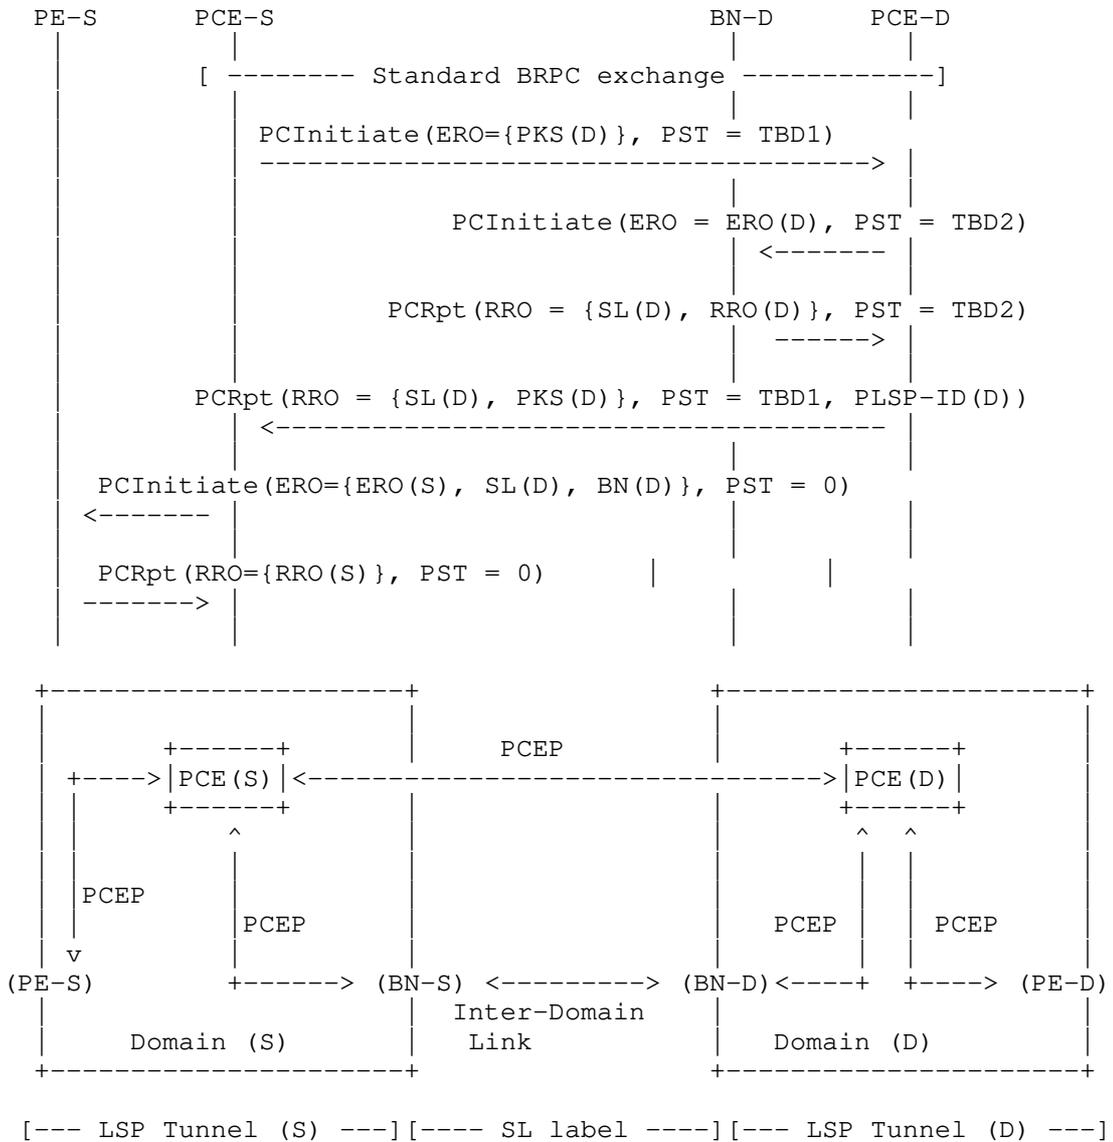
Steps n: Actions performed at the source domain(1) by PCE(1)

Once PCE(1) received the PCRpt message from PCE(2) with the RRO containing the label SL(2), it sends a PCInitiate message to PCC node S with ERO equal to {ERO(1), [LK(2), SL(2)]}, PST = 0 and End-Points Object = {S, BN-ex(1)}. This time, the PST is equal to 0 as the PCC S does not need to return a Stitching Label SL, i.e. it is the head-end of the inter-domain path. Standard PCRpt message is sent back to PCE(1) by the PCC node S.

3.2. Example

In the figure below, two different domains S and D are interconnected through BN respectively BN-S and BN-D. PE-S and PE-D are edge routers. All routers in the figure are connected to their respective PCE through PCEP protocol. In this example, PCE(S) would setup an inter-domain path between PE-S and PE-D acting as source and destination of the tunnel. Intermediate routers between (PE-S, BN-S), (BN-D and PE-D) as well as RSVP-TE messages are not represented to simplify the figure. But they are all presents. The following notation is used in the figure (note that, in this example, we use the PKS for the sake of simplicity):

- o PKS(D) = Path Key corresponding to the path from BN(D) to PE-D
- o ERO(D) = Explicit Route Object corresponding to the path from BN(D) to PE-D retrieves from PKS(D)
- o RRO(D) = Record Route Object of local LSP tunnel(D) from BN(D) to PE-D
- o SL(D) = Stitching Label for local LSP tunnel from BN(D) to PE-D
- o ERO(S) = Explicit Route Object corresponding to the path from PE-S to BN(S)
- o RRO(S) = Record Route Object of local LSP tunnel(S) from PE-S to BN(S)



Example of inter-domain path setup between two domains

3.3. Inter-domain LSP setup procedure completion failure

In case of error during LSP setup, PCRpt and or PCErr messages MUST be used to signal the problem to the neighbor PCE domain backward. In particular, if new PST values defined in this memo are not

supported by the neighbor PCE or the PCC, the PCE, receptively the PCC, MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = 1 (Unsupported path setup type) to its neighbor PCE. If a PCE(i) receives a PCInitiate message from its peer PCE(i-1) without PST set to TBD1 or PST set to a value different from TBD1, it MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = 1 (Unsupported path setup type) to its peer PCE(i-1).

If a PCC or a PCE don't return an RRO or an RRO without the Stitching Label SL with the IP address of the associated link following a PCInitiate message with PST set to TBD1, the PCE MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = TBD5 (No Mandatory Stitching Label is present in the RRO).

In case of completion failure, the PCE(i) MUST propagate the PCErr message up to the PCE(1). In turn, PCE(1) MUST send a PCInitiate message (R flag set in the SRP Object as per draft pce initiated lsp [RFC8281]) to delete this inter-domain path to its neighbor PCEs. PCE(i) MUST propagate the PCInitiate message and remove their local LSP tunnel by means of PCInitiate message to their PCC BN-en(i) and send back PCRpt message to PCE(i-1).

In case of error in domain(i+1), PCE(i) MAY add the AS number of domain(i+1) in the RRO to identify the faulty domain.

4. Hierarchical PCInitiate procedure

This section describes how to setup inter-domain paths than cross several different domains by using a Hierarchical method which is compatible to inter-domain path computation as describe in [RFC6805].

4.1. Mode of operation

This section describes how PCInitiate and PCRpt messages are combined between PCE in order to setup inter-domain paths between a source domain(1) to a destination domain(n). S and D are respectively the source and destination of the inter-domain path. Domain(1) and domain(n) are different and connected through 0 or more intermediate domains denoted domain(i) with $i = (2, n-1)$. Domains are directly connected when $n = 2$.

First, the Parent PCE contacts its Child PCE as per [RFC6805] in order to compute the inter-domain path from S to D, where S and D are respectively a node in the domain(1) and domain(n). Path Key confidentiality as per RFC5520 [RFC5520] SHOULD be used to obfuscate the detailed ERO(i) of the different domains(i). The resulting ERO

is of the form (S, PKS(1), BN-ex(1), ..., BN-en(i), PKS(i), BN-ex(i), ..., BN-en(n), PKS(n), D) when Path Key is used and of the form {S, R(1,1), ..., R(1,k), BN-ex(1), ..., BN-en(i), R(i,1), ..., R(i,l), BN-ex(i), ..., BN-en(n), R(n,1), ..., R(n,m), D} otherwise.

The complete procedure with Path Key follow the different steps described below:

Step 1: Initialization

Parent PCE sends a PCInitiate message to child PCE(n) with an ERO = {PKS(n)} and End-Points = {BN-en(n), D}. Then, PCE(n) retrieves the ERO from the PKS(n) if necessary and sends to BN-en(n) a PCInitiate message with the ERO(n) = {BN-en(n), ..., D}, PST = TBD2 and End-Points Object = {BN-en(n), D} in order to inform the PCC BN-en(n) that this local LSP tunnel(n) is part of an inter-domain path. When the PCC BN-en(n) received the PCInitiate message from its PCE(n), it setup the local LSP tunnel from entry BN-en(n) to D by means of RSVP-TE signaling with the given ERO(n). Once the tunnel setup, it chooses a free label for the Stitching Label SL(n) and add a new entry in its MPLS L(F)IB with this SL(n) label. Then, it sends a PCRpt message to its PCE(n) with an RRO equal to {[LK(n), SL(n)], RRO(n)} and PLSP-ID(n). Once PCE(n) receives the PCRpt from the PCC BN-en(n) with the RRO, PLSP-ID and PST = TBD2, it sends to its Parent PCE a PCRpt containing the RRO equal to {[LK(n), SL(n)]} and PLSP-ID(n). PCE(n) MAY add PKS(n) in the RRO.

Steps i: Actions performed for all intermediate domains(i), for i = n-1 to 2

1. Parent PCE sends a PCInitiate message to Child PCE(i) with PST = TBD1, ERO = {PKS(i), [LK(i+1), SL(i+1)]} and End-Points = {BN-en(i), BN-ex(i)}
2. Then, PCE(i) retrieves the ERO from the PKS(i) if necessary and sends to the PCC BN-en(i) a PCInitiate message with ERO = {ERO(i), [LK(i+1), SL(i+1)]}, PST = TBD2 and End-Points Object = {BN-en(i), BN-ex(i)} in order to inform the PCC BN-en(i) that this local LSP tunnel(i) is part of an inter-domain path.
3. When the PCC BN-en(i) received the PCInitiate message from its PCE(i), it setup the local LSP tunnel from BN-en(i) to BN-ex(i) by means of RSVP-TE signaling with the given ERO(i).
4. Egress Control mechanism, as per RFC4003 section 2.1 [RFC4003], is used to instruct the egress node of domain(i), i.e. BN-ex(i) to forward packets belonging to this tunnel with the Stitching Label. Both Label Stitching and IP address of outgoing interface

are carried in the ERO = {..., [LK(i+1), SL(i+1)]} as the last SubObject in conformance to [RFC4003]. So that, BN-ex(i) installs in its MPLS L(F)IB the SWAP instruction to label SL(i+1) with forward to LK(i+1) instead of the usual POP instruction.

5. Once the tunnel setup, PCC BN-en(i) chooses a free label for the Stitching Label SL(i) and add a new entry in its MPLS L(F)IB with this SL(i) label. Then, it sends a PCRpt message to its PCE(i) with an RRO equal to {[LK(i), SL(i)], RRO(i)} and PLSP-ID(i).
6. Once PCE(i) receives the PCRpt from the PCC BN-en(i) with the RRO and PST = TBD2, it sends to its Parent PCE a PCRpt message containing the RRO equal to {[LK(i), SL(i)]} and the PLSP-ID(i). PCE(i) MAY add PKS(i) in the RRO.
7. Once Parent PCE receives the PCRpt from the Child PCE(i), it stores the corresponding PLSP-ID for this inter-domain tunnel part

Steps n: Actions performed to the source domain(1)

Finally, Parent PCE sends a last PCInitiate message to Child PCE(1) with PST = TBD1, ERO = {PKS(1), [LK(2), SL(2)]} and End-Points = {S, BN-ex(1)}. In turn, Child PCE(1) sends a PCInitiate message to PCC node S with ERO equal to {ERO(1), [LK(2), SL(2)]}, PST = 0 and End-Points Object = {S, BN-ex(1)}. This time, the PST is equal to 0 as the PCC S does not need to return a Stitching Label SL, i.e. it is the head-end of the inter-domain path. Standard PCRpt message is sent back to PCE(1) by the PCC node S. In turn, Child PCE(1) send a final PCRpt message to the Parent PCE with the PSLP-ID(1). PCE(1) MAY adds {S, BN-ex(1)} in the RRO as loose path.

4.2. Inter-domain LSP setup procedure completion failure

In case of error during LSP setup, PCRpt and or PCErr messages MUST be used to signal the problem to the Parent PCE. In particular, if new PST values defined in this memo are not supported by the Child PCE or the PCC, the Child PCE, receptively the PCC, MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = 1 (Unsupported path setup type) to its Parent PCE. If Child PCE(i) receives a PCInitiate message from its Parent PCE without PST set to TBD1 or PST set to a value different from TBD1, it MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = 1 (Unsupported path setup type) to its Parent PCE.

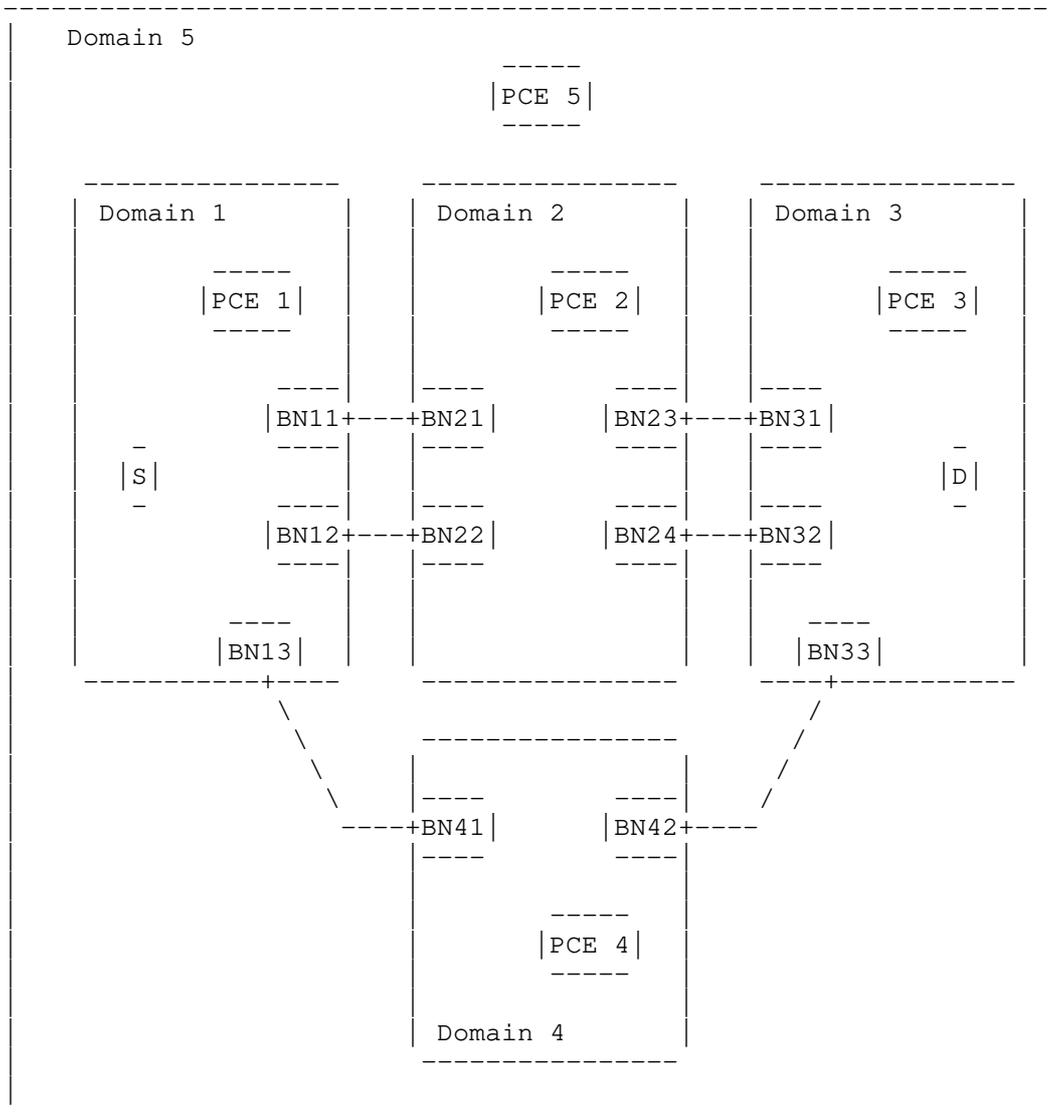
If a Child PCE or a PCC don't return an RRO or an RRO without the Stitching Label SL with the IP address of the associated link

following a PCInitiate message with PST set to TBD1, the Parent PCE, respectively the Child PCE, MUST return a PCErr message with Error-Type = 21 (Traffic engineering path setup error) and Error-Value = TBD5 (No Mandatory Stitching Label is present in the RRO).

In case of completion failure, the Parent PCE MUST send a PCInitiate message (R flag set in the SRP Object as per draft pce initiated lsp [RFC8281]) to delete this inter-domain path to the Child PCEs that already setup their respective part of the inter-domain tunnel. Child PCE(i) MUST remove their local LSP tunnel by means of PCInitiate message with R flag set to 1 to their PCC BN-en(i) and send back PCRpt message to the Parent PCE.

4.3. Example for Stateful H-PCE Sticking procedure

Taking the sample hierarchical domain topology example from [RFC6805] as the reference topology for the entirety of this section.



Hierarchical domain topology from RFC6805

Section 3.3.1 of [I-D.ietf-pce-stateful-hpce] describes the per-domain stitched LSP mode and list all the steps needed. To support SL based stitching, using the reference architecture described in Figure above, the steps are modified as follows (note that we do not use PKS in this example for simplicity):

Step 1: initialization

The P-PCE (PCE5) is requested to initiate a LSP. Steps 4 to 10 of section 4.6.2 of [RFC6805] are executed to determine the end to end path, which are broken into per-domain LSPs e.g. {S-BN41, BN41-BN33, BN33-D}

Step 2: LSP (BN33-D) at PCE3:

1. The P-PCE (PCE5) sends the initiate request to the child PCE (PCE3) via PCInitiate message for LSP (BN33-D) with ERO=(BN33..D) and PST = TBD1
2. The PCE3 further propagates the initiate message to BN33 with the ERO and PST = TBD2/TBD3 based on setup type
3. BN33 initiates the setup of the LSP as per the path and reports to the PCE3 the LSP status ("GOING-UP")
4. The PCE3 further reports the status of the LSP to the P-PCE (PCE5)
5. The node BN33 notifies the LSP state to PCE3 when the state is "UP" it also carry the stitching label (SL33) in RRO as (SL33,BN33..D)
6. The PCE3 further reports the status of the LSP to the P-PCE (PCE5) as well as carry the stitching label (SL33) in RRO as (LK33,SL33,BN33..D)

Step 3: LSP (BN41-BN33) at PCE4

1. The P-PCE (PCE5) sends the initiate request to the child PCE (PCE4) via PCInitiate message for LSP (BN41-BN33) with ERO=(BN41..BN42,LK33,SL33,BN33) and PST = TBD1
2. The PCE4 further propagates the initiate message to BN41 with the ERO and PST = TBD2/TBD3 based on setup type. In case of RSVP_TE, the node BN41 encode the stitching label SL33 as part of the ERO to make sure the node BN42 uses the label SL33 towards node BN33. In case of SR, the label SL33 is part of the label stack pushed at node BN41
3. BN41 initiates the setup of the LSP as per the path and reports to the PCE4 the LSP status ("GOING-UP")
4. The PCE4 further reports the status of the LSP to the P-PCE (PCE5)

5. The node BN41 notifies the LSP state to PCE4 when the state is "UP" it also carry the stitching label (SL41) in RRO as (LK41,SL41,BN41..BN33)
6. The PCE4 further reports the status of the LSP to the P-PCE (PCE5) as well as carry the stitching label (SL41) in RRO as (LK41,SL41,BN41..BN33)

Step 3: LSP (S-BN41) for PCE1

1. The P-PCE (PCE5) sends the initiate request to the child PCE (PCE1) via PCInitiate message for LSP (S-BN41) with ERO=(S..BN13,LK41,SL41,BN41)
2. The PCE1 further propagates the initiate message to node S with the ERO. In case of RSVP_TE, the node S encode the stitching label SL41 as part of the ERO to make sure the node BN13 uses the label SL41 towards node BN41. In case of SR, the label SL41 is part of the label stack pushed at node S
3. S initiates the setup of the LSP as per the path and reports to the PCE1 the LSP status ("GOING-UP")
4. The PCE1 further reports the status of the LSP to the P-PCE (PCE5)
5. The node S notifies the LSP state to PCE1 when the state is "UP"
6. The PCE1 further reports the status of the LSP to the P-PCE (PCE5)

In this way, per-domain LSP are stitched together using the stitching label (SL). The per-domain LSP MUST be setup from the destination domain towards the source domain one after the other.

Once the per-domain LSP is setup, the entry BN chooses a free label for the Stitching Label SL and add a new entry in its MPLS L(F)IB with this SL label. The SL from the destination domain is propagated to adjacent transit domain, towards the source domain at each step. This happens through the entry BN to C-PCE to the P-PCE and vice-versa. In case of RSVP-TE, the entry BN further propagates the SL label to the exit BN via RSVP-TE. In case of SR, the SL label is pushed as part of the SR label stack.

5. Inter-domain LSP Management

This section describe how inter-domain LSPs could be manage.

5.1. Identification of inter-domain tunnels

First, in order to manage inter-domain tunnels composed by the stitching or nesting of local tunnels, it is important to identify them. For this purpose, PLSP-ID managed by PCEs are combined to one provided by PCCs to form global identifier as follow:

- o PCE(i) in the Backward Recursive method or the Child PCE in Hierarchical method MUST create a new unique PLSP-ID for this inter-domain LSP part and MUST send it in the PCRpt message, to the PCE(i-1), respectively the Parent PCE. In addition this new PLSP-ID MUST be associated to the one received from the PCC that instantiate the local tunnel part for further reference.
- o In Hierarchical mode, Parent PCE MUST store and associate the different PLSP-ID(i)s received from the different Child PCE(i)s in order to identify the different part of the inter-domain paths.
- o In Backward Recursive method, PCE(i) MUST store and associate its PLSP-ID(i) and the PLSP-ID(i+1) it received from the PCE(i+1). PCE(n) i.e. the last one in the chain, don't need to perform such association.

Further reference to the inter-domain tunnel will use this PLSP-ID(i). In Backward Recursive method, PCE(i) MUST replace the PLSP-ID(i) by PLSP-ID(i+1) in the PCUpd, PCRpt or PCinitiate message before propagating it to PCE(i+1) and PCE(i) MUST replace the PLSP-ID(i+1) by PLSP-ID(i) in the PCRpt message before propagating it to the PCE(i-1). In Hierarchical method, Parent PCE MUST use the corresponding PLSP-ID(i) of the Child PCE(i).

5.2. Inter-domain association group

In case of failure, a PCE(i) will received PCRpt messages from its PCCs and neighbors PCE(i+1) to synchronize the Inter-domain LSPs. In addition, it may received PCInitiate messages from its previous neighbors PCE(i-1) to re-initiate inter-domain tunnel part. As the PCE(i) may loose the PLSP-ID association, a new association group (within Association Object) is used to ease the association of the different part of the inter-domain tunnel: the local parts and the PCE to PCE parts. The use of the Association Object is MANDATORY in the Backward Recursive method and OPTIONAL in the Hierarchical method.

For that purpose, a new Inter-Domain Association Type with value TBD4 is defined. The first PCE in the Backward Recursive chain (the one which received the initial request) MUST send the PCInitiate message with an Association Object as follows:

- o Association Type field MUST be set to new value TBD4
- o Association ID MUST be set to a unique value. In case of Association ID field is too short or wraps, the first PCE MAY use the Extended Association ID to increase the number of association groups. The Association ID is managed locally by the PCE and does not need to be coordinated with neighbor or remote PCEs.
- o IPV4 or IPv6 association source MUST be set to the IP address which identifies PCE(1) in domain(1).
- o The Global Association Source TLV MUST be present and set with the ASN number of domain(1). It allows to create a globally unique association scope without putting constraint on operator's IP association source. Thus the IP Association Source is associated with the Global Association source to form a unique identifier.
- o Extended Association ID MAY be present and MANDATORY if association ID is too short or wraps.

Subsequent PCE(i) for $i = 2$ to n , MUST send this Association Object as is to the local PCC and the neighbor PCE(i+1).

In case of error with the association group, PCErr message MUST be raised with Error = 26 (Association Error) and Error value set accordingly. A new Error value TBD6 is defined to identify association of inter-domain LSPs.

In Hierarchical method, parent PCE MAY act as initiator of the Association and send to the Child PCEs an Association Object that follows the same rules as for the Backward Recursive method. In turn, Child PCEs MUST propagate the Association Object to the local PCCs as is.

5.3. Inter-domain LSP management

For the Backward Recursive method, each domain manages their respective local LSP tunnel part of an inter-domain path independently of each other. In particular, Stitching Label(i) is managed by domain(i) and is of interest of domain(i-1) only. Thus, Stitching Label SL(i) is not supposed to be propagated to other domains. The same behavior apply to PLSP-ID(i). In Hierarchical method, the Parent PCE MUST ensure the correct distribution of

Stitching Label SL(i) to Child PCE(i-1). The PLSP-ID(i) is kept for the usage of the Parent PCE and thus is not propagated. Only the Association Object defined in section 5.2 is propagated if it is present.

If a PCE(i) needs to modify its local LSP tunnel(i) with a PCUpd message to the PCC BN-en(i), once PCRpt message received by the PCC BN-en(i), it MUST send a new PCRpt message to its neighbor PCE(i-1) in Backward Recursive method, respectively to Parent PCE in Hierarchical method, to advertise PCE(i-1) of the modification. In this case PLSP-ID(i) is used to identify the inter-domain tunnel. PCE(i-1), respectively the Parent PCE, MUST propagate the PCRpt message if the modification implies the previous domain e.g. if the PCRpt indicates that the Stitching Label SL(i) has changed.

PCE(1), respectively Parent PCE, could modify the inter-domain path. For that purpose, it MUST send a PCUpd message to its neighbor PCEs, respectively Child PCE, using the PLSP-ID it received. Each PCE(i) MUST process PCUpd message the same way they process PCInitiate message as defined in section 3.1 for Backward Recursive method and in section 4.1 for Hierarchical method.

In case a failure appears in domain(i), e.g. tunnel becoming down, PCE(i) MUST send a PCRpt message to its neighbor PCE(i-1), respectively its Parent PCE to advertise it of the problem in its local part of the inter-domain path. Once PCE(1), respectively Parent PCE, receives this PCRpt message indicating that the tunnel is down, it is up to the PCE(1), respectively Parent PCE to take appropriate correction e.g. start a new path computation to update the ERO.

5.4. Modification of inter-domain LSP

Modification of local LSP tunnel, BN-en(i) and BN-ex(i) is left for further study.

5.5. Removal of inter-domain LSP

Deletion of inter-domain LSP is only possible by the inter-domain tunnel initiator i.e. PCE(1). For Backward Recursive method, PCInitiate message with R flag set to 1, PLSP-ID set accordingly to section 5.1 and the Association Object with R flag set to 1, is sent by PCE(1) to PCE(n) through PCE(i) and processed the same way as described in section 3.1. For Hierarchical method, PCInitiate message with R flag set to 1 is sent by the Parent PCE to each Child PCE(i) with corresponding PLSP-ID(i) and processed accordingly to section 4.1. Each domain PCE(i) is responsible to delete its part of the tunnel and PCC MUST remove the Stitching label SL in its L(F)IB in addition

to the tunnel when it receives the PCInitiate message with the R flag set to 1 and corresponding PLSP-ID. The Association Group MUST also be removed by the PCC and PCE(i).

6. Applicability

The newly introduce Stitching Label SL serves to stitch or nest part of local LSP tunnels to form an inter-domain path. Each domain is free to decide if the tunnel is stitched or nested and how the tunnel is enforced e.g. tough RSVP-TE or Segment Routing. However, the Stitching Label principle is only compatible with MPLS data plane. At the peering point, the Border Node BN-ex(i) MUST encapsulated the packet with the Stitching Label i.e. the MPLS label prior to send them to the next Border Node BN-en(i+1). Thus, only RSVP-TE and Segment Routing over MPLS technology are detailed in the following sections.

6.1. RSVP-TE

In case of RSVP-TE, the Border Node BN-ex(i) needs to received the Stitching Label through the RSVP-TE message and install in its L(F)IB a SWAP instruction to the Stitching Label and forward it to the next Border Node BN-en(i+1). For that purpose, the Egress Control mechanism, as per RFC4003 section 2.1 [RFC4003], is RECOMMENDED to instruct the Border Node BN-ex(i) of this action. Other mechanisms to program the L(F)IB could be used e.g. NetConf.

As the Stitching Label could serves to stitch or nest tunnels, a domain(i) may decided to nest the incoming local LSP tunnel into a higher hierarchy of tunnel for Traffic Engineering purpose. A PCE(i) may also decided to group local LSP tunnels part of inter-domain paths into a higher hierarchical tunnel to carry all these local LSP tunnels from one BN-en(i) to one BN-ex(i).

6.2. Segment Routing

To use Segment Routing instead of RSVP-TE to setup the local LSP tunnels as defined in draft pce segment routing [I-D.ietf-pce-segment-routing], PCE(i) MUST send PCInitiate message with PST = TBD3 instead of TBD2 to advertise their respective PCC that the local LSP tunnels is enforce by means of Segment Routing.

Stitching Label SL(i) will be inserted in the label stack in order to become the top label in the stack when the packet reach BN-en(i+1): Thus, the Stitching Label SL(i) serves as entry FEC for BN-en(i+1) to identify the packets that follow the next Segment Path. For that purpose, BN-en(i+1) MUST install in its MPLS L(F)IB an instruction to replace the incoming Stitching Label SL(i) by the label stack given

by the ERO(i+1) plus the Stitching Label SL(i+1). When a packet reaches BN-ex(i), the last label in the stack before the label SL(i+1) corresponds to a SID that allows to reach BN-en(i+1).

However, BN-ex(i) needs to know how to send the packets to BN-en(i+1), in particular when there are multiple interfaces between Border Nodes. Similar to the Egress Control mechanism used with RSVP-TE, it is RECOMMENDED to use the inter-domain SID defined as per draft Egress Peer Engineering [I-D.ietf-idr-bgpls-segment-routing-epe] for that purpose. The inter-domain SID is announced by BN-ex(i) to PCE(i) through BGP-LS for each interface that connect BN-ex(i) to neighbors BN-en(i+1). Thus, the label stack will end with {BN-ex(i) SID, Inter-Domain SID, SL(i+1)} and processes as follows:

- o Penultimate router pops its node SID, and sends the packet to the next node designated by the top label in the label stack i.e. the node SID of BN-ex(i)
- o BN-ex(i) pops its node SID and looks up the next label in the stack, i.e. the inter-domain SID which corresponds to the interface to BN-en(i+1). BN-ex(i) pops again this inter-domain SID and send the packet to BN-en(i) through the interface that correspond to the inter-domain SID.
- o BN-en(i+1) pops the Stitching Label SL(i+1) and replaces it by the sub-sequent label stack.

Other mechanisms, e.g. NetConf, could be used to configure the inter-domain SID on exit Border Nodes.

6.3. Mixing technology

During the instantiation procedure, if PCE(i) decides to reuse a local tunnel which is not yet part of an inter-domain tunnel, it SHOULD send a PCUpd message with PST = TBD2 to the PCC BN-en(i) in order to request a Stitching Label SL(i) and new ERO(i) to include the Stitching Label SL(i+1) and the associated link to the previous ERO.

[RFC8453] describes framework for Abstraction and Control of TE Networks (ACTN), where each Physical Network Controller (PNC) is equivalent to C-PCE and P-PCE is the Multi-Domain Service Coordinator (MDSC). The Per domain stitched LSP as per the Hierarchical PCE architecture described in Section 3.3.1 and Section 4.1 of [I-D.ietf-pce-stateful-hpce] is well suited for ACTN. The stitching label (SL) mechanism as described in this document is well suited for ACTN when per domain LSP needs to be stitched to form an E2E tunnel

or a VN Member. It is to be noted that certain VNs require isolation from other clients. The stitching label mechanism described in this document can be applicable to the VN isolation use-case by uniquely identifying the concatenated stitching labels across multi-domain only to a certain VN member or an E2E tunnel.

As each operator is free to enforce the tunnel with its technology choice, it is a local policy decision for PCE(i) to instantiate the local part of the end to end tunnel by either RSVP-TE or Segment Routing. Thus, the PST value (i.e. TBD2 or TBD3) used in the PCinitiate message sends by the PCE(i) to the local PCC is determined by the local policy. How the local policy decision is set in PCE is out of scope of this memo. This flexibility is allowed because the stitching label principle allows to mix (data plane) technologies between domains. For example, a domain(i) could use RSVP-TE while domain(i+1) used Segment Routing, reciprocally. The Stitching Label SL could serve to stitch indifferently Segment Path and RSVP-TE tunnel. Indeed, Stitching Label SL will be part of the label stack in order to become the top label in the stack when reaching the BN-en(i+1). This Stitching Label could be swapped as usual if the next domain uses RSVP-TE tunnel. When the previous domain uses a RSVP-TE tunnel, the Stitching Label will serve as key for the BN-en(i+1) to determine which label stack it must use on top of the packet for a Segment Routing path.

7. IANA Considerations

7.1. Path Setup Type values

[RFC8408] defines the PATH-SETUP-TYPE TLV and requests that IANA creates a registry to manage the value of the PATH_SETUP_TYPE TLV's PST field. IANA is requested to allocate a new code point in the PCEP PATH_SETUP_TYPE TLV PST field registry, as follows:

Value	Description	Reference
TBD1	Inter-Domain Traffic engineering end-to-end path is setup using Backward Recursive method	This Document
TBD2	Inter-Domain Traffic engineering local path is setup using RSVP-TE	This Document
TBD3	Inter-Domain Traffic engineering local path is setup using Segment Routing	This Document

7.2. Association Type value

Draft pce association group [I-D.ietf-pce-association-group] defines the ASSOCIATION Object and requests that IANA creates a registry to manage the value of the Association Type value. IANA is requested to allocate a new code point in the PCEP ASSOCIATION GROUP TLV Association Type field registry, as follows:

Association Type	Description
TBD4	Inter-domain Association Group

7.3. PCEP Error values

IANA is requested to allocate code-points in the PCEP-ERROR Object Error Values registry for a new error-value of Error-Type 21 Invalid traffic engineering path setup and new error-value of Error-Type 26 Association Error:

Error-Type	Error-Value	Description
21	TBD5	Missing Mandatory Stitching Label in RRO
26	TBD6	Error in association of Inter-domain LSPs

8. Security Considerations

No modification of PCE protocol (PCEP) has been requested by this draft which not introduce any issue regarding security. Concerning the PCEP session between PCEs, authors recommend to use the secure version of PCEP as defined in PCEPS [RFC8253] or use any other secure tunnel mechanism e.g. IPsec tunnel to transport PCEP session between PCE.

9. Acknowledgements

The authors want to thanks PCE's WG members, and in particular Dhruv Dhody who greatly contributed to the Hierarchical section of this document.

10. Disclaimer

This work has been performed in the framework of the H2020-ICT-2014 project 5GEx (Grant Agreement no. 671636), which is partially funded by the European Commission. This information reflects the consortium's view, but neither the consortium nor the European Commission are liable for any use that may be done of the information contained therein.

11. References

11.1. Normative References

- [I-D.ietf-pce-association-group]
Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "PCEP Extensions for Establishing Relationships Between Sets of LSPs", draft-ietf-pce-association-group-07 (work in progress), December 2018.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC5441] Vasseur, JP., Ed., 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, DOI 10.17487/RFC5441, April 2009, <<https://www.rfc-editor.org/info/rfc5441>>.
- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.

- [RFC8408] Sivabalan, S., Tantsura, J., Minei, I., Varga, R., and J. Hardwick, "Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages", RFC 8408, DOI 10.17487/RFC8408, July 2018, <<https://www.rfc-editor.org/info/rfc8408>>.

11.2. Informative References

- [I-D.dong-pce-discovery-proto-bgp]
Dong, J., Chen, M., Dhody, D., Tantsura, J., Kumaki, K., and T. Murai, "BGP Extensions for Path Computation Element (PCE) Discovery", draft-dong-pce-discovery-proto-bgp-07 (work in progress), July 2017.
- [I-D.ietf-idr-bgpls-segment-routing-epe]
Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", draft-ietf-idr-bgpls-segment-routing-epe-17 (work in progress), October 2018.
- [I-D.ietf-pce-segment-routing]
Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", draft-ietf-pce-segment-routing-16 (work in progress), March 2019.
- [I-D.ietf-pce-stateful-hpce]
Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., King, D., and O. Dios, "Hierarchical Stateful Path Computation Element (PCE).", draft-ietf-pce-stateful-hpce-06 (work in progress), October 2018.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc3473>>.
- [RFC4003] Berger, L., "GMPLS Signaling Procedure for Egress Control", RFC 4003, DOI 10.17487/RFC4003, February 2005, <<https://www.rfc-editor.org/info/rfc4003>>.

- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, DOI 10.17487/RFC4206, October 2005, <<https://www.rfc-editor.org/info/rfc4206>>.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, DOI 10.17487/RFC5150, February 2008, <<https://www.rfc-editor.org/info/rfc5150>>.
- [RFC5520] Bradford, R., Ed., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, DOI 10.17487/RFC5520, April 2009, <<https://www.rfc-editor.org/info/rfc5520>>.
- [RFC6805] King, D., Ed. and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, DOI 10.17487/RFC6805, November 2012, <<https://www.rfc-editor.org/info/rfc6805>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/info/rfc8253>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.

Authors' Addresses

Olivier Dugeon
Orange Labs
2, Avenue Pierre Marzin
Lannion 22307
France

Email: olivier.dugeon@orange.com

Julien Meuric
Orange Labs
2, Avenue Pierre Marzin
Lannion 22307
France

Email: julien.meuric@orange.com

Young Lee
Huawei Technologies
5340 Legacy Drive, Building 3
Plano TX 75023
USA

Email: leeyoung@huwaei.com

Daniele Ceccarelli
Ericsson
Torshamnsgatan, 48
Stockholm
Sweden

Email: daniele.ceccarelli@ericsson.com