

Network Working Group
Internet Draft
Category: Informational
Created: November 14, 2007
Expires: May 14, 2008

S. Yasukawa
NTT
A. Farrel
Old Dog Consulting

PCC-PCE Communication Requirements for Point to Multipoint
Multiprotocol Label Switching Traffic Engineering (MPLS-TE)

[draft-yasukawa-pce-p2mp-req-04.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

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

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

Abstract

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multi-Protocol 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). Since P2MP TE LSP routes are sometimes complex to compute, and given the use of PCE in MPLS networks it is likely that PCE will be used in P2MP MPLS-TE networks.

Generic requirements for a communication protocol between Path Computation Clients (PCCs) and PCEs are presented in "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 traffic engineering.

[draft-yasukawa-pce-p2mp-req-04.txt](#)

November 2007

Conventions used in this document

Although this document is 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 description of requirements.

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 VPNs [[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 [[PCE-P2MP-APP](#)]. 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. PCC-PCE Communication Requirements for P2MP MPLS Traffic Engineering

This section sets out additional requirements not covered in [[RFC4657](#)] specific to P2MP MPLS TE.

[2.1. 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.

[draft-yasukawa-pce-p2mp-req-04.txt](#)

November 2007

[2.1.1. Indication of P2MP Path Computation Request](#)

Although the presence of certain parameters (such as a list of more than one destination) MAY be used 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.

[2.1.2. Indication of P2MP Objective Functions](#)

[RFC4657] includes the requirement to be able to specify the objective functions to be applied by a PCE during path computation and for the PCE to indicate which objective functions were applied.

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, and that these objective functions will need to be defined and allocated codepoints in a separate document.

[2.1.3. Non-Support of P2MP Path Computation](#)

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.

[2.1.4. 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 2.1.1](#). Such legacy implementations will not be able to make use of the new reason code described in [Section 2.1.3](#).

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 is the parameter described in [Section 2.1.1](#).

[2.1.5. Specification of Destinations](#)

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

[2.1.6. Indication of P2MP Paths](#)

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

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

By default, the path returned by the PCE SHOULD use the compacted format.

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

[2.1.7. Multi-Message Requests and Responses](#)

A single P2MP LSP may have very 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.

[2.1.8. 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.

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.

[2.1.9. Path Modification and Path Diversity](#)

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 on 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.

[2.1.10. Reoptimization of P2MP TE LSPs](#)

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 2.1.9](#).

Because P2MP LSPs are more complex it is often the case that small optimization improvements can be made after changes in network resource availability. But 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 [[PCE-POLICY](#)].

Path reoptimization responses SHOULD indicate which of the routes (as supplied according to [Section 2.1.6](#)) have been modified from the paths supplied on the request.

[2.1.11. Addition and Removal of Destinations from Existing Paths](#)

A variation of path modification described in [Section 2.1.9](#) 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.

To support this function for leaf additions it MUST be possible to make the following indications on a path computation request:

- The path of an existing P2MP LSP (as described in [Section 2.1.9](#)).
- 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 on a path computation request a cost-benefit reoptimization threshold such that the tree and/or a new path to any individual destination is not supplied unless a certain improvement is made. Compare with [Section 2.1.10](#).

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 2.1.10](#), 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) on 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 2.1.6](#)) have been modified from the paths supplied on the request as described in [Section 2.1.10](#).

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 [[PCE-POLICY](#)].

[2.1.12](#). 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. That is, 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.

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

[2.1.13. Capabilities Exchange](#)

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

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 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 2.1.2](#).

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

[3. Manageability Considerations](#)

[3.1. 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 2.1.13](#).

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.

[3.2. Information and Data Models](#)

PCECP protocol extensions to support P2MP MPLS TE MUST 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 MUST provide the ability to control and monitor all aspects of PCECP relevant to P2MP MPLS TE path computation.

[3.3. 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 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.

[3.4. 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.

[3.5. 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 MUST include the requisite information. If the TE extensions to the routing protocol are used, these SHOULD be as

described in [[TE-NODE-CAP](#)].

[3.6](#). Impact on Network Operation

The use of a PCE to compute P2MP paths is not expected to have significant impact on network operations. But 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.

[4](#). 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.

[5](#). IANA Considerations

This document makes no requests for IANA action.

[6](#). Acknowledgments

Thanks to Dean Cheng, Young Lee, and Quintin Zhao for their comments and suggestions on this document.

[7](#). References

[7.1](#). Normative Reference

[RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", [RFC 2119](#), March 1997.

- [RFC4657] Ash, J., and Le Roux, J.L., "Path Computation Element (PCE) Communication Protocol Generic Requirements", [RFC 4657](#), September 2006.
- [PCE-POLICY] Bryskin, I., Papadimitriou, D., and Berger, L., "Policy-Enabled Path Computation Framework", [draft-ietf-pce-policy-enabled-path-comp](#), work in progress.

7.2. Informative Reference

- [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] S. Yasukawa, Editor "Signaling Requirements for Point-to-Multipoint Traffic Engineered MPLS LSPs", [RFC4461](#), April 2006.
- [RFC4655] Farrel, A., Vasseur, J.P., and Ash, G., "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC4674] J.L. Le Roux, Editor, "Requirements for Path Computation Element (PCE) Discovery", [RFC 4674](#), October 2006.
- [PCE-P2MP-APP] S. Yasukawa et al., "Applicability of the Path Computation Element to Point-to-Multipoint Traffic Engineering", [draft-yasukawa-pce-p2mp-app](#), work in progress.
- [RFC4834] Morin, T., "Requirements for Multicast in Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", [RFC 4834](#), April 2007.
- [RFC4875] Aggarwal, R., Papadimitriou, D., and Yasukawa, S., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 4875](#), May 2007.
- [TE-NODE-CAP] Vasseur, J.P, and Le Roux, J.L., Editors, "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", [draft-ietf-ccamp-te-node-cap](#), work in progress.

[draft-yasukawa-pce-p2mp-req-04.txt](#)

November 2007

8. Authors' Addresses

Seisho Yasukawa
NTT Corporation (R&D Strategy Department)
3-1, Otemachi 2-Chome Chiyodaku, Tokyo 100-8116 Japan
Email: s.yasukawa@hco.ntt.co.jp

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

9. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

10. Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an

"AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.