

PCE Working Group
Internet Draft
Category: Standards track

Xian Zhang
Haomian Zheng
Huawei Technologies
Oscar Gonzales de Dios
Victor Lopez
Telefonica I+D

Expires: August 13, 2019

February 13, 2019

**Extensions to Path Computation Element Protocol (PCEP) to Support
Resource Sharing-based Path Computation**

[draft-zhang-pce-resource-sharing-09](#)

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 13, 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.

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](#)].

Abstract

Resource sharing in a network means two or more Label Switched Paths (LSPs) use common piece(s) of resource along their paths. This can help save network resource and is useful in scenarios such as LSP recovery or when two LSPs do not need to be active at the same time. A Path Computation Element (PCE) is responsible for path computation with such requirement. The resource-sharing-based path computation with better efficiency can be achieved together with the association object in PCEP.

This document extends the Path Computation Element Protocol (PCEP) in order to support resource sharing-based path computation, which is a special case in the association path computation.

Table of Contents

| | | |
|----------------------|--|--------------------|
| 1. | Introduction and Motivation | 3 |
| 2. | Motivation | 4 |
| 2.1. | Single Domain Use Case | 4 |
| 2.2. | Multiple Domains Use Case | 7 |
| 2.3. | Bulk Path Computation Use Case | 8 |
| 3. | Extensions to PCEP | 10 |
| 3.1. | Association group and type | 10 |
| 3.2. | Resource Sharing TLV | 10 |
| 3.3. | Processing Rules | 11 |
| 4. | Security Considerations | 12 |
| 5. | IANA Considerations | 13 |
| 5.1. | Association Object Type Indicators | 13 |
| 6. | References | 14 |

| | |
|---|--------------------|
| 6.1. Normative References | 14 |
| 6.2. Informative References | 14 |
| 7. Authors' Addresses | 15 |

[1. Introduction and Motivation](#)

A Path Computation Element (PCE) provides an alternative way for providing path computation function, and it is especially useful in the scenarios where complex constraints and/or a demanding amount of computation resource are required [[RFC4655](#)]. The development of PCE standardization has evolved from stateless to stateful. A stateful PCE has access to the LSP database information of the network(s) it serves as a computation engine [[RFC8231](#)]. Unless specified, this document assumes a PCE mentioned is a stateful PCE (either passive or active).

Resource sharing denotes that two or more Label Switched Paths (LSPs) share common piece(s) of resource, (such as a common time slot of a link in an Optical Transport Network (OTN)). This is usually useful in the scenario where only one LSP is active and the benefit herein is to save network resources. A simple example of this is dynamically calculating a LSP for an existing LSP undergoing a link failure. Note that the resource sharing can be worked out using a stateless PCE, but the mechanism may be complex and is out the scope of this draft.

This document considers the following requirement: new LSP may request for resource sharing with one or multiple existing LSPs. Furthermore, if there is resource sharing between new LSP and existing LSP, the two LSPs cannot exist simultaneously, the new LSP will replace the existing LSP(s).

In a single domain, this is a common requirement in the recovery cases especially in order to increase traffic resilience against failure while reducing the amount of network resource used for recovery purpose [[RFC4428](#)].

The current protocol supporting the communication between a PCE and a Path Computation Client (PCC), i.e. PCE Protocol (PCEP), allows for re-optimization of an existing LSP [[RFC5440](#)]. This is achieved by setting R bit in the Request Parameter (RP) object, together with some additional information if applicable, in the Path Computation Request (PCReq) message sent from a PCC to the PCE. To support this type of resource sharing, a PCC needs to ask a PCE to compute a new path with the constraints of sharing resource with one or multiple existing LSPs. It is worth noting the "resource sharing" in this

draft not only means one LSP re-using the same link(s) of another LSP, but also the same slice of bandwidth in TDM networks. This may occur when an LSP is required for re-routing, or online re-optimization. Current PCEP specifications do not provide such function. More specifically, this draft describes the resource sharing issue during the procedure when a new LSP is required to replace an existing LSP, which can be used together with Make-before-break (MBB) described in [[RFC3209](#)].

There are a few objects which indicate the resource sharing/disjoint relationships, such as SRLG and ASSOCIATE. However, these objects are used to describe the relationship with two simultaneous LSPs, instead of a new one and an old one, which is different with the object proposed in this draft.

As mentioned in [[RFC8231](#)], the PLSP-ID is unique during a PCEP session between PCC and PCE. Such identification is helpful in supporting the above resource sharing requirement for standardization of stateful PCEs. With a unique identifier, the configuration of PCCs is greatly simplified. Instead of determining all the resources to be shared, the PCC could request resource sharing directly from PCE.

The resource sharing can also be required in an inter-layer PCEP session. This is similar to the previous requirement. However, it is more complex and therefore deserves a more detailed explanation here.

In a multi-layer network, Label Switched Paths (LSPs) in a lower layer are used to carry higher-layer LSPs across the lower-layer network [[RFC5623](#)]. Therefore, the resource sharing constraints in the higher layer might actually relate to the resource sharing in the lower layer. Thus, it is useful to consider how this can be achieved and whether additional extensions are needed using the models defined in [[RFC5623](#)].

In the next sections, use cases are provided to show what information needs to be exchanged to fulfill these requirements. This memo then provides extensions to PCEP to enable this function.

[2. Motivation](#)

[2.1. Single Domain Use Case](#)

Figure 1 shows a single domain network with a stateful PCE. Assume a working LSP (N1-N2-N3) exists in the network, when there is failure on the link N2-N3, it is desired to set up a restoration path for

this working LSP. Suppose N1 serves as the PCC and sends a request to the stateful PCE for such an LSP. Before sending the request, N1 may need to check what policy should be applied for the path re-computation. For example, it might value resource sharing and prefer to share as much resource with the working LSP as possible and specify this policy in the PCReq message. If resources are shared between the old and new LSPs, there will be some 'interruption' when the traffic is switched from the old LSP to the new LSP. Here the resources to be shared mean the LSP information, which includes the node, link and corresponding SRLG information, etc.

On the other hand, in some scenarios there are different policies, for example the LSP should be restored without any interruption with best effort. An example can be found in Fig. 1 without failure on N2-N3 link, instead, an online re-optimization is needed for the working LSP (N1-N2-N3) from the stateful PCE. In such cases, the best choice is to set up a backup LSP for the working LSP with totally separate routing (for example N1-N5-N4-N3), and move the traffic to that backup LSP. After that the working LSP can be torn down, which will not result in any interruption during the optimization procedure. This can actually be implemented with existing PCEP mechanism. However, if there is no such separate path, existing PCEP will reply error. A secondary option for this case is to set up an LSP and complete such re-optimization with resource sharing, even if some interruption introduced. Given the resource from the LSP to be interrupted, there may be some solutions instead of Path Compute error due to the lack of resource.

A simple illustration is provided in Figure 1:

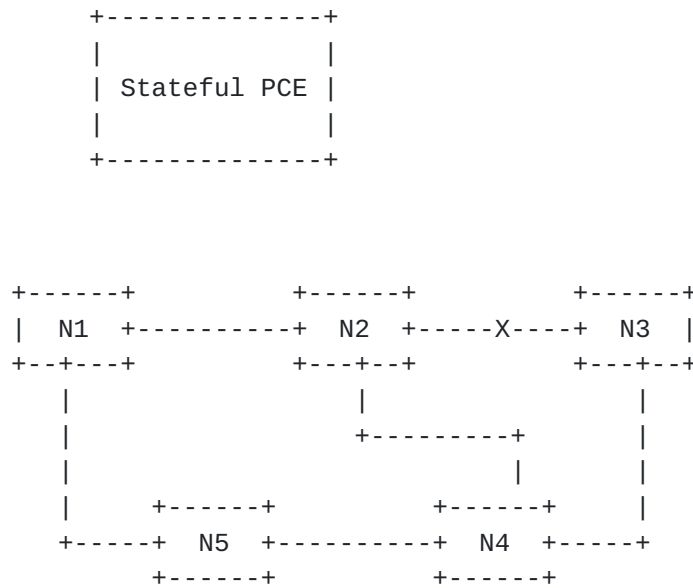


Figure 1: A Single Domain Example

Available recovery paths computed by the stateful PCE:

LSP1: N1-N2-N4-N3

LSP2: N1-N5-N4-N3

If resource sharing is preferred, the stateful PCE will reply with LSP1 information. Instead, if PCC prefer to have less interruption, PCE will reply with LSP2 information.

Another piece of information that needs to be conveyed to the PCE is the information about the working path LSP. Note this simple use case assumes end-to-end recovery. But in order to be applicable to use cases such as shared mesh protection purpose, where the head-end or tail-end nodes may be different, this information is necessary in the message exchange between PCCs and PCEs, so that the stateful PCE knows which LSP the path computation request wants to share the resource.

Besides, parameter changes during the resource sharing computation also need to be considered. For example, the bandwidth of the request LSP may be different with the existing LSP, while resource sharing is still preferred by the PCC. PCE should consider the sharing request together with the policy and available resource(s) in the network. Details can be found in [Section 3.3](#).

2.2. Multiple Domains Use Case

Figure 2 shows a two-layer network example, with each layer managed by a PCE. As Discussed in [Section 3 of \[RFC5623\]](#), there are three models for inter-layer path computation. They are single PCE computation, multiple PCE with inter-PCE communication and multiple PCE without inter-PCE communication, respectively. For the single PCE computation, the process would be similar to that of the use case in [Section 2.1](#). Thus, this model is not discussed further.




Figure 2: A Two-layer Network Example

An inter-layer path computation example is shown in Fig. 2, assume a LSP (LSP1: H2-H3) has been established already, visible as H2-H3 from view of higher-layer PCE and H2-L1-L2-H3 from the global view (or from the view of lower-layer PCE). A new request comes at H2 to establish a new LSP (LSP2: from H2 to H5), given the constraint it can share resource with LSP1. This requirement is possible if only one of the LSPs needs to be active and resource sharing is the target.

If multiple PCE with inter-PCE communication model is employed, the path computation request sent by H2 to higher-layer PCE will be

forwarded to lower-layer PCE since there is no resource readily available in the higher layer. So it leaves the lower-layer PCE to compute a path in the lower layer in order to support the higher layer request. In this case, lower-layer PCE is required to compute a path between H2 and H5 under the constraint that it can share the resource with that of the LSP1. At this moment the lower-layer PCE has the knowledge on the explicit routing that LSP1 go through (H2-L1-L2-H3), and therefore can map the lower layer LSP with the higher-layer one. So when lower-layer PCE computes the path for LSP2, it can consider the resource used by LSP1 as available with higher priority. For example, lower-layer PCE may choose H2-L1-L2-L4-H5 as the computation result. On the other hand, if the path computation policy is to have a separate path with LSP1, the lower-layer PCE may choose H2-L1-L3-L4-H5.

During this procedure higher-layer PCE can only use LSP1 information (such as its five-tuple LSP information) as the information, an issue to solve is how lower-layer PCE can resolve this information to the actual resource usage in its own layer, i.e. lower layer. This could be solved by edge LSR L1 reporting this higher-lower layer LSP correlation to the lower-layer PCE as part of the LSP information during the LSP state synchronization process. If needed, it can be later updated when there is a change in this information. Alternatively, the lower-layer PCE can get this information from other sources, such as network management system, where this information should be stored.

If multiple PCE without inter-PCE communication model is employed, the path computation request in the lower layer will be initiated the border LSR node, i.e., L1. The process would be similar to that of the previous scenario. A point worth noting is that the border LSR node may be able to resolve the higher layer LSP information itself, such as mapping it to the corresponding LSP in the lower layer, in this way lower-layer PCE does not need to perform this function. Otherwise, the mapping method mentioned above can still be used.

2.3. Bulk Path Computation Use Case

There is a potential need of resource sharing during the bulk path computation, especially the processing of the "sticky resources" in [RFC7399]. It would be useful to specify the resources that could be shared among different paths, i.e., the bandwidth information. Considering the H-PCE architecture in [H-PCE], when the parent PCE may ask for a single path across a few domains, such request may become a bulk path computation to a certain child PCE. Figure 3

shows and example of 3 domains. The parent PCE will select one of these path for establishment.




Figure 3: Bulk Request example with Hierarchical PCE

A 3-domain example is shown in Figure 3, with hierarchical PCE architecture. In this example node A/B/C belong to domain 1, node D/E/G/H belong to domain 2 and node J/K/L belong to domain 3. Inter-domain links are B-D/C-G between domain 1 and 2, and H-J/H-K between domain 2 and 3. Given a path computation request from A to L, a bulk request from P-PCE would be helpful to understand whether it is possible to have different combination on the inter-domain links. However, the resources on some specific links become 'sticky' and have to be indicated as 'sharing allowed' to avoid unnecessary resource competition. For example, both the route A-B-D-E-H-J-L and A-C-G-E-H-K-L are qualified, but these routes are competing the resource on the link E-H and cannot be established simultaneously, so there must be one route failed to be reported P-PCE. Given the indication of allowing the share on the link E-H, both of these route can be reported for P-PCE's decision, and there will not be any competition as the P-PCE understand only one path needs to be set up.

3. Extensions to PCEP

3.1. Association group and type

According to the definition in [[ietf-pce-association-group](#)], the association group is used to associate multiple LSPs into one group for further path computation considerations, such as disjointness and resource sharing. An association ID will be used to identify the resource sharing group. An association type that described disjointness has been defined in [[ietf-pce-association-diversity](#)]. In this draft, a new association type is defined as:

Association type = TBD1 ("Sharing Association Type").

A sharing group should have multiple LSPs. The number of LSPs and the criteria for how LSPs share among each other are implementation dependent. Local path computation policies apply to different PCE and PCC, some examples can be found in [section 2](#).

3.2. Resource Sharing TLV

The PCEP Resource Sharing group MUST carry the following TLV. It MAY be carried within a PCReq message from the network element (or other PCCs) so as to indicate the desired resource sharing requirements to be applied by the stateful PCE during path computation.

```

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

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type = [TBD2]                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Flags                               |B|S|N|L|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Optional TLVs                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Currently the following flags have been defined:

- * L (Link share) bit: when set, this flag indicates that the PCE should prioritize the links that shared by existing LSPs within the sharing group for path computation.
- * N (Node share) bit: when set, this flag indicates that the PCE should prioritize the nodes that shared by existing LSPs within the sharing group for path computation.
- * S (SRLG share) bit: when set, this flag indicates that the PCE should set the SRLG (Shared Risk Link Group) of the computed LSP to the same as existing LSPs within the sharing group for path computation.
- * B (Bandwidth share) bit: when set, this flag indicates that the PCE should prioritize the bandwidth to be shared by LSPs request within the sharing group for bulk path computation.

Optional TLVs may be needed to indicate the LSP(s) with which the resource is shared. If multiple LSPs are required, the PCE may need to consider different sharing policies, which is implementation dependent and may result in a different computing result. The selection policy among multiple computation result is out of the scope of this draft.

3.3. Processing Rules

To request a path allowing sharing resource with one or multiple existing LSPs, a PCC includes a Resource Sharing TLV in the association group object in any kind of path computation request message, such as the PCReq, PCUpd or PCInitiate messages specified in [[RFC8231](#)] and [[RFC8281](#)].

On receipt of a PCEP message with a Resource Sharing TLV, a stateful PCE MUST proceed as follows:

- If the Resource Sharing TLV is unknown/unsupported, the PCE will follow procedures defined in [[RFC5440](#)]. That is, the PCE sends a PCErr message with error type 3 or 4 (Unknown / Not supported object) and error value 1 or 2 (unknown / unsupported object class / object type), and the related path computation request is discarded.
- If Resource Sharing TLV are unknown/unsupported and the P bit is set, the PCE MUST send a PCErr message with error type 3 or 4 (Unknown / Not supported object) and error value 4 (Unrecognized/Unsupported parameter), and the related path computation request MUST be discarded as defined in [[RFC5440](#)].

- If the resource sharing TLV is extracted correctly, the PCE MUST apply the requested resource sharing requirement.

The procedure of setting flags follows the rules defined in [Section 3.1](#). The RSO flags may be locally configured on the requesting nodes via external entities, such as a network management system or the entity that impose the resource sharing requirement.

It is worth noting that the Resource Sharing TLV can be used together with other path indication objects like IRO/XRO, with difference objectives. The first difference is, the use of Resource Sharing TLV is to setup an alternative path, instead a new path. It is also dependent on the knowledge of PCC, e.g., if the PCC have a full knowledge of the path information and have strong preference on the route, it may send the request message with IRO message to specify the route. On the other hand, if the PCC does not know how the path should go but just want to set up a new LSP to replace the old one, it may use the Resource Sharing TLV instead of IRO. The second difference is, resource Sharing TLV is a loose requirement. For example, if the constraint specified in IRO/XRO in an A-Z path computation request cannot be satisfied, the reply message from PCE to PCC would be unsuccessful. However it is still possible to have a path from the A-Z. If the target node/link/SRLG/Bandwidth is set in Resource Sharing TLV rather than IRO, the PCE may feedback a path that from A-Z that not sharing the target specified in Resource Sharing TLV.

4. Security Considerations

Security of PCEP is discussed in [\[RFC5440\]](#) and [\[RFC6952\]](#). The extensions in this document do not change the fundamentals of security for PCEP.

However, the introduction of the Resource Sharing TLV in association group object provides a vector that may be used to probe for information from a network. For example, a PCC that wants to discover the path of an LSP with which it is not involved can issue a request message with a Resource sharing TLV and may be able to get back quite a lot of information about the path of the LSP through issuing multiple such requests for different endpoints and analyzing the received results. To protect against this, a PCE should be configured with access and authorization controls such that only authorized PCCs (for example, those within the network) can make computation requests, only specifically authorized PCCs can make requests for resource sharing, and such requests relating to specific LSPs are further limited to a select few PCCs. How such access controls and authorization is managed is outside the scope of

this document, but it will at the least include Access Control Lists.

Furthermore, a PCC must be aware that setting up an LSP that share resources with another LSP may be a way of attacking the other LSP, for example by depriving it of the resources it needs to operate correctly. Thus it is important that, both in PCEP and the associated signaling protocols, only authorized resource sharing is allowed.

5. IANA Considerations

5.1. Association Object Type Indicators

This document defines a new association type, with the following information:

| Object Class | Name | Object Type | Reference |
|--------------|---------------|------------------|-----------------|
| ----- | | | |
| TBA1 | Sharing-group | Association Type | [this document] |

5.2 PCEP TLV Definitions

This document defines the following TLVs to support the resource sharing scenario:

| Value | Name | Reference |
|-------|----------------------|-----------------|
| ----- | | |
| TBA2 | Resource-sharing TLV | [this document] |

IANA is requested to allocate the following bit numbers in the flag spaces of Resource-sharing TLV:

| Bit | Flag name | Reference |
|-----|-----------------|-----------------|
| 0 | Link Share | [this document] |
| 1 | Node Share | [this document] |
| 2 | SRLG Share | [this document] |
| 3 | Bandwidth Share | [this document] |

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", [RFC 2119](#), March 1997.
- [RFC4655] Farrel, A., Vasseur, J.-P., and Ash, J., "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC5440] Vasseur, J.-P., and Le Roux, J.L., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC8231] Crabbe, E., Medved, J., Minei, I., and R. Varga, "PCEP Extensions for Stateful PCE", [RFC8231](#), June 2017.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", [RFC 8281](#), October 2017.
- [ietf-pce-association-group] Minei, I., Crabbe E., Sivabalan S., Ananthakrishnan H., Dhody D., Tanaka Y., "PCEP Extensions for Establishing Relationships Between Sets of LSPs", work in Progress.
- [ietf-pce-association-diversity] Litkowski, S., Sivabalan, S., Barth, C., Dhody, D., "Path Computation Element communication Protocol extension for signaling LSP diversity constraint", Work in Progress.

6.2. Informative References

- [RFC4428] Papadimitriou, D., Mannie., E., "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", [RFC4428](#), March 2006.
- [RFC5623] Oki., E., Takeda, T., Le Roux, J.L., Farrel, A., "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", [RFC5623](#), September 2009.
- [RFC6952] Jethanandani, M., Patel, K., Zheng, L., "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", [RFC6952](#), May 2013.

[H-PCE] Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., King, D., Gonzalez de Dios, O., "Hierarchical Stateful Path Computation Element (PCE)", [draft-ietf-pce-stateful-hpce-06](#), October 2018.

7. Authors' Addresses

Xian Zhang
Huawei Technologies

Email: zhang.xian@huawei.com

Haomian Zheng
Huawei Technologies

Email: zhenghaomian@huawei.com

Oscar Gonzalez de Dios
Telefonica I+D/gCTIO
Distrito Telefonica
E-28050 Madrid, Spain
EMail: oscar.gonzalezdedios@telefonica.com

Victor Lopez
Telefonica I+D/gCTIO
Distrito Telefonica
E-28050 Madrid, Spain
EMail: victor.lopezalvarez@telefonica.com

Contributor's Address :

Dhruv Dhody
Huawei Technologies
Email: dhruv.dhody@huawei.com

Igor Bryskin
Huawei Technologies
Email: Igor.Bryskin@huawei.com