

Internet Engineering Task Force (IETF)
Request for Comments: 5862
Category: Informational
ISSN: 2070-1721

S. Yasukawa
NTT Corporation
A. Farrel
Old Dog Consulting
June 2010

Path Computation Clients (PCC) - Path Computation Element (PCE)
Requirements for Point-to-Multipoint MPLS-TE

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.

Extensions to the MPLS and GMPLS signaling and routing protocols have been made in support of point-to-multipoint (P2MP) Traffic Engineered (TE) Label Switched Paths (LSPs). The use of PCE in MPLS networks is already established, and since P2MP TE LSP routes are sometimes complex to compute, it is likely that PCE will be used for P2MP LSPs.

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". This document complements the generic requirements and presents a detailed set of PCC-PCE communication protocol requirements for point-to-multipoint MPLS/GMPLS 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/rfc5862>.

Copyright Notice

Copyright (c) 2010 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.

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. The intention is that the PCE is used to compute the path of Traffic Engineered Label Switched Paths (TE LSPs) within Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

Requirements for point-to-multipoint (P2MP) MPLS TE LSPs are documented in [RFC4461], and signaling protocol extensions for setting up P2MP MPLS TE LSPs are defined in [RFC4875]. P2MP MPLS TE networks are considered in support of various features, including layer 3 multicast virtual private networks [RFC4834].

Path computation for P2MP TE LSPs presents a significant challenge, and network optimization of multiple P2MP TE LSPs requires considerable computational resources. PCE offers a way to offload such path computations from Label Switching Routers (LSRs).

The applicability of the PCE-based path computation architecture to P2MP MPLS TE is described in a companion document [RFC5671]. No further attempt is made to justify the use of PCE for P2MP MPLS TE within this document.

This document presents a set of PCC-PCE communication protocol (PCECP) requirements for P2MP MPLS traffic engineering. It supplements the generic requirements documented in [RFC4657].

2. Conventions Used in This Document

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

3. PCC-PCE Communication Requirements for P2MP MPLS Traffic Engineering

This section sets out additional requirements specific to P2MP MPLS TE that are not covered in [RFC4657].

3.1. PCC-PCE Communication

The PCC-PCE communication protocol MUST allow requests and replies for the computation of paths for P2MP LSPs.

This requires no additional messages, but requires the addition of the parameters described in the following sections to the existing PCC-PCE communication protocol messages.

3.1.1. Indication of P2MP Path Computation Request

R1: Although the presence of certain parameters (such as a list of more than one destination) MAY be used by a protocol specification to allow an implementation to infer that a Path Computation Request is for a P2MP LSP, an explicit parameter SHOULD be placed in a conspicuous place within a Path Computation Request message to allow a receiving PCE to easily identify that the request is for a P2MP path.

3.1.2. Indication of P2MP Objective Functions

R2: [RFC4657] includes the requirement to be able to specify the objective functions to be applied by a PCE during path computation.

This document makes no change to that requirement, but it should be noted that new and different objective functions will be used for P2MP computation. Definitions for core objective functions can be found in [RFC5541] together with usage procedures. New objective functions for use with P2MP path computations will need to be defined and allocated codepoints in a separate document.

3.1.3. Non-Support of P2MP Path Computation

R3: PCEs are not required to support P2MP path computation. Therefore, it MUST be possible for a PCE to reject a P2MP Path Computation Request message with a reason code that indicates no support for P2MP path computation.

3.1.4. Non-Support by Back-Level PCE Implementations

It is possible that initial PCE implementations will be developed without support for P2MP path computation and without the ability to recognize the explicit parameter described in Section 3.1.1. Such legacy implementations will not be able to make use of the new reason code described in Section 3.1.3.

R4: Therefore, at least one parameter required for inclusion in a P2MP Path Computation Request message MUST be defined in such a way as to cause automatic rejection as unprocessable or unrecognized by a back-level PCE implementation without requiring any changes to that PCE. It is RECOMMENDED that the parameter that causes this result be the parameter described in Section 3.1.1.

3.1.5. Specification of Destinations

R5: Since P2MP LSPs have more than one destination, it MUST be possible for a single Path Computation Request to list multiple destinations.

3.1.6. Indication of P2MP Paths

R6: The Path Computation Response MUST be able to carry the path of a P2MP LSP.

P2MP paths can be expressed as a compressed series of routes as described in [RFC4875]. The Path Computation Response MUST be able to carry the P2MP path as either a compressed path (but not necessarily using the identical encoding as described in [RFC4875]), or as a non-compressed path comprising a series of source-to-leaf point-to-point (P2P) paths (known as S2L sub-paths).

R7: By default, the path returned by the PCE SHOULD use the compressed format.

The request from the PCC MAY allow the PCC to express a preference for receiving a compressed or non-compressed P2MP path in the response.

3.1.7. Multi-Message Requests and Responses

R8: A single P2MP LSP may have many destinations, and the computed path (tree) may be very extensive. In these cases, it is possible that the entire Path Computation Request or Response cannot fit within one PCE message. Therefore, it MUST be possible for a single request or response to be conveyed by a sequence of PCE messages.

Note that there is a requirement in [RFC4657] for reliable and in-order message delivery, so it is assumed that components of the sequence will be delivered in order and without missing components.

3.1.8. Non-Specification of Per-Destination Constraints and Parameters

[RFC4875] requires that all branches of a single P2MP LSP have the same characteristics, and achieves this by not allowing the signaling parameters to be varied for different branches of the same P2MP LSP.

R9: It MUST NOT be possible to set different constraints, traffic parameters, or quality-of-service requirements for different destinations of a P2MP LSP within a single computation request.

3.1.9. Path Modification and Path Diversity

R10: No changes are made to the requirement to support path modification and path diversity as described in [RFC4657]. Note, however, that a consequence of this requirement is that it MUST be possible to supply an existing path in a Path Computation Request. This requirement is unchanged from [RFC4657], but it is a new requirement that such paths MUST be able to be P2MP paths. The PCC MUST be able to supply these paths as compressed paths or as non-compressed paths (see Section 3.1.6) according to the preference of the PCC.

3.1.10. Reoptimization of P2MP TE LSPs

R11: Reoptimization MUST be supported for P2MP TE LSPs as described for P2P LSPs in [RFC4657]. To support this, the existing path MUST be supplied as described in Section 3.1.9.

Because P2MP LSPs are more complex, it is often the case that small optimization improvements can be made after changes in network resource availability. However, re-signaling any LSP introduces risks to the stability of the service provided to the customer and the stability of the network, even when techniques like make-before-break [RFC3209] are used. Therefore, a P2MP Path Computation Request SHOULD contain a parameter that allows

the PCC to express a cost-benefit reoptimization threshold for the whole LSP, as well as per destination. The setting of this parameter is subject to local policy at the PCC and SHOULD be subject to policy at the PCE [RFC5394].

Path reoptimization responses SHOULD indicate which of the routes (as supplied according to Section 3.1.6) have been modified from the paths supplied in the request.

3.1.11. Addition and Removal of Destinations from Existing Paths

A variation of path modification described in Section 3.1.9 is that destinations may be added to, or removed from, existing P2MP TE LSPs.

In the case of the addition of one or more destinations, it is necessary to compute a path for a new branch of the P2MP LSP. It may be desirable to recompute the whole P2MP tree, to add the new branch as a simple spur from the existing tree, or to recompute part of the P2MP tree.

R12: To support this function for leaf additions, it MUST be possible to make the following indications in a Path Computation Request:

- The path of an existing P2MP LSP (as described in Section 3.1.9).
- Which destinations are new additions to the tree.
- Which destinations of the existing tree must not have their paths modified.

It MAY also be possible to indicate in a Path Computation Request a cost-benefit reoptimization threshold, such that the addition of new leaves will not cause reoptimization of the existing P2MP tree unless a certain improvement is made over simply grafting the new leaves to the existing tree. (Compare with Section 3.1.10.)

In the case of the deletion of one or more destinations, it is not necessary to compute a new path for the P2MP TE LSP, but such a computation may yield optimizations over a simple pruning of the tree. The recomputation function in this case is essentially the same as that described in Section 3.1.10, but note that it MAY be possible to supply the full previous path of the entire P2MP TE LSP (that is, before the deletion of the destinations) in the Path Computation Request.

For both addition and deletion of destinations, the Path Computation Response SHOULD indicate which of the routes (as supplied according to Section 3.1.6) have been modified from the paths supplied in the Path Computation Request, as described in Section 3.1.10.

Note that the selection of all of these options is subject to local policy at the PCC and SHOULD be subject to policy at the PCE [RFC5394].

3.1.12. Specification of Applicable Branch Nodes

For administrative or security reasons, or for other policy reasons, it may be desirable to limit the set of nodes within the network that may be used as branch points for a given LSP, i.e., to provide to the path computation a limiting set of nodes that can be used as branches for a P2MP path computation, or to provide a list of nodes that must not be used as branch points.

R13: The PCC MUST be able to specify in a Path Computation Request a list of nodes that constitutes a limiting superset of the branch nodes for a P2MP path computation.

A PCC MUST be able to specify in a Path Computation Request a list of nodes that must not be used as branch nodes for a P2MP path computation.

3.1.13. Capabilities Exchange

PCE capabilities exchange forms part of PCE discovery [RFC4674], but may also be included in the PCECP message exchanges [RFC4657].

R14: The ability to perform P2MP path computation and the objective functions supported by a PCE SHOULD be advertised as part of PCE discovery. In the event that the PCE's ability to perform P2MP computation is not advertised as part of PCE discovery, the PCECP MUST allow a PCC to discover which PCEs with which it communicates support P2MP path computation, and which objective functions specific to P2MP path computation are supported by each PCE.

The list of objective functions is assumed to be coordinated with those that can be requested as described in Section 3.1.2.

These requirements do not represent a change to [RFC4657], except to add more capabilities and objective functions.

3.1.14. Path-Tree Diversity

Section 3.1.9 sets out the requirement to be able to request multiple diverse paths. Additionally, with P2MP trees, it may be that only parts of the tree can be, or need to be, diverse.

R15: The PCC SHOULD be able to request a PCE to compute a secondary P2MP path tree with partial path diversity for specific leaves or a specific S2L sub-path.

4. Manageability Considerations

4.1. Control of Function and Policy

PCE implementations MAY provide a configuration switch to allow support of P2MP MPLS TE computations to be enabled or disabled. When the level of support is changed, this SHOULD be re-advertised as described in Section 3.1.13.

Support for, and advertisement of support for, P2MP MPLS TE path computation MAY be subject to policy, and a PCE MAY hide its P2MP capabilities from certain PCCs by not advertising them through the discovery protocol and not reporting them to the specific PCCs in any PCECP capabilities exchange. Further, a PCE MAY be directed by policy to refuse a P2MP path computation for any reason including, but not limited to, the identity of the PCC that makes the request.

4.2. Information and Data Models

PCECP protocol extensions to support P2MP MPLS TE SHOULD be accompanied by MIB objects for the control and monitoring of the protocol and the PCE that performs the computations. The MIB objects MAY be provided in the same MIB module as used for general PCECP control and monitoring or MAY be provided in a new MIB module.

The MIB objects SHOULD provide the ability to control and monitor all aspects of PCECP relevant to P2MP MPLS TE path computation.

4.3. Liveness Detection and Monitoring

No changes are necessary to the liveness detection and monitoring requirements as already embodied in [RFC4657]. However, it should be noted that, in general, P2MP computations are likely to take longer than P2P computations. The liveness detection and monitoring features of the PCECP SHOULD take this into account.

4.4. Verifying Correct Operation

There are no additional requirements beyond those expressed in [RFC4657] for verifying the correct operation of the PCECP. 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 P2MP LSP, the topology graph must include information about the P2MP signaling and branching capabilities of each LSR in the network.

Whatever means is 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 be as described in [RFC5073].

4.6. Impact on Network Operation

The use of a PCE to compute P2MP paths is not expected to have significant impact on network operations. However, it should be noted that the introduction of P2MP support to a PCE that already provides P2P path computation might change the loading of the PCE significantly, and that might have an impact on the network behavior, especially during recovery periods immediately after a network failure.

5. Security Considerations

P2MP computation requests do not raise any additional security issues for the PCECP, as there are no new messages and no new PCC-PCE relationships or transactions introduced.

Note, however, that P2MP computation requests are more CPU-intensive and also use more link bandwidth. Therefore, if the PCECP was susceptible to denial of service attacks based on the injection of spurious Path Computation Requests, the support of P2MP path computation would exacerbate the effect.

It would be possible to consider applying different authorization policies for P2MP Path Computation Requests compared to other requests.

6. Acknowledgments

Thanks to Dean Cheng, Young Lee, Quintin Zhao, Daniel King, Fabien Verhaeghe, and Francis Dupont for their comments and suggestions on this document.

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.
- [RFC4657] Ash, J., Ed., and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, December 2008.
- [RFC5671] Yasukawa, S. and A. Farrel, Ed., "Applicability of the Path Computation Element (PCE) to Point-to-Multipoint (P2MP) MPLS and GMPLS Traffic Engineering (TE)", RFC 5671, October 2009.

7.2. Informative References

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC4461] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4674] Le Roux, J., Ed., "Requirements for Path Computation Element (PCE) Discovery", RFC 4674, October 2006.
- [RFC4834] Morin, T., Ed., "Requirements for Multicast in Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4834, April 2007.

- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5073] Vasseur, J., Ed., and J. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, December 2007.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", RFC 5541, June 2009.

Authors' Addresses

Seisho Yasukawa
NTT Corporation
9-11, Midori-Cho 3-Chome
Musashino-Shi, Tokyo 180-8585
JAPAN
EMail: yasukawa.seisho@lab.ntt.co.jp

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

