Internet Engineering Task Force (IETF)
Request for Comments: 6457
Category: Informational
ISSN: 2070-1721

T. Takeda, Ed.
NTT
A. Farrel
Old Dog Consulting

December 2011

PCC-PCE Communication and PCE Discovery Requirements for Inter-Layer Traffic Engineering

Abstract

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

MPLS and GMPLS networks may be constructed from layered client/server networks. It is advantageous for overall network efficiency to provide end-to-end traffic engineering across multiple network layers. PCE is a candidate solution for such requirements.

Generic requirements for a communication protocol between Path Computation Clients (PCCs) and PCEs are presented in RFC 4657, "Path Computation Element (PCE) Communication Protocol Generic Requirements". Generic requirements for a PCE discovery protocol are presented in RFC 4674, "Requirements for Path Computation Element (PCE) Discovery".

This document complements the generic requirements and presents detailed sets of PCC-PCE communication protocol requirements and PCE discovery protocol requirements for inter-layer traffic engineering.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6457.

Table of Contents

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 network may comprise multiple layers. These layers may represent the separation of technologies (e.g., Packet Switch Capable (PSC), Time Division Multiplex (TDM), lambda switch capable (LSC)) into GMPLS regions [RFC3945], the separation of data plane switching

granularity levels (e.g., PSC-1 and PSC-2 or Virtual Circuit 4 (VC4) and VC12) into GMPLS layers [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 upper-layer LSPs. The network topology formed by lower-layer LSPs and advertised to the higher layer is called a "Virtual Network Topology (VNT)" [RFC5212].

In layered networks under the operation of Multiprotocol Label Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) protocols, it is important to provide mechanisms to allow global optimization of network resources. That is, to take 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 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].

Inter-layer traffic engineering is included in the scope of the PCE architecture [RFC4655], and PCE can provide a suitable mechanism for resolving inter-layer path computation issues. The applicability of the PCE-based path computation architecture to inter-layer traffic engineering is described in [RFC5623].

This document presents sets of requirements for communication between Path Computation Clients (PCCs) and PCEs using the PCE Communication Protocol (PCEP) and for PCE discovery for inter-layer traffic engineering. It supplements the generic requirements documented in [RFC4657], [RFC4674], and the framework provided in [RFC5623].

1.1. Terminology

LSP: Label Switched Path.

LSR: Label Switching Router.

PCC: A Path Computation Client is any client entity (component, application or network node) requesting a path computation to be performed by a Path Computation Element.

PCE: A Path Computation Element is an entity that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: A PCE Communication Protocol is a protocol for communication between PCCs and PCEs.

Although this requirements document is informational and not a protocol specification, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] for clarity of requirement specification.

2. Motivation for PCE-Based Inter-Layer Path Computation

[RFC4206] defines a way to signal an MPLS or a GMPLS LSP with an explicit route in a higher layer of a network that includes hops traversed by LSPs in lower layers of the network. The computation of end-to-end paths across layers is called "inter-layer path computation".

An LSR in the higher layer might not have information on the topology of lower layers, particularly in an overlay or augmented model; hence, it might not be able to compute an end-to-end path across layers.

PCE-based inter-layer path computation consists of relying on one or more PCEs to compute an end-to-end path across layers. This could rely on a single PCE path computation where the PCE 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 as a multiple PCE computation, where each member of a set of PCEs has information about the topology of one or more layers, but not all layers, and they collaborate to compute an end-to-end path.

Consider a two-layer network where the higher-layer network is a packet-based IP/MPLS or GMPLS network and the lower-layer network is a GMPLS-controlled optical network. An ingress LSR in the higher-layer network tries to set up an LSP to an egress LSR also in the higher-layer network across the lower-layer network, and it needs a path in the higher-layer network. However, suppose that there is no TE link between border LSRs, which are located on the boundary between the higher-layer and lower-layer networks, and that the ingress LSR does not have topology visibility in the lower layer. If a single-layer path computation is applied for the higher layer, the path computation fails. On the other hand, inter-layer path computation is able to provide a route in the higher layer and a suggestion that a lower-layer LSP be set up between border LSRs, considering both layers as TE topologies.

Further discussion of the application of PCE to inter-layer path computation can be found in [RFC5623].

3. PCC-PCE Communication and Discovery Requirements for Inter-Layer Traffic Engineering

This section provides additional requirements specific to the problems of multi-layer TE that are not covered in [RFC4657] or [RFC4674].

3.1. PCC-PCE Communication

PCEP MUST allow requests and replies for inter-layer path computation.

This requires no additional messages, but it implies the following additional constraints to be added to PCEP.

3.1.1. Control of Inter-Layer Path Computation

A request from a PCC to a PCE MUST support the inclusion of an optional indication of whether inter-layer path computation is allowed. In the absence of such an indication, the default is that inter-layer path computation is not allowed.

3.1.2. Control of the Type of Path to Be Computed

The PCE computes and returns a path to the PCC that the PCC can use to build a higher-layer or lower-layer LSP once converted to an Explicit Route Object (ERO) for use in RSVP - Traffic Engineering (RSVP-TE) signaling. There are two options [RFC5623].

- Option 1: Mono-Layer Path. The PCE computes a "mono-layer" path, i.e., a path that includes only TE links from the same layer.

- Option 2: Multi-Layer Path. The PCE computes a "multi-layer" path, i.e., a path that includes TE links from distinct layers [RFC4206].

It may be necessary or desirable for a PCC to control the type of path that is produced by a PCE. For example, a PCC may know that it is not possible, for technological or policy reasons, to signal a multi-layer path and that a mono-layer path is required, or the PCC may know that it does not wish the layer border node to have control of path computation. In order to make this level of control possible, PCEP MUST allow the PCC to select the path types to be computed, and that may be returned, by choosing one or more from the following list:

- A mono-layer path that is specified by strict hop(s). The path may include virtual TE link(s).

- A mono-layer path that includes loose hop(s).

- A multi-layer path that can include the path (as strict or loose hops) of one or more lower-layer LSPs not yet established.

The path computation response from a PCE to a PCC MUST report the type of path computed, and where a multi-layer path is returned, PCEP MUST support the inclusion, as part of end-to-end path, of the path of the lower-layer LSPs to be established.

If a response message from a PCE to PCC carries a mono-layer path that is specified by strict hops but includes virtual TE link(s), includes loose hop(s), or carries a multi-layer path that can include the complete path of one or more lower-layer LSPs not yet established, the signaling of the higher-layer LSP may trigger the establishment of the lower-layer LSPs (triggered signaling). The triggered signaling may increase the higher-layer connection setup latency. An ingress LSR for the higher-layer LSP, or a PCC, needs to know whether or not triggered signaling is required.

A request from a PCC to a PCE MUST allow indicating whether or not triggered signaling is acceptable.

A response from a PCE to a PCC MUST allow indicating whether or not the computed path requires triggered signaling.

Note that a PCE may not be able to distinguish virtual TE links from regular TE links. In such cases, even if a request from a PCC to a PCE indicates that triggered signaling is not acceptable, a PCE may choose virtual TE links in path computation. Therefore, when a network uses virtual TE links and a PCE is not able to distinguish virtual TE links from regular TE links, a PCE MAY choose virtual TE links even if a request from a PCC to a PCE indicates triggered signaling is not acceptable.

Also, note that an ingress LSR of a higher-layer or lower-layer LSP may be present in multiple layers. Thus, even when a mono-layer path is requested or supplied, PCEP MUST be able to indicate the required/provided path layer.

3.1.3. Communication of Inter-Layer Constraints

A request from a PCC to a PCE MUST support the inclusion of constraints for a multi-layer path. This includes control over which network layers may, must, or must not be included in the computed path. Such control may be expressed in terms of the switching types of the layer networks.

Furthermore, it may be desirable to constrain the number of layer boundaries crossed (i.e., the number of adaptations in the sense used in [RFC5212] performed on the end-to-end path), so PCEP SHOULD include a constraint or objective function to minimize or cap the number of adaptations on a path and a mechanism to report that number when a path is supplied.

The path computation request MUST also allow for different objective functions to be applied within different network layers. For example, the path in a packet-network may need to be optimized for least delay using the IGP metric as a measure of delay, while the path in an underlying TDM network might be optimized for fewest hops.

3.1.4. Adaptation Capability

The concept of adaptation is used here as introduced in [RFC5212].

It MUST be possible for the path computation request to indicate the desired adaptation function at the end points of the lower-layer LSP that is being computed. This will be particularly important where the ingress and egress LSR participate in more than one layer network but may not be capable of all associated adaptations.

3.1.5. Cooperation between PCEs

When each layer is in scope of a different PCE, which only has access to the topology information of its layer, the PCEs of each layer need to cooperate to perform inter-layer path computation. In this case, communication between PCEs is required for inter-layer path computation. A PCE that behaves as a client is defined as a PCC [RFC4655].

PCEP MUST allow requests and replies for multiple PCE inter-layer path computation.

3.1.6. Inter-Layer Diverse Paths

PCEP MUST allow for the computation of diverse inter-layer paths. A request from a PCC to a PCE MUST support the inclusion of multiple path requests, with the desired level of diversity at each layer (link, node, Shared Risk Link Group (SRLG)).

3.2. Capabilities Advertisements for PCE Discovery

In the case where there are several PCEs with distinct capabilities available, a PCC has to select one or more appropriate PCEs. For that purpose, the PCE discovery mechanism MAY support the disclosure of some detailed PCE capabilities. A PCE MAY (to be consistent with the above text and RFC 4674) be able to advise the following PCE capabilities related to inter-layer path computation:

- Support for inter-layer path computation
- Support for mono-layer/multi-layer paths
- Support for inter-layer constraints
- Support for adaptation capability
- Support for inter-PCE communication
- Support for inter-layer diverse path computation

3.3. Supported Network Models

PCEP SHOULD allow several architectural alternatives for interworking between MPLS- and GMPLS-controlled networks: overlay, integrated, and augmented models [RFC3945] [RFC5145] [RFC5146].

4. Manageability Considerations

4.1. Control of Function and Policy

An individual PCE MAY elect to support inter-layer computations and advertise its capabilities as described in the previous sections. PCE implementations MAY provide a configuration switch to allow support of inter-layer path computations to be enabled or disabled. When the level of support is changed, this SHOULD be re-advertised.

However, a PCE MAY also elect to support inter-layer computations, but not to advertise the fact, so that only those PCCs configured to know of the PCE and its capabilities can use it.

Support for, and advertisement of support for, inter-layer path computation MAY be subject to policy and a PCE MAY hide its inter-layer capabilities from certain PCCs by not advertising them through the discovery protocol and not reporting them to the specific PCCs in any PCEP capabilities exchange. Further, a PCE MAY be directed by policy to refuse an inter-layer path computation request for any reason including, but not limited to, the identity of the PCC that makes the request.

A further discussion of policy-enabled path computation can be found in [RFC5394].

4.2. Information and Data Models

PCEP extensions to support inter-layer computations MUST be accompanied by MIB objects for the control and monitoring of the protocol and of the PCE that performs the computations. The MIB objects MAY be provided in the same MIB module as used for general PCEP control and monitoring [PCEP-MIB] or MAY be provided in a new MIB module.

The MIB objects MUST provide the ability to control and monitor all aspects of PCEP relevant to inter-layer path computation.

4.3. Liveness Detection and Monitoring

No changes are necessary to the liveness detection and monitoring requirements as already embodied in [RFC4657]. It should be noted, however, that inter-layer path computations might require extended cooperation between PCEs (as is also the case for inter-AS (Autonomous System) and inter-area computations), and so the liveness detection and monitoring SHOULD be applied to each PCEP communication and aggregated to report the behavior of an individual PCEP request to the originating PCC.

In particular, where a request is forwarded between multiple PCEs, neither the PCC nor the first PCE can monitor the liveness of all PCE-PCE connections or of the PCEs themselves. In this case, suitable performance of the original PCEP request relies on each PCE operating correct monitoring procedures and correlating any failures back to the PCEP requests that are outstanding. These requirements are no different from those for any cooperative PCE usage, and they are expected already to be covered by general, and by inter-AS and inter-area, implementations. Such a procedure is specified in [RFC5441]. In addition, [RFC5886] specifies mechanisms to gather various state metrics along the path computation chain.

4.4. Verifying Correct Operation

There are no additional requirements beyond those expressed in [RFC4657] for verifying the correct operation of the PCEP. Note that verification of the correct operation of the PCE and its algorithms is out of scope for the protocol requirements, but a PCC MAY send the same request to more than one PCE and compare the results.

4.5. Requirements on Other Protocols and Functional Components

A PCE operates on a topology graph that may be built using information distributed by TE extensions to the routing protocol operating within the network. In order that the PCE can select a suitable path for the signaling protocol to use to install the inter-layer LSP, the topology graph must include information about the inter-layer signaling and forwarding (i.e., adaptation) capabilities of each LSR in the network.

Whatever means are used to collect the information to build the topology graph, the graph MUST include the requisite information. If the TE extensions to the routing protocol are used, these SHOULD satisfy the requirements as described in [RFC5212].

4.6. Impact on Network Operation

This section examines the impact on network operations of the use of a PCE for inter-layer traffic engineering. It does not present any further requirements on the PCE or PCC, for the PCEP or for deployment.

The use of a PCE to compute inter-layer paths is not expected to have significant impact on network operations if the upper-layer traffic engineering practices are aware of the frequent changes that might occur in the VNT. It should also be noted that the introduction of inter-layer support to a PCE that already provides mono-layer path computation might change the loading of the PCE and that might have an impact on the network behavior especially during recovery periods immediately after a network failure.

On the other hand, it is envisioned that the use of inter-layer path computation will have significant benefits to the operation of a multi-layer network including improving the network resource usage and enabling a greater number of higher-layer LSPs to be supported.

5. Security Considerations

Inter-layer traffic engineering with PCE may raise new security issues when PCE-PCE communication is used 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 formal introduction of a VNT Manager component, as described in [RFC5623], provides the basis for the application of inter-layer security and policy.

It is expected that solutions for inter-layer protocol extensions will address these issues in detail.

6. Acknowledgments

We would like to thank Kohei Shiomoto, Ichiro Inoue, Dean Cheng, Meral Shirazipour, Julien Meuric, and Stewart Bryant for their useful comments. Thanks to members of ITU-T Study Group 15, Question 14 for their constructive comments during the liaison process.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

7.2. Informative References

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC4657] Ash, J., Ed., and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[RFC4674] Le Roux, J., Ed., "Requirements for Path Computation Element (PCE) Discovery", RFC 4674, October 2006.

[RFC5145] Shiomoto, K., Ed., "Framework for MPLS-TE to GMPLS Migration", RFC 5145, March 2008.

[RFC5146] Kumaki, K., Ed., "Interworking Requirements to Support Operation of MPLS-TE over GMPLS Networks", RFC 5146, March 2008.

[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, July 2008.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, December 2008.

[RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, September 2009.

[PCEP-MIB] A. Koushik, and E. Stephan, "PCE communication protocol (PCEP) Management Information Base", Work in Progress, July 2010.

[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, April 2009.

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

Contributing Authors

Eiji Oki
University of Electro-Communications
Tokyo, Japan
EMail: oki@ice.uec.ac.jp

Jean-Louis Le Roux
France Telecom R&D,
Av Pierre Marzin,
22300 Lannion, France
EMail: jeanlouis.leroux@orange-ftgroup.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460, JAPAN
EMail: ke-kumaki@kddi.com

Authors' Addresses

Tomonori Takeda (editor)
NTT
3-9-11 Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
EMail: takeda.tomonori@lab.ntt.co.jp

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk