

PCE Working Group
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

D. Dhody
X. Zhang
Huawei Technologies
July 4, 2014

Stateful Path Computation Element (PCE) Inter-domain Considerations
draft-dhody-pce-stateful-pce-interdomain-00

Abstract

A stateful Path Computation Element (PCE) maintains information about Label Switched Path (LSP) characteristics and resource usage within a network in order to provide traffic engineering path calculations for its associated Path Computation Clients (PCCs). Furthermore, PCEs are used to compute shortest constrained Traffic Engineering Label Switched Paths (TE LSPs) in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks across multiple domains.

This document describes general considerations for the deployment of stateful PCE(s) in inter-domain scenarios including inter-area and inter-AS.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 5, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	2
2. Overview	3
2.1. LSP State Synchronization	4
3. Stateful PCE Deployments	4
3.1. Single Stateful PCE, Multiple Domains	4
3.2. Multiple Stateful PCE, Multiple Domains	5
3.2.1. Per Domain Path Computation	6
3.2.2. Backward-Recursive PCE-based Computation	7
3.2.3. Hierarchical PCE	7
4. Other Considerations	8
4.1. Delegation	8
5. Security Considerations	9
6. Manageability Considerations	9
6.1. Control of Function and Policy	9
6.2. Information and Data Models	9
6.3. Liveness Detection and Monitoring	9
6.4. Verify Correct Operations	9
6.5. Requirements On Other Protocols	9
6.6. Impact On Network Operations	10
7. IANA Considerations	10
8. Acknowledgments	10
9. References	10
9.1. Normative References	10
9.2. Informative References	10
Appendix A. Contributor Addresses	12

1. Introduction

The Path Computation Element communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients' (PCCs) requests.

[I-D.ietf-pce-stateful-pce-app] describes general considerations for a stateful PCE deployment and examines its applicability and benefits, as well as its challenges and limitations through a number of use cases. [I-D.ietf-pce-stateful-pce] describes a set of extensions to PCEP to provide stateful control. A stateful PCE has access to not only the information carried by the network's Interior

Gateway Protocol (IGP), but also the set of active paths and their reserved resources for its computations. The additional state allows the PCE to compute constrained paths while considering individual LSPs and their interactions.

The ability to compute shortest constrained TE LSPs in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks across multiple domains has been identified as a key motivation for PCE development. In this context, a domain is a collection of network elements within a common sphere of address management or path computational responsibility such as an Interior Gateway Protocol (IGP) area or an Autonomous Systems (AS).

This document presents general considerations for stateful PCE(s) deployment in multi-domain scenarios.

2. Overview

A stateful PCE maintains two sets of information for use in path computation. The first is the Traffic Engineering Database (TED) which includes the topology and resource state in the network. The second is the LSP State Database (LSP-DB), in which a PCE stores attributes of all active LSPs in the network, such as their paths through the network, bandwidth/resource usage, switching types and LSP constraints. This state information allows the PCE to compute constrained paths while considering individual LSPs and their inter-dependency. [I-D.ietf-pce-stateful-pce] applies equally to MPLS-TE and GMPLS LSPs and distinguishes between an active and a passive stateful PCE. A passive stateful PCE uses LSP state information to optimize path computations but does not actively update LSP state. In contrast, an active stateful PCE may issue recommendations to the network. For example, an active stateful PCE may update LSP parameters for those LSPs that have been delegated, by its PCCs, the control over to the PCE.

The capability to compute the routes of end-to-end inter-domain MPLS-TE LSPs is expressed as requirements in [RFC4105] and [RFC4216] and may be realized by PCE(s). PCEs may use one of the following mechanisms to compute end-to-end paths:

- o a per-domain path computation technique [RFC5152];
- o a Backward-Recursive PCE-based Computation (BRPC) mechanism [RFC5441];
- o a Hierarchical PCE mechanism [RFC6805];

This document examines the stateful PCE inter-domain considerations for all of these mechanisms.

2.1. LSP State Synchronization

The population of the LSP-DB using information received from PCCs (ingress LSR) is supported by the stateful PCE extensions defined in [I-D.ietf-pce-stateful-pce] , i.e., via LSP state report messages.

The inter-domain LSP state is synchronised to the ingress-PCE from the ingress LSR (PCC), but this PCC cannot synchronise to other PCEs (in transit or egress domains), thus other mechanism must be investigated for this purpose.

3. Stateful PCE Deployments

There are multiple models to perform PCE-based inter-domain path computation:


- o A single PCE;
- o Multiple PCE;
 - * without inter-PCE communication;
 - * with inter-PCE communication;

This section describe stateful PCE considerations for each of these deployment models.

3.1. Single Stateful PCE, Multiple Domains

In this model, inter-domain path computation is performed by a single stateful PCE that has topology visibility into all domains. The inter-domain LSP state is synchronised to this PCE from the ingress LSR (PCC) itself. This PCC may also choose to delegate control over this LSP to the PCE. Thus this model is similar to a single domain in all aspects.

Following figure show an example of inter-area case comprising of Area 0,1 and 2. A single stateful PCE is deployed for all areas.




In this model PCE has visibility into the topology of all domains as well as the state of all active LSPs including inter-domain LSPs. This model is thus well suited to take advantage of all stateful PCE capabilities.

It should be noted that a single PCE may not be possible because of administrative and confidentiality concerns.

3.2. Multiple Stateful PCE, Multiple Domains

In this model, there is at least one PCE per domain, and each PCE has topology visibility restricted to its own domain. The inter-domain LSP state is synchronised to the ingress-PCE from the ingress LSR (PCC), but this PCC cannot synchronise to other PCEs (in transit or egress domains). This PCC may also choose to delegate control over this LSP to the Ingress-PCE, which may issue inter-domain path computation or re-optimization request to other PCEs. An inter-domain LSP that originates in the domain, is synchronised to the PCE in that domain. But a mechanism is needed to synchronize state of inter-domain LSP that do not originate in the domain. In other words, inter-domain LSP state should also be synchronised to transit and egress PCEs.

Following figure show an example of inter-AS case comprising of AS 100 and AS 200. A stateful PCE is deployed per AS.



In order to conceal the information, a PCE may use path-key based confidentiality mechanisms as per [RFC5520].

This section further describes considerations with respect to each of the inter-domain path computation techniques.

3.2.1. Per Domain Path Computation

The per domain path computation technique [RFC5152] is based on Multiple PCE Path Computation without Inter-PCE Communication Model as described in [RFC4655]. It defines a method where the path is computed during the signaling process (on a per-domain basis). The entry Boundary Node (BN) of each domain is responsible for performing the path computation for the section of the LSP that crosses the domain, or for requesting that a PCE for that domain computes that piece of the path.

The ingress LSR would synchronise the the state to the ingress PCE, further the entry boundary nodes should synchronize the state of inter-domain LSP to transit and egress PCEs. Note that the BN on the path of an LSP can probably see the path (through the Record Route object in RSVP-TE signaling [RFC3209]) and knows the bandwidth reserved for the LSP. Thus each entry BN along the path could be made responsible to synchronise the LSP state to the transit/egress PCE(s).

Since the stateful PCE(s) do not communicate during this inter-domain path computation technique and each entry BN would perform path computation via Path Computation Request (PCReq) and Reply (PCRep) messages, a passive stateful PCE is well suited for this case.

In case of delegation to the ingress PCE (active stateful PCE), it would be capable of loose path computation only and make updates to the ingress LSR with this limited visibility. The entry BN would perform path computation via Path Computation Request and Reply

messages (and thus rely on the passive stateful mode). Thus the inter-domain LSP is delegated only to the ingress PCE.

3.2.2. Backward-Recursive PCE-based Computation

The BRPC [RFC5441] technique is based on Multiple PCE Path Computation with Inter-PCE Communication Model as described in [RFC4655]. It involves cooperation and communication between PCEs in order to compute an optimal end-to-end path across multiple domains. The sequence of domains to be traversed may be known before the path computation, but it can also be used when the domain path is unknown and determined during path computation.

As described in Section 3.2.1, the entry boundary nodes may synchronize the state of inter-domain LSPs to transit and egress PCEs. An alternative approach may be for each PCE to synchronise the state along the path across domains, i.e., each PCE would report the state to the next PCE(s) in the adjacent domain along the domain sequence of the inter-domain path. A mechanism similar to LSP-DB backup [I-D.palle-pce-stateful-pce-lspdb-sync] may be utilized for this purpose.

Some path segment in the end to end path may also be hidden via path-key as per [RFC5520] during state synchronization.

In case of passive path computation request to the ingress PCE from the ingress LSR the BRPC path computation procedure is applied to compute end-to-end path by using PCReq and PCRep messages among stateful PCE(s) in passive mode.

In case of delegation to the ingress PCE (active stateful PCE), the ingress PCE may trigger the end-to-end path computation via the same BRPC procedure using the path computation request and reply messages among stateful PCE(s) in passive mode. For re-optimization or update the ingress PCE still rely on the same BRPC procedure triggered by the ingress PCE. Ultimately the inter-domain LSP is delegated to the ingress PCE and only the ingress PCE can issue updates to the inter-domain LSP. It may trigger E2E path re-optimization with help of transit/egress PCE using the BRPC procedure.

3.2.3. Hierarchical PCE

In H-PCE [RFC6805] architecture, the parent PCE is used to compute a multi-domain path based on the domain connectivity information. The parent PCE may be requested to provide a end-to-end path or only the sequence of domains.

As described in Section 3.2.1 and Section 3.2.2, the entry boundary nodes may synchronize the state of inter-domain LSP to transit and egress child PCEs. If the parent PCE provides the sequence of domains and BRPC procedure is used to get the E2E path, each PCE may be responsible to synchronise the state along the path across domains similar to Section 3.2.2. An alternative approach may be for ingress PCE to synchronise LSP state with the Parent PCE and it may further synchronise the state to the child PCE(s) along the path across domains, i.e. parent PCE would report the state to the child PCE(s) along the domain sequence.

Some path segment in the end to end path may also be hidden via path-key as per [RFC5520] during state synchronization.

In case of passive path computation request to the ingress PCE from the ingress LSR, the H-PCE path computation procedure is applied to compute sequence of domains or end-to-end path by using PCReq and PCRep messages among stateful PCE(s) in passive mode.

In case of delegation to the ingress PCE (active stateful PCE), the ingress PCE may trigger the H-PCE path computation via the same procedure using the PCReq and PCRep messages among stateful PCE(s) in passive mode. For re-optimization or update the ingress PCE still rely on the same H-PCE procedure triggered by the ingress PCE. Ultimately the inter-domain LSP is delegated to the ingress PCE and only the ingress PCE can issue updates to the inter-domain LSP. It may trigger E2E path re-optimization with help of parent and child PCEs using the H-PCE procedure.

4. Other Considerations

4.1. Delegation

As noted in this document, the inter-domain LSP is delegated to the ingress PCE and only the ingress PCE can issue updates to the inter-domain LSP. The ingress PCE is responsible to trigger E2E path re-optimization.

Thus the ingress PCE can recommend updation for all aspects of the inter-domain LSP including the segment of path in another domain (which it may have computed with the help of other cooperating PCEs). These interaction between PCEs for the inter-domain path computation are done using PCReq/PCRep messages (i.e., in a passive mode).

The transit/egress PCE cannot update any attribute of the inter-domain LSP on its own as it may not have any interaction with the ingress LSR. A mechanism may be developed for transit/egress PCE to inform the ingress PCE to trigger E2E re-optimization and choose to

update the inter-domain LSP based on the result. Also the ingress PCE may use combination of local information and events along with some external mechanism (management / monitoring interface) to trigger E2E path re-optimization.

Though Ingress PCE can recommend update for path segments in other domains, the entry boundary node of that domain can apply policy control during signalling as explained in [RFC4105] and [RFC4216].

5. Security Considerations

The security considerations are as per [RFC5440] and [I-D.ietf-pce-stateful-pce]. Any multi-domain operation necessarily involves the exchange of information across domain boundaries. This may represent a significant security and confidentiality risk especially when the domains are controlled by different commercial entities. PCEP allows individual PCEs to maintain confidentiality of their domain path information by using path-keys [RFC5520].

6. Manageability Considerations

6.1. Control of Function and Policy

Mechanisms defined in this document do not imply any new control of function and policy requirements.

6.2. Information and Data Models

[I-D.ietf-pce-pcep-mib] describes the PCEP MIB, there are no new MIB Objects for this document.

6.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

6.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440].

6.5. Requirements On Other Protocols

Mechanisms defined in this document do not imply any new requirements on other protocols.

6.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [RFC5440].

7. IANA Considerations

This is an informational document and has no IANA considerations.

8. Acknowledgments

TBD.

9. References

9.1. Normative References

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-09 (work in progress), June 2014.

9.2. 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.

[RFC4105] Le Roux, J., Vasseur, J., and J. Boyle, "Requirements for Inter-Area MPLS Traffic Engineering", RFC 4105, June 2005.

[RFC4216] Zhang, R. and J. Vasseur, "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005.

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

[RFC5152] Vasseur, JP., Ayyangar, A., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.

- [RFC5441] Vasseur, JP., 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.
- [RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.
- [RFC6805] King, D. and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, November 2012.
- [I-D.ietf-pce-stateful-pce-app]
Zhang, X. and I. Minei, "Applicability of a Stateful Path Computation Element (PCE)", draft-ietf-pce-stateful-pce-app-02 (work in progress), June 2014.
- [I-D.ietf-pce-pcep-mib]
Koushik, K., Emile, S., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Protocol (PCEP) Management Information Base", draft-ietf-pce-pcep-mib-08 (work in progress), April 2014.
- [I-D.palle-pce-stateful-pce-lspdb-sync]
Palle, U., Dhody, D., and X. Zhang, "LSP-DB Synchronization between Stateful PCEs", draft-palle-pce-stateful-pce-lspdb-sync-02 (work in progress), January 2014.

Appendix A. Contributor Addresses

Udayasree Palle
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: udayasree.palle@huawei.com

Avantika
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: avantika.sushilkumar@huawei.com

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruv.ietf@gmail.com

Xian Zhang
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R.China

EMail: zhang.xian@huawei.com