PCE Working Group
Internet-Draft
Intended status: Standards Track

Young Lee (Editor)
Fatai Zhang
Huawei
Ramon Casellas
CTTC
Oscar Gonzalez de Dios
Telefonica I+D
Zafar Ali
Cisco Systems

March 7, 2019

Path Computation Element (PCE) Protocol Extensions for Stateful PCE Usage in GMPLS-controlled Networks

draft-ietf-pce-pcep-stateful-pce-gmpls-10

Abstract

The Path Computation Element (PCE) facilitates Traffic Engineering (TE) based path calculation in large, multi-domain, multi-region, or multi-layer networks. The PCE communication Protocol (PCEP) has been extended to support stateful PCE functions where the PCE retains information about the paths already present in the network, but those extensions are technology-agnostic. This memo provides extensions required for PCEP so as to enable the usage of a stateful PCE capability in GMPLS-controlled networks.

Table of Contents

1. Introduction

[RFC4655] presents the architecture of a Path Computation Element (PCE)-based model for computing Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering Label Switched Paths (TE LSPs). To perform such a constrained computation, a PCE stores the network topology (i.e., TE links and nodes) and resource information (i.e., TE attributes) in its TE Database (TED). Such a PCE is usually referred as a stateless PCE. To request path computation services to a PCE, [RFC5440] defines the PCE communication Protocol (PCEP) for interaction between a Path Computation Client (PCC) and a PCE, or between two PCEs. PCEP as specified in [RFC 5440] mainly focuses on MPLS networks and the PCEP extensions needed for GMPLS-controlled networks are provided in [PCEP-GMPLS].

Stateful PCEs are shown to be helpful in many application scenarios, in both MPLS and GMPLS networks, as illustrated in [RFC8051]. Further discussion of concept of a stateful PCE can be found in [RFC7399]. In order for these applications to able to exploit the capability of stateful PCEs, extensions to PCEP are required.

[RFC8051] describes how a stateful PCE can be applicable to solve various problems for MPLS-TE and GMPLS networks and the benefits it brings to such deployments.

[RFC8231] provides the fundamental extensions needed for stateful PCE to support general functionality, but leaves out the specification for technology-specific objects/TLVs. This document focuses on the extensions that are necessary in order for the deployment of stateful PCEs in GMPLS-controlled networks.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in

BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Context of Stateful PCE and PCEP for GMPLS

This document is built on the basis of Stateful PCE [RFC8231] and PCEP for GMPLS [PCEP-GMPLS].

There are two types of LSP operation for Stateful PCE.

For Active Stateful PCE, PCUpd message is sent from PCE to PCC to update the LSP state for the LSP delegated to PCE. Any changes to the delegated LSPs generate a PCRpt message by the PCC to PCE to convey the changes of the LSP. Any modifications to the Objects/TLVs that are identified in this document to support GMPLS technology-specific attributes will be carried in the PCRpt and PCUpd messages.

For Passive Stateful PCEs, PCReq/PCRep messages are used to convey path computation instructions. GMPLS-technology specific Objects and TLVs are defined in [PCEP-GMPLS], so this document just points at that work and only adds the stateful PCE aspects where applicable. Passive Stateful PCE makes use of PCRpt messages when reporting LSP State changes sent by PCC to PCEs. Any modifications to the Objects/TLVs that are identified in this document to support GMPLS technology-specific attributes will be carried in the PCRpt message.

[PCEP-GMPLS] defines GMPLS-technology specific Objects/TLVs and this document makes use of these Objects/TLVs without modifications where applicable. Some of these Objects/TLVs may require modifications to incorporate stateful PCE element where applicable.

4. Main Requirements

This section notes the main functional requirements for PCEP extensions to support stateful PCE for use in GMPLS-controlled networks, based on the description in [RFC8051]. Many requirements are common across a variety of network types (e.g., MPLS-TE networks and GMPLS networks) and the protocol extensions to meet the requirements are already described in [RFC8231]. This document does not repeat the description of those protocol extensions. This document presents protocol extensions for a set of requirements which are specific to the use of a stateful PCE in a GMPLS-controlled network.

The basic requirements are as follows:

- o Advertisement of the stateful PCE capability. This generic requirement is covered in Section 5.4. of [RFC8231]. This document assumes that STATEFUL-PCE-CAPABILITY TLV can be used for GMPLS Stateful PCE capability and therefore does not provide any further extensions.

- LSP delegation is already covered in Section 5.7. of [RFC8231]. Section 2.2. of this document does not provide any further extensions.

- Active LSP update is covered in Section 6.2 of [RFC8231]. Section 4.1. of this document provides extension for its application in GMPLS-controlled networks.

- LSP state synchronization and LSP state report. This is a generic requirement already covered in Section 5.6. of [RFC8231]. However, there are further extensions required specifically for GMPLS-controlled networks and discussed in Section 4.2.

5. PCEP Extensions

5.1. LSP Update in GMPLS-controlled Networks

[RFC8231] defines the Path Computation LSP Update Request (PCUpd) message to enable to update the attributes of an LSP. However, that document does not define technology-specific parameters.

A key element of the PCUpd message is the attribute-list construct defined in [RFC5440] and extended by many other PCEP specifications.

For GMPLS purposes we note that the BANDWIDTH object used in the attribute-list is defined in [PCEP-GMPLS]. Furthermore, additional TLVs are defined for the LSPA object in [PCEP-GMPLS] and MAY be included to indicate technology-specific attributes. There are other technology-specific attributes that need to be conveyed in the <intended-attribute-list> of the <path> construct in the PCUpd message. Note that these path details in the PCUpd message are the same as the <attribute-list> of the PCRep message. See Section 4.2 for the details.

5.2. LSP Synchronization in GMPLS-controlled Networks

PCCs need to report the attributes of LSPs to the PCE to enable stateful operation of a GMPLS network. This process is known as LSP state synchronization. The LSP attributes include bandwidth, associated route, and protection information etc., are stored by the PCE in the LSP database (LSP-DB). Note that, as described in [RFC8231], the LSP state synchronization covers both the bulk reporting of LSPs at initialization as well the reporting of new or modified LSP during normal operation. Incremental LSP-DB synchronization may be desired in a GMPLS-controlled network and it is specified in [RFC8232].

[RFC8231] describes mechanisms for LSP synchronization using the Path Computation State Report (PCRpt) message, but does not cover reporting of technology-specific attributes. As stated in [RFC8231], the <path> construct is further composed of a compulsory Explicit Route Object (ERO) and a compulsory attribute-list and an optional Record Route Object (RRO). In order to report LSP states in GMPLS networks, this specification allows the use within a PCRpt message both of technology- and GMPLS-specific attribute objects and TLVs defined in [PCEP-GMPLS] as follows:

- o Include Route Object (IRO)/ Exclude Route Object (XRO) Extensions to support the inclusion/exclusion of labels and label sub-objects for GMPLS. (See Section 2.6 and 2.7 in [PCEP-GMPLS])

- o END-POINTS (Generalized END-POINTS Object Type. See Section 2.5 in [PCEP-GMPLS])

- o BANDWIDTH (Generalized BANDWIDTH Object Type. See Section 2.3 in [PCEP-GMPLS])

- o LSPA (PROTECTION ATTRIBUTE TLV, See Section 2.8 in [PCEP-GMPLS].

The END-POINTS object SHOULD be carried within the attribute-list to specify the endpoints pertaining to the reported LSP. The XRO object MAY be carried to specify the network resources that the reported LSP avoids and a PCE SHOULD consider avoid these network resources during the process of re-optimizing after this LSP is delegated to the PCE. To be more specific, the <attribute-list> is updated as follows using the notations of [RFC5511]:

```
<attribute-list> ::= [<END-POINTS>]
                     [<LSPA>]
                     [<BANDWIDTH>]
                     [<metric-list>]
                     [<IRO>]
                     [<XRO>]

<metric-list>::= <METRIC>[<metric-list>]
```

If the LSP being reported protects another LSP, the PROTECTION-ATTRIBUTE TLV [PCEP-GMPLS] MUST be included in the LSPA object to describe its attributes and restrictions. Moreover, if the status of the protecting LSP changes from non-operational to operational, the PCC SHOULD synchronize the state change of the LSPs to the stateful PCE using a PCRpt message. This use case arises, for example, when the protecting LSP becomes operational due to the failure of the primary LSP.

5.3. Modification of Existing PCEP Messages and Procedures

One of the advantages mentioned in [RFC8051] is that the stateful nature of a PCE simplifies the information conveyed in PCEP messages, notably between PCC and PCE, since it is possible to refer to PCE managed state for active LSPs. To be more specific, with a stateful PCE, it is possible to refer to an LSP with a unique identifier in the scope of the PCC-PCE session and thus use such identifier to refer to that LSP. Note this is also applicable to packet networks.

5.3.1. Modification for LSP Re-optimization

The Request Parameters (RP) object on a Path Computation Request (PCReq) message carries the R bit. When set, this indicates that the PCC is requesting re-optimization of an existing LSP. Upon receiving such a PCReq, a stateful PCE SHOULD perform the re-optimization in the following cases:

- o The existing bandwidth and route information of the LSP to be re-optimized is provided in the PCReq message using the BANDWIDTH object and the ERO.

- o The existing bandwidth and route information is not supplied in the PCReq message, but can be found in the PCE’s LSP-DB. In this case, the LSP MUST be identified using an LSP identifier carried in the PCReq message, and that fact requires that the LSP identifier was previously supplied either by the PCC in a PCRpt message or by the PCE in a PCRep message. [RFC8231] defines how this is achieved using a combination of the per-node LSP identifier (PLSP-ID) and the PCC’s address.

If no LSP state information is available to carry out re-optimization, the stateful PCE should report the error "LSP state information unavailable for the LSP re-optimization" (Error Type = TBD1, Error value= TBD2).

5.3.2. Modification for Route Exclusion

[RFC5521] defines a mechanism for a PCC to request or demand that specific nodes, links, or other network resources are excluded from paths computed by a PCE. A PCC may wish to request the computation of a path that avoids all link and nodes traversed by some other LSP.

To this end this document defines a new sub-object for use with route exclusion defined in [RFC5521]. The LSP exclusion sub-object is 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 (TBD3)  |    Length     |  Attributes   |     Flag      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//                    Symbolic Path Name                       //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

X bit and Attribute fields are defined in [RFC5521].

Type: Subobject Type for an LSP exclusion sub-object. Value of TBD3. To be assigned by IANA.

Length: The Length contains the total length of the subobject in bytes, including the Type and Length fields.

Flags: This field may be used to further specify the exclusion constraint with regard to the LSP. Currently, no values are defined.

Symbolic Path Name: This is the identifier given to an LSP and is unique in the context of the PCC address as defined in [RFC8231].

Reserved: MUST be transmitted as zero and SHOULD be ignored on receipt.

This sub-object is OPTIONAL in the exclude route object (XRO) and can be present multiple times. When a stateful PCE receives a PCReq message carrying this sub-object, it SHOULD search for the identified LSP in its LSP-DB and then exclude from the new path computation all resources used by the identified LSP. If the stateful PCE cannot recognize one or more of the received LSP identifiers, it should send an error message PCErr reporting "The LSP state information for route exclusion purpose cannot be found"

(Error-type = TBD1, Error-value = TBD4). Optionally, it may provide with the unrecognized identifier information to the requesting PCC using the error reporting techniques described in [RFC5440].

5.3.3. Modification for SRP Object to indicate Bi-directional LSP

The format of the SRP object is defined in [RFC8231]. The object is used in PCUpd and PCInit messages for GMPLS.

This document defines a new flag to be carried in the Flags field of the SRP object. This flag indicates a bidirectional co-routed LSP setup operation initiated by the PCE as follows:

- o B (Bidirectional LSP -- 1 bit): If set to 0, it indicates a request to create a uni-directional LSP. If set to 1, it indicates a request to create a bidirectional co-routed LSP.

The bit position is TBD5 as assigned by IANA (see Section 5.3)

5.4. Object Encoding

Note that, as is stated in Section 7 of [RFC8231], the P flag and the I flag of the PCEP objects used on PCUpd and PCRpt messages SHOULD be set to 0 on transmission and SHOULD be ignored on receipt since these flags are exclusively related to path computation requests.

6. IANA Considerations

6.1. New PCEP Error Codes

IANA is requested to make the following allocation in the "PCEP-ERROR Object Error Types and Values" registry.

Error Type	Meaning	Reference
TBD1	LSP state information missing	[This.I-D]
Error-value TBD2:	LSP state information unavailable for the LSP re-optimization	[This.I-D]
Error-value TBD4:	LSP state information for route exclusion purpose cannot be found	[This.I-D]

6.2. New Subobject for the Exclude Route Object

IANA maintains the "PCEP Parameters" registry containing a subregistry called "PCEP Objects". This registry has a subregistry for the XRO (Exclude Route Object) listing the sub-objects that can be carried in the XRO. IANA is requested to assign a further sub-object that can be carried in the XRO as follows:

Value	Description	Reference
TBD3	LSP identifier sub-object	[This.I-D]

6.3. New "B" Flag in the SRP Object

IANA maintains a subregistry, named the "SRP Object Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry, to manage the Flag field of the SRP object.

IANA is requested to make an assignment from this registry as follows:

Bit	Description	Reference
TDB5	Bi-directional co-routed LSP	[This.I-D]

7. Manageability Considerations

The description and functionality specifications presented related to stateful PCEs should also comply with the manageability specifications covered in Section 8 of [RFC4655]. Furthermore, a further list of manageability issues presented in [RFC8231] should also be considered.

Additional considerations are presented in the next section.

7.1. Requirements on Other Protocols and Functional Components

When the detailed route information is included for LSP state synchronization (either at the initial stage or during LSP state report process), this requires the ingress node of an LSP carry the RRO object in order to enable the collection of such information.

8. Security Considerations

This draft provides additional extensions to PCEP so as to facilitate stateful PCE usage in GMPLS-controlled networks, on top of [RFC8231]. The PCEP extensions to support GMPLS-controlled networks should be considered under the same security as for MPLS networks, as noted in [RFC7025]. Therefore, the security considerations elaborated in [RFC5440] still apply to this draft. Furthermore, [RFC8231] provides a detailed analysis of the additional security issues incurred due to the new extensions and possible solutions needed to support for the new stateful PCE capabilities and they apply to this document as well.

9. Acknowledgement

We would like to thank Adrian Farrel and Cyril Margaria for the useful comments and discussions.

10. References

10.1. 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, JL., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC8174] B. Leiba, "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", RFC 8174, May 2017.

[RFC8231] Crabbe, E., Medved, J., Varga, R., Minei, I., "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, September 2017.

[PCEP-GMPLS] Margaria, C., Gonzalez de Dios, O., Zhang, F., "PCEP extensions for GMPLS", draft-ietf-pce-gmpls-pcep-extensions, work in progress.

10.2. Informative References

[RFC5511] A. Farrel, "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

[RFC8051] Zhang, X., Minei, I., et al, "Applicability of Stateful Path Computation Element (PCE) ", RFC 8051, January 2017.

[RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", RFC 8232, September 2017.

11. Contributors' Address

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972645
Email: zhang.xian@huawei.com

Dhruv Dhody
Huawei Technology
India

Email: dhruv.ietf@gmail.com

Yi Lin
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972914
Email: yi.lin@huawei.com

Authors' Addresses

Young Lee (Editor)
Huawei
5340 Legacy Drive, Suite 170
Plano, TX 75023
US

Phone: +1 469 278 5838
EMail: leeyoung@huawei.com

Fatai Zhang
Huawei
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
P.R. China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

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

Phone:
Email: ramon.casellas@cttc.es

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain

Phone: +34 913374013
Email: ogondio@tid.es

Zafar Ali
Cisco Systems
Email: zali@cisco.com