

Internet Engineering Task Force (IETF)
Request for Comments: 8282
Category: Standards Track
ISSN: 2070-1721

E. Oki
Kyoto University
T. Takeda
NTT
A. Farrel
Juniper Networks
F. Zhang
Huawei Technologies Co., Ltd.
December 2017

Extensions to the Path Computation Element Communication Protocol (PCEP)
for Inter-Layer MPLS and GMPLS Traffic Engineering

Abstract

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

MPLS and GMPLS networks may be constructed from layered service networks. It is advantageous for overall network efficiency to provide end-to-end traffic engineering across multiple network layers through a process called inter-layer traffic engineering. PCE is a candidate solution for such requirements.

The PCE Communication Protocol (PCEP) is designed as a communication protocol between Path Computation Clients (PCCs) and PCEs. This document presents PCEP extensions for inter-layer traffic engineering.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8282>.

Copyright Notice

Copyright (c) 2017 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	4
2.	Overview of PCE-Based Inter-Layer Path Computation	4
3.	Protocol Extensions	4
3.1.	INTER-LAYER Object	5
3.2.	SWITCH-LAYER Object	8
3.3.	REQ-ADAP-CAP Object	9
3.4.	New Metric Types	10
3.5.	SERVER-INDICATION Object	11
4.	Procedures	11
4.1.	Path Computation Request	11
4.2.	Path Computation Reply	12
4.3.	Stateful PCE and PCE Initiated LSPs	13
5.	Updated Format of PCEP Messages	14
6.	Manageability Considerations	15
7.	IANA Considerations	16
7.1.	New PCEP Objects	16
7.2.	New Registry for INTER-LAYER Object Flags	17
7.3.	New Metric Types	17
8.	Security Considerations	18
9.	References	18
9.1.	Normative References	18
9.2.	Informative References	19
	Acknowledgments	21
	Contributors	21
	Authors' Addresses	22

1. Introduction

The Path Computation Element (PCE) defined in [RFC4655] is an entity that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) may make requests to a PCE for paths to be computed, and a PCE may initiate or modify services in a network by supplying new paths [RFC8231] [RFC8281].

A network may comprise multiple layers. These layers may represent separation of technologies (e.g., packet switch capable (PSC), time division multiplex (TDM), and lambda switch capable (LSC)) [RFC3945]; separation of data-plane switching granularity levels (e.g., Virtual Circuit 4 (VC4) and VC12) [RFC5212]; or a distinction between client and server networking roles (e.g., commercial or administrative separation of client and server networks). In this multi-layer network, Label Switched Paths (LSPs) in lower layers are used to carry higher-layer LSPs. The network topology formed by lower-layer LSPs and advertised as traffic engineering links (TE links) in the higher layer is called a Virtual Network Topology (VNT) [RFC5212]. Discussion of other ways that network layering can be supported such that connectivity in a higher-layer network can be provided by LSPs in a lower-layer network is provided in [RFC7926].

It is important to optimize network resource utilization globally, i.e., taking into account all layers, rather than optimizing resource utilization at each layer independently. This allows better network efficiency to be achieved. This is what we call inter-layer traffic engineering. This includes mechanisms allowing the computation of end-to-end paths across layers (known as inter-layer path computation) and mechanisms for control and management of the VNT by setting up and releasing LSPs in the lower layers [RFC5212].

PCE can provide a suitable mechanism for resolving inter-layer path computation issues. The framework for applying the PCE-based path computation architecture to inter-layer traffic engineering is described in [RFC5623].

The PCE communication protocol (PCEP) is designed as a communication protocol between PCCs and PCEs and is defined in [RFC5440]. A set of requirements for PCEP extensions to support inter-layer traffic engineering is described in [RFC6457].

This document presents PCEP extensions for inter-layer traffic engineering that satisfy the requirements described in [RFC6457].

1.1. Requirements Language

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.

2. Overview of PCE-Based Inter-Layer Path Computation

[RFC4206] defines a way to signal a higher-layer LSP, which has an explicit route that includes hops traversed by LSPs in lower layers. The computation of end-to-end paths across layers is called inter-layer path computation.

A Label Switching Router (LSR) in the higher layer might not have information on the lower-layer topology, particularly in an overlay or augmented model [RFC3945]; hence, it may not be able to compute an end-to-end path across layers.

PCE-based inter-layer path computation consists of using one or more PCEs to compute an end-to-end path across layers. This could be achieved by relying on a single PCE that has topology information about multiple layers and can directly compute an end-to-end path across layers considering the topology of all of the layers. Alternatively, the inter-layer path computation could be performed using multiple cooperating PCEs where each PCE has information about the topology of one or more layers (but not all layers) and where the PCEs collaborate to compute an end-to-end path.

As described in [RFC5339], a hybrid node may advertise a single TE link with multiple switching capabilities. Normally, those TE links exist at the layer/region boarder. In this case, a PCE needs to be capable of specifying the server-layer path information when the server-layer path information is required to be returned to the PCC.

[RFC5623] describes models for inter-layer path computation in more detail. It introduces the Virtual Network Topology Manager (VNTM), a functional element that controls the VNT, and sets out three distinct models (and a fourth hybrid model) for inter-layer control involving a PCE, triggered signaling, and a Network Management System (NMS).

3. Protocol Extensions

This section describes PCEP extensions for inter-layer path computation. Four new objects are defined: the INTER-LAYER object, the SWITCH-LAYER object, the REQ-ADAP-CAP object, and the SERVER-INDICATION object. Also, two new metric types are defined.

3.1. INTER-LAYER Object

The INTER-LAYER object is optional and can be used in Path Computation Request (PCReq) and Path Computation Reply (PCRep) messages, and also in Path Computation State Report (PCRpt), Path Computation Update Request (PCUpd), and Path Computation LSP Initiate Request (PCInitiate) messages.

In a PCReq message, the INTER-LAYER object indicates whether inter-layer path computation is allowed, the type of path to be computed, and whether triggered signaling (hierarchical LSPs per [RFC4206] or stitched LSPs per [RFC5150] depending on physical network technologies) is allowed. When the INTER-LAYER object is absent from a PCReq message, the receiving PCE MUST process as though inter-layer path computation had been explicitly disallowed (I-bit set to zero -- see below).

In a PCRep message, the INTER-LAYER object indicates whether inter-layer path computation has been performed, the type of path that has been computed, and whether triggered signaling is used.

When a PCReq message includes more than one request, an INTER-LAYER object is used per request. When a PCRep message includes more than one path per request that is responded to, an INTER-LAYER object is used per path.

The applicability of this object to PCRpt and PCUpd messages is the same as for other objects on those messages as described in [RFC8231]. The applicability of this object to the PCInitiate message is the same as for other objects on those messages as described in [RFC8281]. These messages use the <attribute-list> as defined in [RFC5440] and extended by further PCEP extensions, so the <attribute-list> as extended in Section 5 can be used to include the INTER-LAYER object on these messages.

INTER-LAYER Object-Class is 36.

Inter-layer Object-Type is 1.

The format of the INTER-LAYER object body is shown in Figure 1.

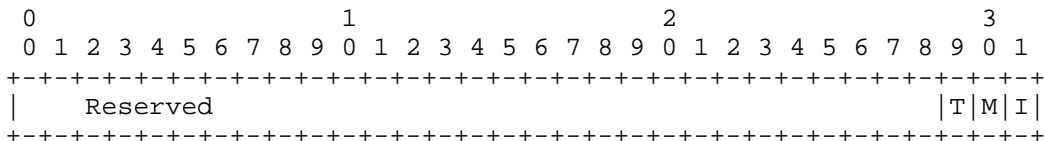


Figure 1: The INTER-LAYER Object

I flag (1 bit): The I flag is used by a PCC in a PCReq message to indicate to a PCE whether an inter-layer path is allowed. When the I flag is set (one), the PCE MAY perform inter-layer path computation and return an inter-layer path. When the flag is clear (zero), the path that is returned MUST NOT be an inter-layer path.

The I flag is used by a PCE in a PCRep message to indicate to a PCC whether the path returned is an inter-layer path. When the I flag is set (one), the path is an inter-layer path. When it is clear (zero), the path is contained within a single layer because either inter-layer path computation was not performed or a mono-layer path (without any virtual TE link and without any loose hop that spans the lower-layer network) was found notwithstanding the use of inter-layer path computation.

M flag (1 bit): The M flag is used by a PCC in a PCReq message to indicate to a PCE whether a mono-layer path or multi-layer path is requested. When the M flag is set (one), a multi-layer path is requested. When it is clear (zero), a mono-layer path is requested.

The M flag is used by a PCE in a PCRep message to indicate to a PCC whether a mono-layer path or multi-layer path is returned. When the M flag is set (one), a multi-layer path is returned. When the M flag is clear (zero), a mono-layer path is returned.

If the I flag is clear (zero), the M flag has no meaning and MUST be ignored.

[RFC6457] describes two sub-options for mono-layer path.

- o A mono-layer path that is specified by strict hops. The path may include virtual TE links.
- o A mono-layer path that includes loose hops that span the lower-layer network.

The choice of this sub-option can be specified by the use of the O flag in the Request Parameter (RP) object specified in [RFC5440].

T flag (1 bit): The T flag is used by a PCC in a PCReq message to indicate to a PCE whether triggered signaling is allowed. When the T flag is set (one), triggered signaling is allowed. When it is clear (zero), triggered signaling is not allowed.

The T flag is used by a PCE in a PCRep message to indicate to a PCC whether triggered signaling is required to support the returned path. When the T flag is set (one), triggered signaling is required. When it is clear (zero), triggered signaling is not required.

Note that triggered signaling is used to support hierarchical [RFC4206] or stitched [RFC5150] LSPs according to the physical attributes of the network layers.

If the I flag is clear (zero), the T flag has no meaning and MUST be ignored.

Note that the I and M flags differ in the following ways. When the I flag is clear (zero), virtual TE links must not be used in path computation. In addition, loose hops that span the lower-layer network must not be specified. Only regular TE links from the same layer may be used.

- o When the I flag is set (one), the M flag is clear (zero), and the T flag is set (one), virtual TE links are allowed in path computation. In addition, when the O flag of the RP object is set, loose hops that span the lower-layer network may be specified. This will initiate lower-layer LSP setup; thus, the inter-layer path is set up even though the path computation result from a PCE to a PCC includes hops from the same layer only.
- o However, when the I flag is set (one), the M flag is clear (zero), and the T flag is clear (zero), since triggered signaling is not allowed, virtual TE links that have not been pre-sigaled MUST NOT be used in path computation. In addition, loose hops that span the lower-layer network MUST NOT be specified. Therefore, this is equivalent to the I flag being clear (zero).

Reserved bits of the INTER-LAYER object sent between a PCC and PCE in the same domain MUST be transmitted as zero and SHOULD be ignored on receipt. A PCE that forwards a path computation request to other PCEs MUST preserve the settings of reserved bits in the PCReq messages it sends and in the PCRep messages it forwards to PCCs.

Note that the flags in the PCRpt message indicate the state of an LSP, whereas the flags in the PCUpd and the PCInitiate messages indicate the intended/desired state as determined by the PCE.

3.2. SWITCH-LAYER Object

The SWITCH-LAYER object is optional on a PCReq message and specifies switching layers in which a path MUST, or MUST NOT, be established. A switching layer is expressed as a switching type and encoding type.

When a SWITCH-LAYER object is used on a PCReq, it is interpreted in the context of the INTER-LAYER object on the same message. If no INTER-LAYER object is present, the PCE MUST process the SWITCH-LAYER object as though inter-layer path computation had been explicitly disallowed. In such a case, the SWITCH-LAYER object MUST NOT have more than one LSP Encoding Type and Switching Type with the I flag set.

The SWITCH-LAYER object is optional on a PCRep message, where it is used with the NO-PATH object in the case of unsuccessful path computation to indicate the set of constraints that could not be satisfied.

The SWITCH-LAYER object may be used on a PCRpt message consistent with how properties of existing LSps are reported on that message [RFC8231]. The PCRpt message uses the <attribute-list> as defined in [RFC5440] and extended by further PCEP extensions. This message can use the <attribute-list> as extended in Section 5 to carry the SWITCH-LAYER object. The SWITCH-LAYER object is not used on a PCUpd or PCInitiate messages.

SWITCH-LAYER Object-Class is 37.

Switch-layer Object-Type is 1.

The format of the SWITCH-LAYER object body is shown in Figure 2.

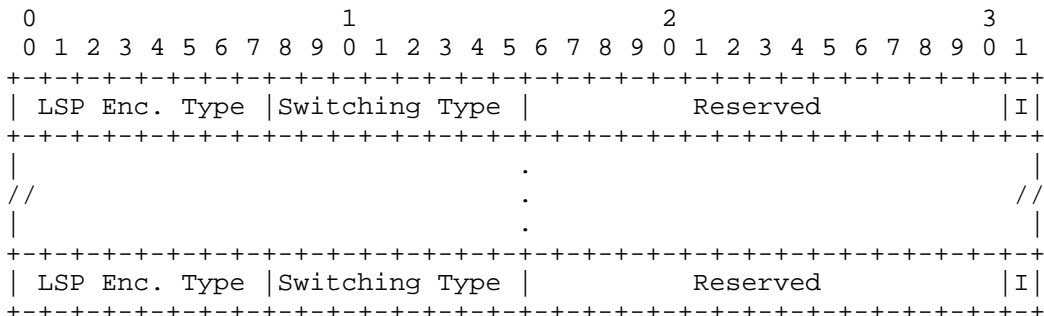


Figure 2: The SWITCH-LAYER Object

Each row indicates a switching type and encoding type that must or must not be used for a specified layer(s) in the computed path.

The format is based on [RFC3471] and has equivalent semantics.

LSP Encoding Type (8 bits): see [RFC3471] for a description of parameters.

Switching Type (8 bits): see [RFC3471] for a description of parameters.

I flag (1 bit): the I flag indicates whether a layer with the specified switching type and encoding type must or must not be used by the computed path. When the I flag is set (one), the computed path MUST traverse a layer with the specified switching type and encoding type. When the I flag is clear (zero), the computed path MUST NOT enter or traverse any layer with the specified switching type and encoding type.

When a combination of switching type and encoding type is not included in the SWITCH-LAYER object, the computed path MAY traverse a layer with that combination of switching type and encoding type.

A PCC may want to specify only a Switching Type and not an LSP Encoding Type. In this case, the LSP Encoding Type is set to zero.

3.3. REQ-ADAP-CAP Object

The REQ-ADAP-CAP object is optional and is used to specify a requested adaptation capability for both ends of the lower-layer LSP. The REQ-ADAP-CAP object is used in a PCReq message for inter-PCE communication, where the PCE that is responsible for computing higher-layer paths acts as a PCC to request a path computation from a PCE that is responsible for computing lower-layer paths.

The REQ-ADAP-CAP object is used in a PCRep message in case of unsuccessful path computation (in this case, the PCRep message also contains a NO-PATH object, and the REQ-ADAP-CAP object is used to indicate the set of constraints that could not be satisfied).

The REQ-ADAP-CAP object MAY be used in a PCReq message in a mono-layer network to specify a requested adaptation capability for both ends of the LSP. In this case, it MAY be carried without an INTER-LAYER object.

The applicability of this object to PCRpt and PCUpd messages is the same as for other objects on those messages as described in [RFC8231]. The applicability of this object to the PCInitiate

message is the same as for other objects on those messages as described in [RFC8281]. These messages use the <attribute-list> as defined in [RFC5440] and extended by further PCEP extensions. These messages can use the <attribute-list> as extended in Section 5 to carry the REQ-ADAP-CAP object.

REQ-ADAP-CAP Object-Class is 38.

Req-Adap-Cap Object-Type is 1.

The format of the REQ-ADAP-CAP object body is shown in Figure 3.

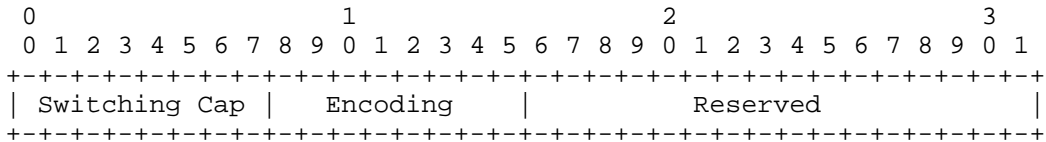


Figure 3: The REQ-ADAP-CAP Object

The format is based on [RFC6001] and has equivalent semantics as the Interface Adjustment Capability Descriptor (IACD) Upper Switching Capability and Lower Switching Capability fields.

Switching Capability (8 bits): see [RFC4203] for a description of parameters.

Encoding (8 bits): see [RFC3471] for a description of parameters.

A PCC may want to specify a Switching Capability, but not an Encoding. In this case, the Encoding MUST be set to zero.

3.4. New Metric Types

This document defines two new metric types for use in the PCEP METRIC object.

IANA has assigned the value 18 to indicate the metric "Number of adaptations on a path".

IANA has assigned the value 19 to indicate the metric "Number of layers on a path".

See Sections 4.1, 4.2, and 4.3 for a description of how these metrics are applied.

3.5. SERVER-INDICATION Object

The SERVER-INDICATION is optional and is used to indicate that path information included in the Explicit Route Object (ERO) is server-layer information, and it specifies the characteristics of the server layer, e.g., the switching capability and encoding of the server-layer path.

The format of the SERVER-INDICATION object body is shown in Figure 4.

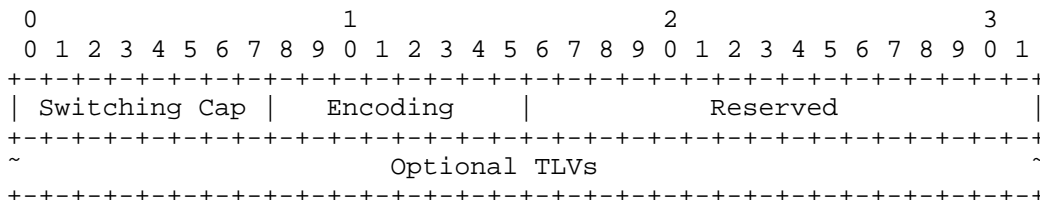


Figure 4: The SERVER-INDICATION Object

SERVER-INDICATION Object-Class is 39.

Server-indication Object-Type is 1.

Switching Capability (8 bits): see [RFC4203] for a description of parameters.

Encoding (8 bits): see [RFC3471] for a description of parameters.

Optional TLVs: Optional TLVs MAY be included within the object to specify more specific server-layer path information (e.g., traffic parameters). Such TLVs will be defined by other documents.

4. Procedures

4.1. Path Computation Request

A PCC requests or allows inter-layer path computation in a PCReq message by including the INTER-LAYER object with the I flag set. The INTER-LAYER object indicates whether inter-layer path computation is allowed, which path type is requested, and whether triggered signaling is allowed.

The SWITCH-LAYER object, which MUST NOT be present unless the INTER-LAYER object is also present, is optionally used to specify the switching types and encoding types that define layers that must, or must not, be used in the computed path. When the SWITCH-LAYER object is used with the INTER-LAYER object I flag clear (zero), inter-layer

path computation is not allowed, but constraints specified in the SWITCH-LAYER object apply. Example usage includes path computation in a single-layer GMPLS network.

The REQ-ADAP-CAP object is optionally used to specify the interface switching capability of both ends of the lower-layer LSP. The REQ-ADAP-CAP object is used in inter-PCE communication, where the PCE that is responsible for computing higher-layer paths makes a request as a PCC to a PCE that is responsible for computing lower-layer paths. Alternatively, the REQ-ADAP-CAP object may be used in the NMS-VNTM model, where the VNTM makes a request as a PCC to a PCE that is responsible for computing lower-layer paths.

The METRIC object is optionally used to specify metric types to be optimized or bounded. When metric type 18 is used, it indicates that path computation MUST minimize or bound the number of adaptations on a path. When metric type 19 is used, it indicates that path computation MUST minimize or bound the number of layers to be involved on a path.

Furthermore, in order to allow different Objective Functions (OFs) to be applied within different network layers, multiple OF objects [RFC5541] MAY be present. In such a case, the first OF object specifies an objective function for the higher-layer network, and subsequent OF objects specify objection functions of the subsequent lower-layer networks.

4.2. Path Computation Reply

In the case of successful path computation, the requested PCE replies to the requesting PCC for the inter-layer path computation result in a PCRep message that MAY include the INTER-LAYER object. When the INTER-LAYER object is included in a PCRep message, the I, M, and T flags indicate semantics of the path as described in Section 3.1. Furthermore, when the C flag of the METRIC object in a PCReq is set, the METRIC object MUST be included in the PCRep to provide the computed metric value, as specified in [RFC5440].

The PCE MAY specify the server-layer path information in the ERO. In this case, the requested PCE replies with a PCRep message that includes at least two sets of ERO information in the path-list: one is for the client-layer path information, and another one is the server-layer path information. When SERVER-INDICATION is included in a PCRep message, it indicates that the path in the ERO is the server-layer path information. The server-layer path specified in the ERO could be loose or strict. On receiving the replied path, the PCC (e.g., NMS and ingress node) can trigger the signaling to set up the LSPs according to the computed paths.

In the case of unsuccessful path computation, the PCRep message also contains a NO-PATH object, and the SWITCH-TYPE object and/or REQ-ADAP-CAP MAY be used to indicate the set of constraints that could not be satisfied.

4.3. Stateful PCE and PCE Initiated LSPs

Processing for stateful PCEs is described in [RFC8231]. That document defines the PCRpt message to allow a PCC to report to a PCE that an LSP already exists in the network and to delegate control of that LSP to the PCE.

When the LSP is a multi-layer LSP (or a mono-layer LSP for which specific adaptations exist), the message objects defined in this document are used on the PCRpt to describe an LSP that is delegated to the PCE so that the PCE may process the LSP.

Furthermore, [RFC8231] defines the PCUpd message to allow a PCE to modify an LSP that has been delegated to it. When the LSP is a multi-layer LSP (or a mono-layer LSP for which specific adaptations exist), the message objects defined in this document are used on the PCUpd to describe the new attributes of the modified LSP.

Processing for PCE-initiated LSPs is described in [RFC8281]. That document defines the PCInitiate message that is used by a PCE to request a PCC to set up a new LSP. When the LSP is a multi-layer LSP (or a mono-layer LSP for which specific adaptations exist), the message objects defined in this document are used on the PCInitiate to describe the attributes of the new LSP.

The new metric types defined in this document can also be used with the stateful PCE extensions. The format of PCEP messages described in [RFC8231] and [RFC8281] uses <attribute-list> (which is extended in Section 5 for the purpose of including the new metrics).

The stateful PCE implementation MAY use the extension of PCReq and PCRep messages as defined in Section 5 to also enable the use of inter-layer parameters during passive stateful operations, using the LSP object.

5. Updated Format of PCEP Messages

Message formats in this section, as those in [RFC5440], are presented using Routing Backus-Naur Format (RBNF) as specified in [RFC5511].

The format of the PCReq message is updated as shown in Figure 5.

```

<PCReq Message> ::= <Common Header>
                    [<svec-list>]
                    <request-list>

where:
  <svec-list> ::= <SVEC>
                 [<svec-list>]

  <request-list> ::= <request> [<request-list>]

  <request> ::= <RP>
               <END-POINTS>
               [<LSP>]
               [<LSPA>]
               [<BANDWIDTH>]
               [<metric-list>]
               [<of-list>]
               [<RRO> [<BANDWIDTH>]]
               [<IRO>]
               [<LOAD-BALANCING>]
               [<INTER-LAYER> [<SWITCH-LAYER>]]
               [<REQ-ADAP-CAP>]

where:
  <of-list> ::= <OF> [<of-list>]
  <metric-list> ::= <METRIC> [<metric-list>]

```

Figure 5: The Updated PCReq Message

The format of the PCRep message is updated as shown in Figure 6.

```

<PCRep Message> ::= <Common Header>
                    <response-list>

where:
  <response-list> ::= <response>[ <response-list>]

  <response> ::= <RP>
                [<LSP>]
                [<NO-PATH>]
                [<attribute-list>]
                [<path-list>]

  <path-list> ::= <path>[ <path-list>]

  <path> ::= <ERO><attribute-list>

where:
  <attribute-list> ::= [ <of-list>]
                     [<LSPA>]
                     [<BANDWIDTH>]
                     [<metric-list>]
                     [<IRO>]
                     [<INTER-LAYER>]
                     [<SWITCH-LAYER>]
                     [<REQ-ADAP-CAP>]
                     [<SERVER-INDICATION>]

  <of-list> ::= <OF>[ <of-list>]
  <metric-list> ::= <METRIC>[ <metric-list>]

```

Figure 6: The Updated PCRep Message

6. Manageability Considerations

Implementations of this specification should provide a mechanism to configure any optional features (such as whether a PCE supports inter-layer computation and which metrics are supported).

A Management Information Base (MIB) module for modeling PCEP is described in [RFC7420]. Systems that already use a MIB module to manage their PCEP implementations might want to augment that module to provide controls and indicators for support of inter-layer features defined in this document and to add counters of messages sent and received containing the objects defined here.

However, the preferred mechanism for configuration is through a YANG model. Work has started on a YANG model for PCEP [PCEP-YANG], and this could be enhanced as described for the MIB module, above.

Additional policy configuration might be provided to allow a PCE to discriminate between the computation services offered to different PCCs.

A set of monitoring tools for the PCE-based architecture are provided in [RFC5886]. Systems implementing this specification and PCE monitoring should consider defining extensions to the mechanisms defined in [RFC5886] to help monitor inter-layer path computation requests.

7. IANA Considerations

IANA maintains a registry called "Path Computation Element Protocol (PCEP) Numbers". Per this document, IANA has carried out actions on subregistries of that registry.

7.1. New PCEP Objects

IANA has made the following assignments in the "PCEP Objects" subregistry.

Object-Class Value	Name	Object-Type	Reference
INTER-LAYER	36	0: Reserved 1: Inter-layer 2-15: Unassigned	RFC 8282
SWITCH-LAYER	37	0: Reserved 1: Switch-layer 2-15: Unassigned	RFC 8282
REQ-ADAP-CAP	38	0: Reserved 1: Req-Adap-Cap 2-15: Unassigned	RFC 8282
SERVER-INDICATION	39	0: Reserved 1: Server-indication	RFC 8282

Figure 7: New PCEP Objects

7.2. New Registry for INTER-LAYER Object Flags

IANA has created a new subregistry to manage the Flag field of the INTER-LAYER object called the "Inter-Layer Object Path Property Bits" registry.

New bit numbers may be allocated only by "IETF Review" [RFC8126]. Each bit should be tracked with the following qualities:

- o Bit number (counting from bit 0 as the most significant bit up to a maximum of bit 31)
- o Capability Description
- o Defining RFC

IANA has populated the registry as follows:

Bit	Flag	Multi-Layer Path Property	Reference
0-28		Unassigned	
29	T	Triggered Signaling Allowed	RFC 8282
30	M	Multi-Layer Requested	RFC 8282
31	I	Inter-Layer Allowed	RFC 8282

Figure 8: New Registry for INTER-LAYER Object Flags

7.3. New Metric Types

Two new metric types are defined in this document for the METRIC object (specified in [RFC5440]). IANA has made the following allocations from the "Metric Object T Field" registry.

Value	Description	Reference
18	Number of adaptations on a path	RFC 8282
19	Number of layers on a path	RFC 8282

Figure 9: New Metric Types

IANA has updated the registry to show the registration procedure of "IETF Review" as already documented in [RFC5440].

8. Security Considerations

Inter-layer traffic engineering with PCE may raise new security issues when PCE-PCE communication is done between different layer networks for inter-layer path computation. Security issues may also exist when a single PCE is granted full visibility of TE information that applies to multiple layers.

The Path-Key-based mechanism defined in [RFC5520] MAY be applied to address the topology confidentiality between different layers.

9. References

9.1. Normative References

- [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>>.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <<https://www.rfc-editor.org/info/rfc3471>>.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, DOI 10.17487/RFC3945, October 2004, <<https://www.rfc-editor.org/info/rfc3945>>.
- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<https://www.rfc-editor.org/info/rfc4203>>.
- [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>>.
- [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>>.

- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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, <<http://www.rfc-editor.org/info/rfc20>>.

9.2. Informative References

- [PCEP-YANG] Dhody, D., Hardwick, J., Beeram, V., and j. jeffrant@gmail.com, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, draft-ietf-pce-pcep-yang-05, June 2017.
- [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>>.

- [RFC5212] Shiomoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, DOI 10.17487/RFC5212, July 2008, <<https://www.rfc-editor.org/info/rfc5212>>.
- [RFC5339] Le Roux, JL., Ed. and D. Papadimitriou, Ed., "Evaluation of Existing GMPLS Protocols against Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 5339, DOI 10.17487/RFC5339, September 2008, <<https://www.rfc-editor.org/info/rfc5339>>.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <<https://www.rfc-editor.org/info/rfc5511>>.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", RFC 5541, DOI 10.17487/RFC5541, June 2009, <<https://www.rfc-editor.org/info/rfc5541>>.
- [RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, DOI 10.17487/RFC5623, September 2009, <<https://www.rfc-editor.org/info/rfc5623>>.
- [RFC5886] Vasseur, JP., Ed., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", RFC 5886, DOI 10.17487/RFC5886, June 2010, <<https://www.rfc-editor.org/info/rfc5886>>.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, DOI 10.17487/RFC6001, October 2010, <<https://www.rfc-editor.org/info/rfc6001>>.
- [RFC6457] Takeda, T., Ed. and A. Farrel, "PCC-PCE Communication and PCE Discovery Requirements for Inter-Layer Traffic Engineering", RFC 6457, DOI 10.17487/RFC6457, December 2011, <<https://www.rfc-editor.org/info/rfc6457>>.

- [RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", RFC 7420, DOI 10.17487/RFC7420, December 2014, <<https://www.rfc-editor.org/info/rfc7420>>.
- [RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.

Acknowledgments

The authors would like to thank Cyril Margaria for his valuable comments. Helpful comments and suggested text were offered by Dhruv Dhody, who also fixed the RBNF. Jonathan Hardwick provided a helpful review as document shepherd.

Contributors

Jean-Louis Le Roux
France Telecom R&D
Av Pierre Marzin
Lannion 22300
France

Email: jeanlouis.leroux@orange.com

Authors' Addresses

Eiji Oki
Kyoto University
Yoshida-honmachi, Sakyo-ku, Kyoto
Japan

Email: oki@i.kyoto-u.ac.jp

Tomonori Takeda
NTT
3-9-11 Midori-cho
Musashino-shi, Tokyo
Japan

Email: tomonori.takeda@ntt.com

Adrian Farrel
Juniper Networks

Email: afarrel@juniper.net

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

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