

Internet Engineering Task Force (IETF)
Request for Comments: 8637
Category: Informational
ISSN: 2070-1721

D. Dhody
Huawei Technologies
Y. Lee
Futurewei Technologies
D. Ceccarelli
Ericsson
July 2019

Applicability of the Path Computation Element (PCE)
to the Abstraction and Control of TE Networks (ACTN)

Abstract

Abstraction and Control of TE Networks (ACTN) refers to the set of virtual network (VN) operations needed to orchestrate, control, and manage large-scale multidomain TE networks so as to facilitate network programmability, automation, efficient resource sharing, and end-to-end virtual service-aware connectivity and network function virtualization services.

The Path Computation Element (PCE) is a component, application, or network node that is capable of computing a network path or route based on a network graph and applying computational constraints. The PCE serves requests from Path Computation Clients (PCCs) that communicate with it over a local API or using the Path Computation Element Communication Protocol (PCEP).

This document examines the applicability of PCE to the ACTN framework.

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 candidates for any level of Internet Standard; see 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/rfc8637>.

Copyright Notice

Copyright (c) 2019 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 2
2. Background Information 3
2.1. Path Computation Element (PCE) 3
2.1.1. Role of PCE in SDN 4
2.1.2. PCE in Multidomain and Multilayer Deployments 4
2.1.3. Relationship to PCE-Based Central Control 5
2.2. Abstraction and Control of TE Networks (ACTN) 5
3. Architectural Considerations 7
3.1. Multidomain Coordination via Hierarchy 7
3.2. Abstraction 8
3.3. Customer Mapping 9
3.4. Virtual Service Coordination 10
4. Interface Considerations 10
5. Realizing ACTN with PCE (and PCEP) 11
6. IANA Considerations 14
7. Security Considerations 15
8. References 16
8.1. Normative References 16
8.2. Informative References 16
Appendix A. Additional Information 21
Acknowledgments 21
Authors' Addresses 22

1. Introduction

Abstraction and Control of TE Networks (ACTN) [RFC8453] refers to the set of virtual network (VN) operations needed to orchestrate, control, and manage large-scale multidomain TE networks so as to facilitate network programmability, automation, efficient resource sharing, and end-to-end virtual service-aware connectivity and network function virtualization services.

The Path Computation Element (PCE) [RFC4655] is a component, application, or network node that is capable of computing a network path or route based on a network graph and applying computational constraints. The PCE serves requests from Path Computation Clients (PCCs) that communicate with it over a local API or using the Path Computation Element Communication Protocol (PCEP).

This document examines the PCE and ACTN architecture and describes how PCE architecture is applicable to ACTN. It also lists the PCEP extensions that are needed to use PCEP as an ACTN interface. This document also identifies any gaps in PCEP that exist at the time of publication of this document.

Further, ACTN, stateful Hierarchical PCE (H-PCE) [PCE-HPCE], and PCE as a central controller (PCECC) [RFC8283] are based on the same basic hierarchy framework and are thus compatible with each other.

2. Background Information

2.1. Path Computation Element (PCE)

The Path Computation Element Communication Protocol (PCEP) [RFC5440] provides mechanisms for Path Computation Clients (PCCs) to request a Path Computation Element (PCE) [RFC4655] to perform path computations.

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.

A stateful PCE [RFC8231] is capable of considering, for the purposes of path computation, not only the network state in terms of links and nodes (referred to as the Traffic Engineering Database or TED), but also the status of active services (previously computed paths), and currently reserved resources, stored in the Label Switched Paths Database (LSP-DB).

[RFC8051] 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.

[RFC8231] 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. [RFC8281] describes the setup, maintenance, and teardown of PCE-initiated LSPs under the stateful PCE model.

[RFC8231] also describes the active stateful PCE. The active PCE functionality allows a PCE to reroute an existing LSP or make changes to the attributes of an existing LSP, or a PCC to delegate control of specific LSPs to a new PCE.

2.1.1. Role of PCE in SDN

Software-Defined Networking (SDN) [RFC7149] refers to a separation between the control elements and the forwarding components so that software running in a centralized system, called a controller, can act to program the devices in the network to behave in specific ways. A required element in an SDN architecture is a component that plans how the network resources will be used and how the devices will be programmed. It is possible to view this component as performing specific computations to place flows within the network given knowledge of the availability of network resources, how other forwarding devices are programmed, and the way that other flows are routed. It is concluded in [RFC7399] that this is the same function that a PCE might offer in a network operated using a dynamic control plane. This is the function and purpose of a PCE, and the way that a PCE integrates into a wider network control system including SDN is presented in Application-Based Network Operation (ABNO) [RFC7491].

2.1.2. PCE in Multidomain and Multilayer Deployments

Computing paths across large multidomain environments requires special computational components and cooperation between entities in different domains capable of complex path computation. The PCE provides an architecture and a set of functional components to address this problem space. A PCE may be used to compute end-to-end paths across multidomain environments using a per-domain path computation technique [RFC5152]. The Backward-Recursive PCE-based path computation (BRPC) mechanism [RFC5441] defines a PCE-based path computation procedure to compute interdomain-constrained MPLS and GMPLS TE networks. However, per-domain technique assumes that the sequence of domains to be crossed from source to destination is known, either fixed by the network operator or obtained by other means. BRPC can work best with a known domain sequence, and it will also work nicely with a small set of interconnected domains. However, it doesn't work well for a large set of interconnected domains.

[RFC6805] describes a Hierarchical PCE (H-PCE) architecture that can be used for computing end-to-end paths for interdomain MPLS Traffic Engineering (TE) and GMPLS Label Switched Paths (LSPs) when the domain sequence is not known. Within the Hierarchical PCE (H-PCE) architecture, the Parent PCE (P-PCE) is used to compute a multidomain path based on the domain connectivity information. A Child PCE (C-PCE) may be responsible for a single domain or multiple domains; it is used to compute the intradomain path based on its domain topology information.

[PCE-HPCE] states the considerations for stateful PCEs in Hierarchical PCE architecture. In particular, the behavior changes and adds to the existing stateful PCE mechanisms (including PCE-initiated LSP setup and active PCE usage) in the context of networks using the H-PCE architecture.

[RFC5623] describes a framework for applying the PCE-based architecture to interlayer (G)MPLS TE. It provides suggestions for the deployment of PCE in support of multilayer networks. It also describes the relationship between PCE and a functional component in charge of the control and management of the Virtual Network Topology (VNT) [RFC5212] called the VNT Manager (VNTM).

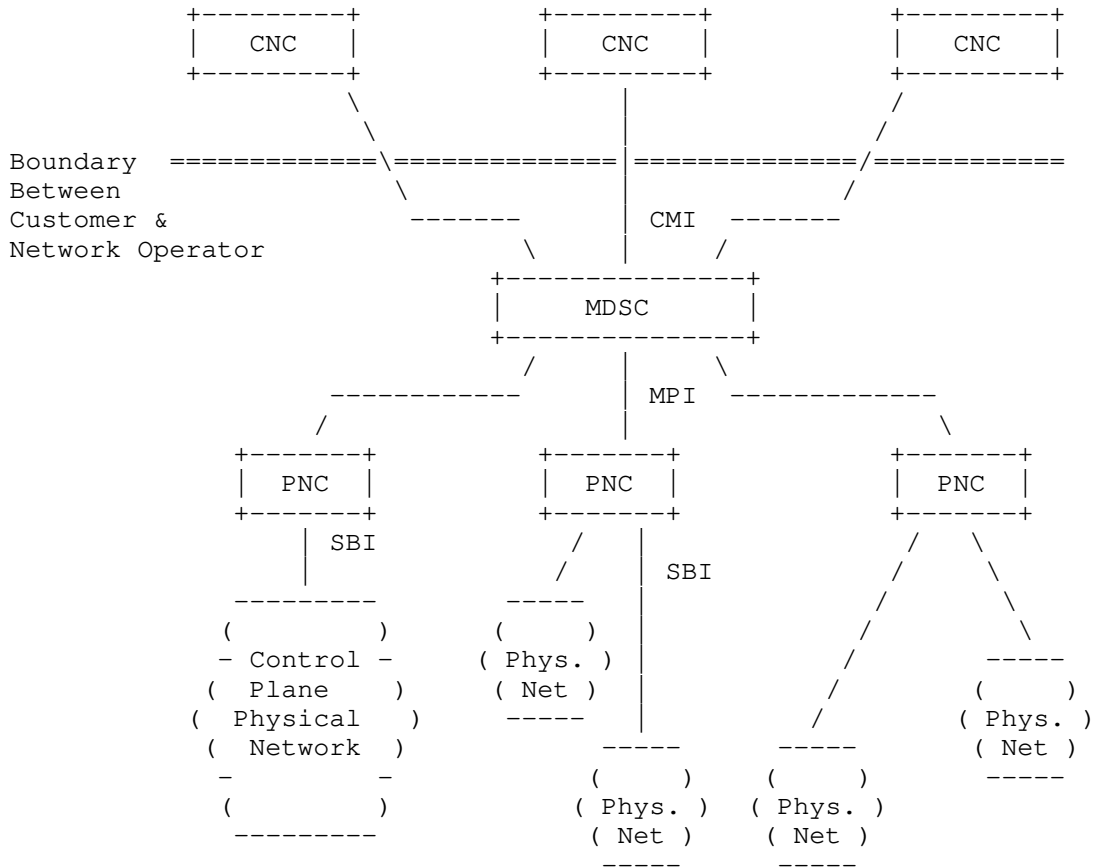
2.1.3. Relationship to PCE-Based Central Control

[RFC8283] introduces the architecture for PCE as a central controller (PCECC); it further examines the motivations and applicability for PCEP as a southbound interface (SBI) and introduces the implications for the protocol. Section 2.1.3 of [RFC8283] describes a hierarchy of PCE-based controllers as per the PCE Hierarchy Framework defined in [RFC6805].

2.2. Abstraction and Control of TE Networks (ACTN)

[RFC8453] describes the high-level ACTN requirements and the architecture model for ACTN, including the entities Customer Network Controller (CNC), Multidomain Service Coordinator (MDSC), and Provisioning Network Controller (PNC) and their interfaces.

The ACTN reference architecture is shown in Figure 1, which is reproduced here from [RFC8453] for convenience. [RFC8453] remains the definitive reference for the ACTN architecture. As depicted in Figure 1, the ACTN architecture identifies a three-tier hierarchy.



CMI - (CNC-MDSC Interface)
 MPI - (MDSC-PNC Interface)
 SBI - (Southbound Interface)

Figure 1: ACTN Hierarchy

There are two interfaces with respect to the MDSC: one north of the MDSC (the CNC-MDSC Interface (CMI)), and one south (the MDSC-PNC Interface (MPI)). A hierarchy of MDSCs is possible with a recursive MPI interface.

[RFC8454] provides an information model for ACTN interfaces.

3. Architectural Considerations

The ACTN architecture [RFC8453] is based on the hierarchy and recursiveness of controllers. It defines three types of controllers (depending on the functionalities they implement). The main functionalities are:

- o Multidomain coordination
- o Abstraction
- o Customer mapping/translation
- o Virtual service coordination

Section 3 of [RFC8453] describes these functions.

It should be noted that this document lists all possible ways in which PCE could be used for each of the above functions, but all functions are not required to be implemented via PCE. Similarly, this document presents the ways in which PCEP could be used as the communications medium between functional components. Operators may choose to use the PCEP for multidomain coordination via stateful H-PCE but alternatively use Network Configuration Protocol (NETCONF) [RFC6241], RESTCONF [RFC8040], or BGP - Link State (BGP-LS) [RFC7752] to get access to the topology and support abstraction function.

3.1. Multidomain Coordination via Hierarchy

With the definition of domain being everything that is under the control of the single logical controller, as per [RFC8453], it is needed both to have a control entity that oversees the specific aspects of the different domains and to build a single abstracted end-to-end network topology in order to coordinate end-to-end path computation and path/service provisioning.

The MDSC in ACTN framework realizes this function by coordinating the per-domain PNCs in a hierarchy of controllers. It also needs to detach from the underlying network technology and express customer concerns by business needs.

[RFC6805] and [PCE-HPCE] describe a hierarchy of PCEs with the Parent PCE coordinating multidomain path computation function between Child PCEs. It is easy to see how these principles align, and thus how the stateful H-PCE architecture can be used to realize ACTN.

The per-domain stitched LSP in the Hierarchical stateful PCE architecture, described in Section 3.3.1 of [PCE-HPCE], is well suited for multidomain coordination function. This includes domain sequence selection, End-to-End (E2E) path computation, and controller-initiated (PCE-initiated) path setup and reporting. This is also applicable to multilayer coordination in case of IP+optical networks.

[PCE-STATE-SYNC] describes the procedures to allow a stateful communication between PCEs for various use cases. The procedures and extensions are also applicable to Child and Parent PCE communication and are thus useful for ACTN as well.

3.2. Abstraction

To realize ACTN, an abstracted view of the underlying network resources needs to be built. This includes global network-wide abstracted topology based on the underlying network resources of each domain. This also includes abstract topology created as per the customer service connectivity requests and represented as a VN slice allocated to each customer.

In order to compute and provide optimal paths, PCEs require an accurate and timely Traffic Engineering Database (TED). Traditionally, this TED has been obtained from a link-state (LS) routing protocol supporting traffic engineering extensions. PCE may construct its TED by participating in the IGP ([RFC3630] and [RFC5305] for MPLS-TE; [RFC4203] and [RFC5307] for GMPLS). An alternative is offered by BGP-LS [RFC7752].

In case of H-PCE [RFC6805], the Parent PCE needs to build the domain topology map of the child domains and their interconnectivity. [RFC6805] and [PCE-INTER-AREA] suggest that BGP-LS could be used as a "northbound" TE advertisement from the Child PCE to the Parent PCE.

[PCEP-LS] proposes another approach for learning and maintaining the Link-State and TE information as an alternative to IGPs and BGP flooding, using PCEP itself. The Child PCE can use this mechanism to transport Link-State and TE information from Child PCE to a Parent PCE using PCEP.

In ACTN, there is a need to control the level of abstraction based on the deployment scenario and business relationship between the controllers. The mechanism used to disseminate information from the PNC (Child PCE) to the MDSC (Parent PCE) should support abstraction. [RFC8453] describes a few alternative approaches of abstraction. The resulting abstracted topology can be encoded using the PCEP-LS mechanisms [PCEP-LS] and its optical network extension

[PCEP-OPTICAL]. PCEP-LS is an attractive option when the operator would wish to have a single control-plane protocol (PCEP) to achieve ACTN functions.

[RFC8453] discusses two ways to build abstract topology from an MDSC standpoint with interaction with PNCs. The primary method is called automatic generation of abstract topology by configuration. With this method, automatic generation is based on the abstraction/summarization of the whole domain by the PNC and its advertisement on the MPI. The secondary method is called on-demand generation of supplementary topology via Path Compute Request/Reply. This method may be needed to obtain further complementary information such as potential connectivity from Child PCEs in order to facilitate an end-to-end path provisioning. PCEP is well suited to support both methods.

3.3. Customer Mapping

In ACTN, there is a need to map customer virtual network (VN) requirements into a network provisioning request to the PNC. That is, the customer requests/commands are mapped by the MDSC into network provisioning requests that can be sent to the PNC. Specifically, the MDSC provides mapping and translation of a customer's service request into a set of parameters that are specific to a network type and technology such that the network configuration process is made possible.

[RFC8281] describes the setup, maintenance, and teardown of PCE-initiated LSPs under the stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network that is centrally controlled and deployed. To instantiate or delete an LSP, the PCE sends the Path Computation LSP Initiate Request (PCInitiate) message to the PCC. As described in [PCE-HPCE], for interdomain LSP in Hierarchical PCE architecture, the initiation operations can be carried out at the Parent PCE. In this case, after Parent PCE finishes the E2E path computation, it can send the PCInitiate message to the Child PCE; the Child PCE further propagates the initiate request to the Label Switching Router (LSR). The customer request is received by the MDSC (Parent PCE), and, based on the business logic, global abstracted topology, network conditions, and local policy, the MDSC (Parent PCE) translates this into a per-domain LSP initiation request that a PNC (Child PCE) can understand and act on. This can be done via the PCInitiate message.

PCEP extensions for associating opaque policy between PCEP peer [ASSOC-POLICY] can be used.

3.4. Virtual Service Coordination

Virtual service coordination function in ACTN incorporates customer service-related information into the virtual network service operations in order to seamlessly operate virtual networks while meeting customers' service requirements.

[PCEP-VN] describes the need for associating a set of LSPs with a VN "construct" to facilitate VN operations in PCE architecture. This association allows the PCEs to identify which LSPs belong to a certain VN.

This association based on VN is useful for various optimizations at the VN level, which can be applied to all the LSPs that are part of the VN slice. During path computation, the impact of a path for an LSP is compared against the paths of other LSPs in the VN. This is to ensure optimization and SLA attainment for the VN rather than for a single LSP. Similarly, during reoptimization, advanced path computation algorithms and optimization techniques can be considered for all the LSPs belonging to a VN/customer and optimize them all together.

4. Interface Considerations

As per [RFC8453], to allow virtualization and multidomain coordination, the network has to provide open, programmable interfaces in which customer applications can create, replace, and modify virtual network resources and services in an interactive, flexible, and dynamic fashion while having no impact on other customers. The two ACTN interfaces are as follows:

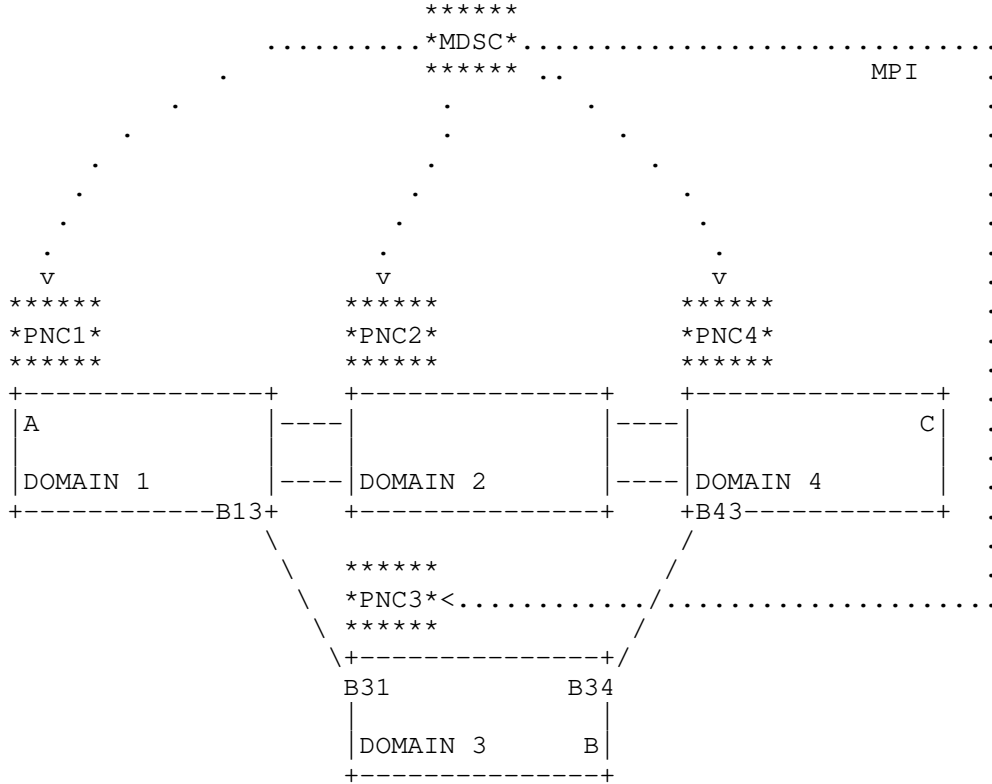
- o The CNC-MDSC Interface (CMI) is an interface between a Customer Network Controller and a Multidomain Service Coordinator. It requests the creation of the network resources, topology, or services for the applications. The MDSC may also report potential network topology availability if queried for current capability from the Customer Network Controller.
- o The MDSC-PNC Interface (MPI) is an interface between a Multidomain Service Coordinator and a Provisioning Network Controller. It communicates the creation request, if required, of new connectivity of bandwidth changes in the physical network via the PNC. In multidomain environments, the MDSC needs to establish multiple MPIs, one for each PNC, as there are multiple PNCs responsible for its domain control.

In the case of a hierarchy of MDSCs, the MPI is applied recursively. From an abstraction point of view, the top-level MDSC, which interfaces the CNC, operates on a higher level of abstraction (i.e., less granular level) than the lower-level MDSCs.

PCEP is especially suitable on the MPI as it meets the requirement and the functions as set out in the ACTN framework [RFC8453]. Its recursive nature is well suited via the multilevel hierarchy of PCE. PCEP can also be applied to the CMI as the CNC can be a path computation client while the MDSC can be a path computation server. Section 5 describes how PCE and PCEP could help realize ACTN on the MPI.

5. Realizing ACTN with PCE (and PCEP)

As per the example in Figure 2, there are 4 domains, each with their own PNC and MDSC on top. The PNC and MDSC need PCE as an important function. The PNC (or Child PCE) already uses PCEP to communicate to the network device. It can utilize the PCEP as the MPI to communicate between controllers too.



```

MDSC -> Parent PCE
PNC  -> Child PCE
MPI  -> PCEP
  
```

Figure 2: ACTN with PCE

- o Building Domain Topology at MDSC: PNC (or Child PCE) needs to have the TED to compute the path in its domain. As described in Section 3.2, it can learn the topology via IGP or BGP-LS. PCEP-LS is also a proposed mechanism to carry link state and traffic engineering information within PCEP. A mechanism to carry abstracted topology while hiding technology-specific information between PNC and MDSC is described in [PCEP-LS]. At the end of this step, the MDSC (or Parent PCE) has the abstracted topology from each of its PNCs (or Child PCE). This could be as simple as a domain topology map as described in [RFC6805], or it can have full topology information of all domains. The latter is not scalable, and thus, an abstracted topology of each domain interconnected by interdomain links is the most common case.

- * **Topology Change:** When the PNC learns of any topology change, the PNC needs to decide if the change needs to be notified to the MDSC. This is dependent on the level of abstraction between the MDSC and the PNC.
- o **VN Instantiate:** When an MDSC is requested to instantiate a VN, the minimal information that is required would be a VN identifier and a set of end points. Various path computation, setup constraints, and objective functions may also be provided. In PCE terms, a VN Instantiate can be considered as a set of paths belonging to the same VN. As described in Section 3.4 and [PCEP-VN], the VN association can help in identifying the set of paths that belong to a VN. The rest of the information, like the endpoints, constraints, and objective function (OF), is already defined in PCEP in terms of a single path.
- * **Path Computation:** As per the example in Figure 2, the VN instantiate requires two end-to-end paths between (A in Domain 1 to B in Domain 3) and (A in Domain 1 to C in Domain 4). The MDSC (or Parent PCE) triggers the end-to-end path computation for these two paths. MDSC can do path computation based on the abstracted domain topology that it already has, or it may use the H-PCE procedures (Section 3.1) using the PCReq and PCRep messages to get the end-to-end path with the help of the Child PCEs (PNC). Either way, the resultant E2E paths may be broken into per-domain paths.
- * **A-B:** (A-B13,B13-B31,B31-B)
- * **A-C:** (A-B13,B13-B31,B31-B34,B34-B43,B43-C)
- * **Per-Domain Path Instantiation:** Based on the above path computation, MDSC can issue the path instantiation request to each PNC via PCInitiate message (see [PCE-HPCE] and [PCEP-VN]). A suitable stitching mechanism would be used to stitch these per-domain LSPs. One such mechanism is described in [PCE-INTERDOMAIN], where PCEP is extended to support stitching in stateful H-PCE context.
- * **Per-Domain Path Report:** Each PNC should report the status of the per-domain LSP to the MDSC via PCRpt message, as per the hierarchy of stateful PCEs ([PCE-HPCE]). The status of the end-to-end LSP (A-B and A-C) is made up when all the per-domain LSPs are reported up by the PNCs.
- * **Delegation:** It is suggested that the per-domain LSPs are delegated to respective PNCs so that they can control the path and attributes based on the conditions of each domain network.

- * State Synchronization: The state needs to be synchronized between the Parent PCE and Child PCE. The mechanism described in [PCE-STATE-SYNC] can be used.
- o VN Modify: MDSC is requested to modify a VN, for example, the bandwidth for VN is increased. This may trigger path computation at MDSC as described in the previous step and can trigger an update to an existing per-intradomain path (via PCUpd message) or the creation (or deletion) of a per-domain path (via PCInitiate message). As described in [PCE-HPCE], this should be done in make-before-break fashion.
- o VN Delete: MDSC is requested to delete a VN, in this case, based on the E2E paths, and the resulting per-domain paths need to be removed (via PCInitiate message).
- o VN Update (based on network changes): Any change in the per-domain LSP is reported to the MDSC (via PCRpt message) as per [PCE-HPCE]. This may result in changes in the E2E path or VN status. This may also trigger a reoptimization leading to a new per-domain path, an update to an existing path, or the deletion of the path.
- o VN Protection: The VN protection/restoration requirements need to be applied to each E2E path as well as each per-domain path. The MDSC needs to play a crucial role in coordinating the right protection/restoration policy across each PNC. The existing protection/restoration mechanism of PCEP can be applied on each path.
- o In case a PNC generates an abstract topology towards the MDSC, the PCInitiate/PCUpd messages from the MDSC to a PNC will contain a path with abstract nodes and links. A PNC would need to take that as an input for path computation to get a path with physical nodes and links. Similarly, a PNC would convert the path received from the device (with physical nodes and links) into an abstract path (based on the abstract topology generated before with abstract nodes and links) and report it to the MDSC.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

Various security considerations for PCEP are described in [RFC5440] and [RFC8253]. Security considerations as stated in Sections 10.1, 10.6, and 10.7 of [RFC5440] continue to apply on PCEP when used as an ACTN interface. Further, this document lists various extensions of PCEP that are applicable; each of them specify various security considerations that continue to apply here.

The ACTN framework described in [RFC8453] defines key components and interfaces for managed traffic-engineered networks. It also lists various security considerations such as request and control of resources, confidentiality of the information, and availability of function, which should be taken into consideration.

As per [RFC8453], securing the request and control of resources, confidentiality of the information, and availability of function should all be critical security considerations when deploying and operating ACTN platforms. From a security and reliability perspective, ACTN may encounter many risks such as malicious attack and rogue elements attempting to connect to various ACTN components (with PCE being one of them). Furthermore, some ACTN components represent a single point of failure and threat vector, and must also manage policy conflicts and eavesdropping of communication between different ACTN components. [RFC8453] further states that all protocols used to realize the ACTN framework should have rich security features, and customer, application, and network data should be stored in encrypted data stores. When PCEP is used as an ACTN interface, the security of PCEP provided by Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525] (unless explicitly set aside in [RFC8253]), is used.

As per [RFC8453], regarding the MPI, a PKI-based mechanism is suggested, such as building a TLS or HTTPS connection between the MDSC and PNCs, to ensure trust between the physical network layer control components and the MDSC. Which MDSC the PNC exports topology information to, and the level of detail (full or abstracted), should also be authenticated, and specific access restrictions and topology views should be configurable and/or policy based. When PCEP is used in MPI, the security functions, as per [RFC8253], are used to fulfill these requirements.

As per [RFC8453], regarding the CMI, suitable authentication and authorization of each CNC connecting to the MDSC will be required. If PCEP is used in CMI, the security functions, as per [RFC8253], can be used to support peer authentication, message encryption, and integrity checks.

8. References

8.1. Normative References

- [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>>.
- [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>>.
- [RFC6805] King, D., Ed. and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, DOI 10.17487/RFC6805, November 2012, <<https://www.rfc-editor.org/info/rfc6805>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.

8.2. Informative References

- [ASSOC-POLICY] Litkowski, S., Sivabalan, S., Tantsura, J., Hardwick, J., and M. Negi, "Path Computation Element communication Protocol extension for associating Policies and LSPs", Work in Progress, draft-ietf-pce-association-policy-05, February 2019.
- [EXP] Casellas, R., Vilalta, R., Martinez, R., Munoz, R., Zheng, H., and Y. Lee, "Experimental Validation of the ACTN architecture for flexi-grid optical networks using Active Stateful Hierarchical PCEs", 19th International Conference on Transparent Optical Networks (ICTON), DOI 10.5281/zenodo.832904, July 2017, <<https://zenodo.org/record/832904>>.
- [PCE-HPCE] Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., and D. King, "Hierarchical Stateful Path Computation Element (PCE).", Work in Progress, draft-ietf-pce-stateful-hpce-10, June 2019.

- [PCE-INTER-AREA] King, D. and H. Zheng, "Applicability of the Path Computation Element to Interarea and Inter-AS MPLS and GMPLS Traffic Engineering", Work in Progress, draft-ietf-pce-inter-area-as-applicability-07, December 2018.
- [PCE-INTERDOMAIN] Dugeon, O., Meuric, J., Lee, Y., and D. Ceccarelli, "PCEP Extension for Stateful Interdomain Tunnels", Work in Progress, draft-dugeon-pce-stateful-interdomain-02, March 2019.
- [PCE-STATE-SYNC] Litkowski, S., Sivabalan, S., Li, C., and H. Zheng, "Inter Stateful Path Computation Element (PCE) Communication Procedures.", Work in Progress, draft-litkowski-pce-state-sync-05, March 2019.
- [PCEP-LS] Dhody, D., Lee, Y., and D. Ceccarelli, "PCEP Extension for Distribution of Link-State and TE Information.", Work in Progress, draft-dhodylee-pce-pcep-ls-13, February 2019.
- [PCEP-OPTICAL] Lee, Y., Zheng, H., Ceccarelli, D., Wang, W., Park, P., and B. Yoon, "PCEP Extension for Distribution of Link-State and TE information for Optical Networks", Work in Progress, draft-lee-pce-pcep-ls-optical-07, March 2019.
- [PCEP-VN] Lee, Y., Zhang, X., and D. Ceccarelli, "Path Computation Element communication Protocol (PCEP) extensions for Establishing Relationships between sets of LSPs and Virtual Networks", Work in Progress, draft-leedhody-pce-vn-association-08, June 2019.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [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>>.

- [RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, DOI 10.17487/RFC5152, February 2008, <<https://www.rfc-editor.org/info/rfc5152>>.
- [RFC5212] Shiimoto, 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>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <<https://www.rfc-editor.org/info/rfc5307>>.
- [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, DOI 10.17487/RFC5441, April 2009, <<https://www.rfc-editor.org/info/rfc5441>>.
- [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>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC7149] Boucadair, M. and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7149, DOI 10.17487/RFC7149, March 2014, <<https://www.rfc-editor.org/info/rfc7149>>.
- [RFC7399] Farrel, A. and D. King, "Unanswered Questions in the Path Computation Element Architecture", RFC 7399, DOI 10.17487/RFC7399, October 2014, <<https://www.rfc-editor.org/info/rfc7399>>.

- [RFC7491] King, D. and A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations", RFC 7491, DOI 10.17487/RFC7491, March 2015, <<https://www.rfc-editor.org/info/rfc7491>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", RFC 8051, DOI 10.17487/RFC8051, January 2017, <<https://www.rfc-editor.org/info/rfc8051>>.
- [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>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/info/rfc8253>>.
- [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, <<https://www.rfc-editor.org/info/rfc8281>>.

- [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <<https://www.rfc-editor.org/info/rfc8283>>.
- [RFC8454] Lee, Y., Belotti, S., Dhody, D., Ceccarelli, D., and B. Yoon, "Information Model for Abstraction and Control of TE Networks (ACTN)", RFC 8454, DOI 10.17487/RFC8454, September 2018, <<https://www.rfc-editor.org/info/rfc8454>>.

Appendix A. Additional Information

In the paper [EXP], the application of the ACTN architecture is presented to demonstrate the control of a multidomain flexi-grid optical network by proposing, adopting, and extending:

- o the Hierarchical active stateful PCE architectures and protocols
- o the PCEP protocol to support efficient and incremental link-state topological reporting, known as PCEP-LS
- o the per-link partitioning of the optical spectrum based on variable-sized allocated frequency slots enabling network sharing and virtualization
- o the use of a model-based interface to dynamically request the instantiation of virtual networks for specific clients/tenants.

The design and implementation of the test bed are reported in order to validate the approach.

Acknowledgments

The authors would like to thank Jonathan Hardwick for the inspiration behind this document. Further thanks to Avantika for her comments with suggested text.

Thanks to Adrian Farrel and Daniel King for their substantial reviews.

Thanks to Yingzhen Qu for RTGDIR review.

Thanks to Rifaat Shekh-Yusef for SECDIR review.

Thanks to Meral Shirazipour for GENART review.

Thanks to Roman Danyliw and Barry Leiba for IESG review comments.

Thanks to Deborah Brungard for being the responsible AD.

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: dhruv.ietf@gmail.com

Young Lee
Futurewei Technologies
5340 Legacy Drive, Suite 173
Plano, TX 75024
United States of America

Email: younglee.tx@gmail.com

Daniele Ceccarelli
Ericsson
Torshamnsgatan, 48
Stockholm
Sweden

Email: daniele.ceccarelli@ericsson.com