

PCE Working Group
Internet Draft
Intended Status: Standard Track

Y. Lee
Huawei
G. Bernstein
Grotto Networking
H. Zheng
D. Dhody
Huawei

Expires: January 2015

July 2, 2014

PCEP Extensions in Support of Transporting Traffic Engineering Data

draft-lee-pce-transporting-te-data-00.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and 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>

This Internet-Draft will expire on January 2, 2009.

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.

Abstract

In order to compute and provide optimal paths, Path Computation Elements (PCEs) require an accurate and timely Traffic Engineering Database (TED). Traditionally this TED has been obtained from a link state routing protocol supporting traffic engineering extensions. This document discusses possible alternatives to TED creation. This document gives architectural alternatives for these enhancements and their potential impacts on network nodes, routing protocols, and PCE.

Table of Contents

1. Introduction.....	3
1.1. TED Creation and Maintenance via IGP-TEs.....	4
2. Alternative TED Creation & Maintenance for a PCE.....	6
2.1. Architecture Options.....	7
2.1.1. Nodes Send TE Info to all PCEs.....	12
2.1.2. Nodes Send TE Info via an Intermediate System.....	12
2.1.3. Nodes Send TE Info to At Least One PCE.....	12
2.2. Nodes Finding PCEs.....	13
2.3. Node TE Information Update Procedures.....	14
2.4. PCE TED Maintenance Procedures.....	14
3. Standardization and Protocol Considerations.....	14
3.1. Architecture Specific Standardization Aspects.....	15
4. Security Considerations.....	16
5. IANA Considerations.....	16
6. Conclusions.....	17
7. Acknowledgments.....	17
8. References.....	17
8.1. Normative References.....	17
8.2. Informative References.....	18
Author's Addresses.....	19
Intellectual Property Statement.....	19
Disclaimer of Validity.....	20

1. Introduction

In Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS), a Traffic Engineering Database (TED) is used in computing paths for connection oriented packet services and for circuits. The TED contains all relevant information that a Path Computation Element (PCE) needs to perform its computations. It is important that the TED be complete and accurate each time, the PCE performs a path computation.

In MPLS and GMPLS, interior gateway routing protocols (IGPs) have been used to create and maintain a copy of the TED at each node running the IGP. One of the benefits of the PCE architecture [RFC4655] is the use of computationally more sophisticated path computation algorithms and the realization that these may need enhanced processing power not necessarily available at each node participating in an IGP.

Section 4.3 of [RFC4655] describes the potential load of the TED on a network node and proposes an architecture where the TED is maintained by the PCE rather than the network nodes. However, it does not describe how a PCE would obtain the information needed to populate its TED. 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 [BGP-LS].

In this document we propose approaches for creating and maintaining the TED directly on a PCE as an alternative to IGPs and BGP flooding and investigate the impact from the PCE, routing protocol, and node perspectives.

New application areas for GMPLS and PCE in optical transport networks include Wavelength Switched Optical Networking (WSON) and Optical Transport Networks (OTN). WSON scenarios can be divided into routing wavelength assignment (RWA) problems where PCE requires detailed information about switching node asymmetries and wavelength constraints as well as detailed up to date information on wavelength usage per link [WSON-Frame]. As more data is anticipated to be made available to PCE with addition of OTN [RFC7062] and Flexi-grid [Flexi-grid] and possible with some optical impairment data [WSON-IMP-Info] even with the minimum set specified in [G.680], the total amount of data requires significantly more information to be held in the TED than is required for other traffic engineered networks.

In some circumstances such additional information could "bog down" the routing protocols on the nodes from a data processing, a storage, or communications perspective. In environments where PCEs

are external to the nodes running the routing protocol, and where the information in the TED is not used by the switching nodes it makes sense to investigate alternative methods to create and maintain the TED at its place of use, i.e., the PCE.

Recent development of a stateful PCE Model [PCE-Initiated] changes the PCE operation from path computation alone to include the support of PCE-initiated LSPs. With a stateful PCE model, it is also noted that LSP-DB is maintained by the PCE. For LSP state synchronization of stateful PCEs in GMPLS networks, the LSP attributes, such as its bandwidth, associated route as well as protection information etc, should be updated by PCCs to PCE LSP database (LSP-DB) [S-PCE-GMPLS]. To support all these recent changes in a stateful PCE model, a direct PCE interface to each PCC has to be supported. Relevant TED information can also be transported from each node to PCE using this PCC-PCE interface. Any resource changes in the node and links can also be quickly updated to PCE using this interface. Convergence time of IGP in GMPLS networks may not be quick enough to support on-line dynamic connectivity required for some applications.

This draft does not advocate that the alternative methods specified in this draft should completely replace the IGP-TE or BGP-LS as the method of creating the TED. The split between the data to be distributed via an IGP and the information conveyed via one of the alternatives in this document depends on the nature of the network situation. One could potentially choose to have some traffic engineering information distributed via an IGP while other more specialized traffic information is only conveyed to the PCEs via an alternative interface discussed here. In addition, the methods specified in this draft is only relevant to a set of architecture options where routing decisions are wholly or partially made in the PCE.

However, the networks that do not support IGP-TE/BGP-LS, the method proposed by this draft may be very relevant.

1.1. TED Creation and Maintenance via IGP-TEs

Routing protocols, in particular, IGP-TEs such as Open Shortest Path First (OSPF) and Intermediate system to intermediate system (IS-IS), take on a number of roles with respect to the control and data planes for IP, MPLS, and GMPLS. In all three technology families the underlying control plane communications technology is IP and hence all utilize the IGPs ability to control and run the IP data plane.

For the IP layer, the IGP directly establishes data plane connectivity. In the MPLS and GMPLS cases separate signaling protocols are used to directly control the data plane connectivity and in these cases the prime purpose of the routing protocol is to furnish network topology and resource status information used by path computation algorithms on the nodes or PCEs. Hence in the IP case the IGP is directly service impacting, while in the MPLS/GMPLS case it is only indirectly service impacting.

The IP layer information and the MPLS/GMPLS data plane layer information may be kept by the IGPs in two different information stores. These are referred to as databases but are not necessarily relational databases. In OSPF the information directly related to IP connectivity (and hence the control communications plane for all three technologies) and non-IP advertisements are kept in the link state database (LSDB), while information related to traffic engineering used by MPLS and GMPLS is kept in a (conceptually) separate TED which can be considered a subset of the LSDB. This TED information is distributed in a different data structure (Opaque LSA [RFC5250]). When we talk about adding additional technology-specific GMPLS information used for path computation we are only talking about adding to the TED and not the IP portion of the LSDB.

There are three main functions performed by an IGP: (a) hello protocol, (b) database synchronization (with neighbors), (c) database updates.

Data Plane Technologies	Hello Protocol	Database Sync & Updates
IP	Establish Control & Data Plane Adjacencies	LSDB
MPLS	Establish Control & Data Plane Adjacencies	LSDB & TED
GMPLS	Establish Control Plane Adjacencies (only)	LSDB & TED

Table 1 Main Functions of an IGP for various technologies

The procedures for maintaining LSDBs and TEDs in IGP-TEs have been very successful and well proven over time. These consist of:

1. Ageing the individual pieces of information in the TED (including discarding them when the information gets too old) to remove stale information from the TED.
2. Originator of the information being required to periodically resend TED information to prevent it from being discarded.
3. Originator of the information sending updates of information as needed, but subject to limits on how many/often these can be sent to keep the TED up-to-date, but to avoid swamping the network.
4. Reliable method for getting this information to other peers (flooding) to ensure that the information is delivered to all participants.
5. An efficient database synchronization mechanism for sharing info with a newly established peer.

From a PCE perspective, however, participating in an IGP, even as a passive receiver of IGP information, can place a significant load on the PCE. The IGP can be quite "chatty" when there are frequent updates to the use of the network, meaning that the PCE must dedicate significant processing to parsing protocol messages and updating the TED. Furthermore, to be truly useful, a PCE implementation would need to support OSPF and IS-IS.

2. Alternative TED Creation & Maintenance for a PCE

Given that nodes, by their position and role in the network, have accurate traffic engineering information concerning their local link ends and switching properties, it seems natural that, if other nodes in the network cannot make use of this information or do not want it, the information should only be conveyed to interested PCEs. In such case the flooding of TE information to all nodes may not be very efficient in terms of memory, CPU, bandwidth, etc.

The benefits of such an approach include:

- o Node: reduced storage demands (doesn't keep the entire TED)
- o Node: reduced processing demands for TED updates and synchronization

- o Control Plane: reduced overall communication demands since the TED is not being updated and maintained on all nodes in the network.
- o PCE: More timely TED updates are possible.
- o Information distribution constraints, such as seen in [Imp-Frame] can be met.

To quantify the previous advantages requires a bit more detail on how such an approach could actually be accomplished. The key pieces needed to implement such an approach include:

- o Multiple PCEs must be supported for robustness and load sharing.
- o Nodes must be able to find a PCE to which to send their traffic engineering information.
- o Nodes must have procedures and a mechanism (protocols) with which to communicate their TE information to a PCE. PCEs must have procedures and a mechanism (protocols) with which to receive this TE information from nodes.
- o Efficient mechanisms must exist in the multi-PCE case to ensure all PCEs have the same TED.

The advantages of using an alternative to IGP-TE comes at the cost of:

- o Additional protocols to be configured and secured. Recall that we still must have an IP IGP for control plane communications.
- o Any new protocols/implementations for alternative TED creation still must support many IGP-TE like features such as removal of stale information, reliable delivery of updates to all participants, recovery after reboots/crashes/upgrades, etc. It should also work along with IGP-TE/BGP-LS TED mechanism with some information in the TED received from existing mechanisms.
- o Node mechanisms to discover PCEs that are capable and willing to accept direct TED updates.

2.1. Architecture Options

There are three general architectural alternatives based on how nodes get their local TED information to the PCEs: (1) Nodes send local information to all PCEs; (2) Nodes send local information to

an intermediate server that will send to all PCEs; (3) Nodes send local information to at least one PCE and have the PCEs share this information with each other. An important functionality that needs to be addressed in each of these approaches is how a new PCE gets initialized in a reasonably timely fashion.

Figures 1-3 show examples of three options for nodes to share local TED information with multiple PCEs. As in the IGP case we assume that switching nodes know their local properties and state including the state of all their local links. In these figures the data plane links are shown with the character "o"; TE information flow from nodes to PCE by the characters "|", "-", "/", or "\"; and PCE to PCE TE information, if any, by the character "i".

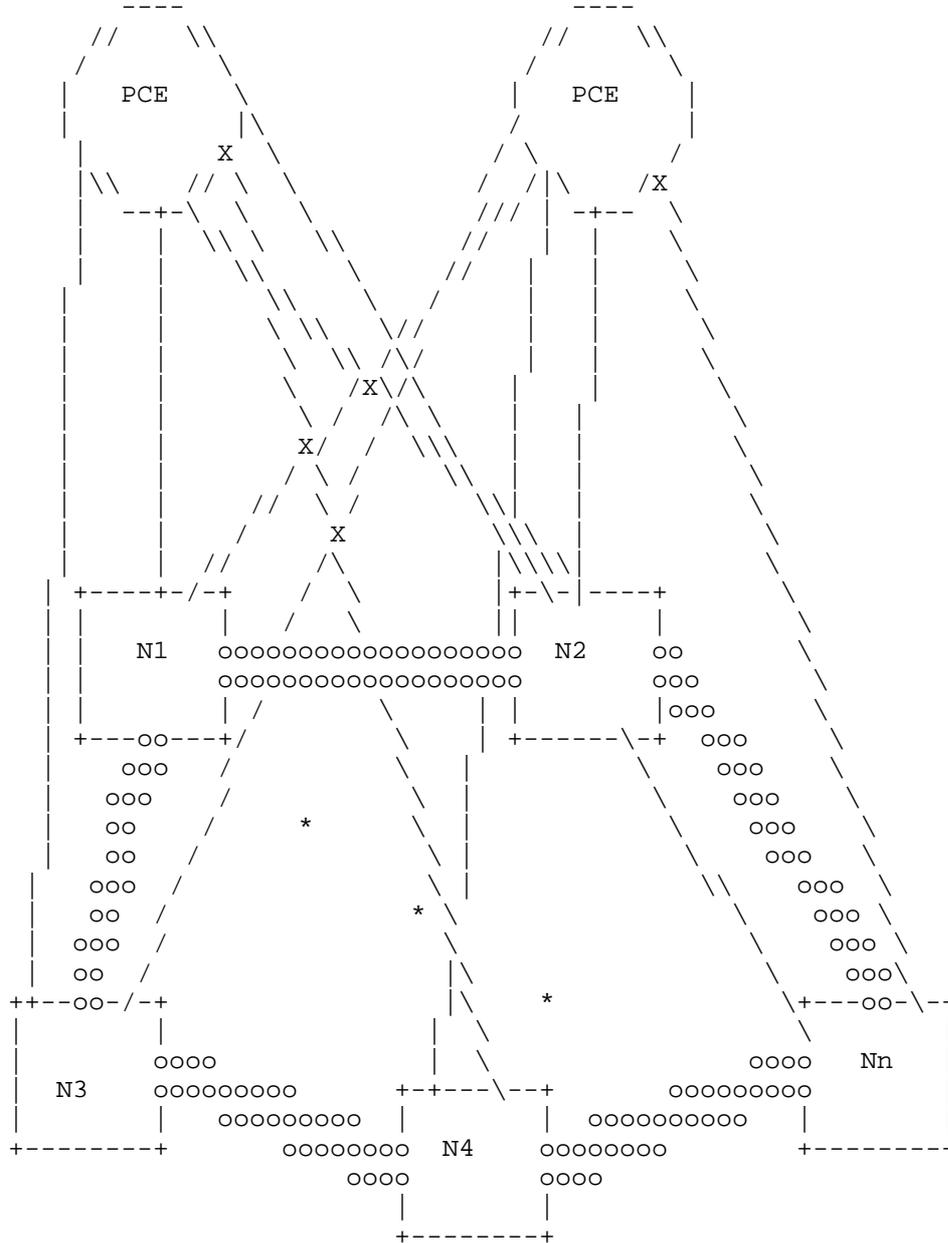


Figure 1 . Nodes send local TE information directly to all PCEs

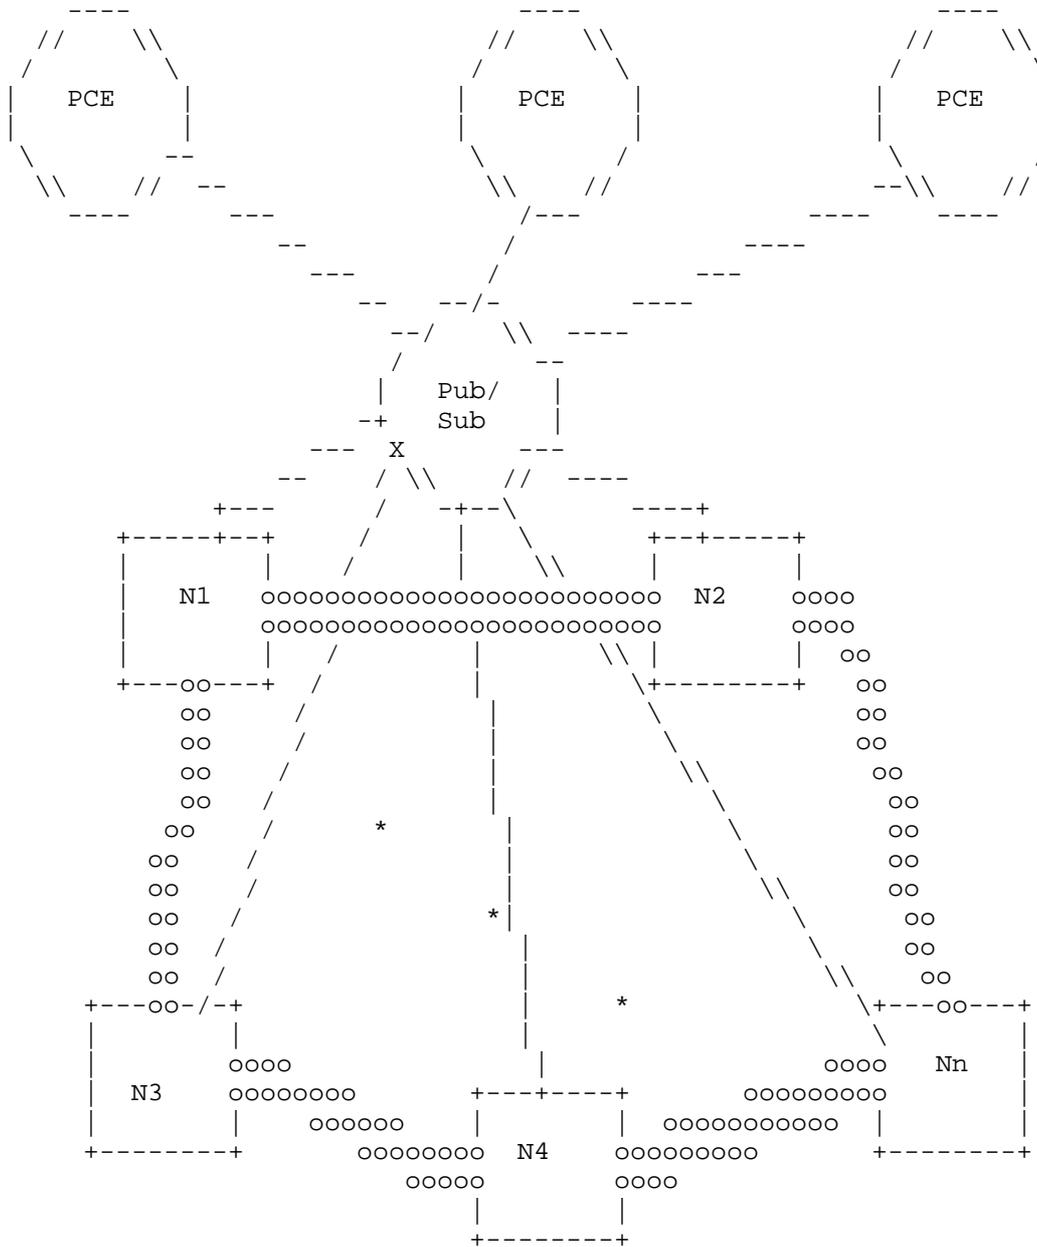


Figure 2 . Nodes send local TE information to PCEs via an intermediary (publish/subscribe)server

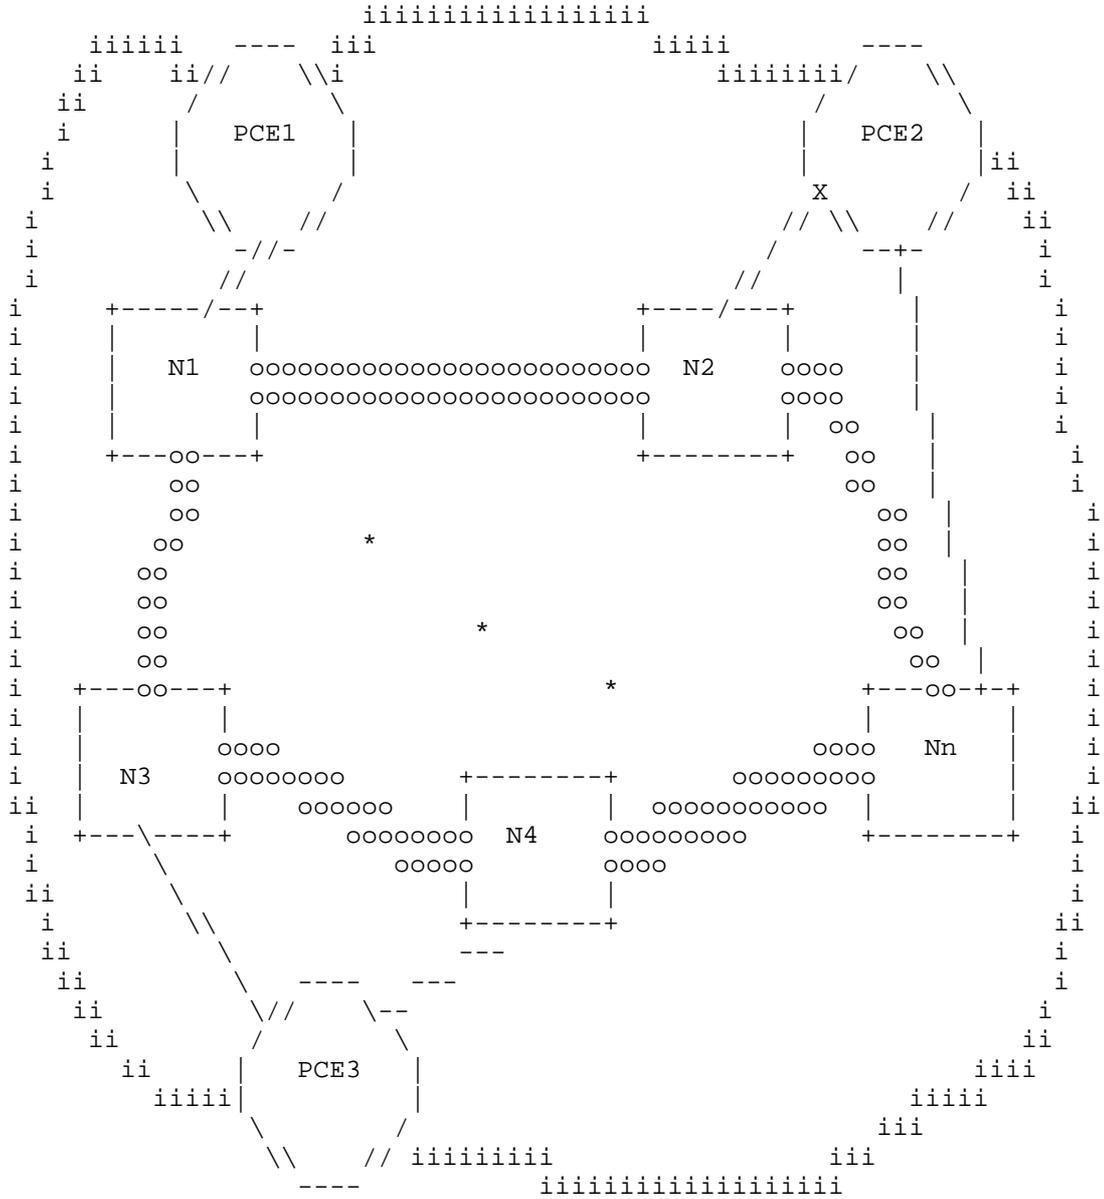


Figure 3 . Nodes send local TE information to at least one PCE and have the PCEs share TED information

2.1.1. Nodes Send TE Info to all PCEs

Architectural alternative 1 shown in Figure 1, illustrates nodes sending their local TE information to all PCEs within their domain. As the number of PCEs grows we have scaling concerns. However, if we are only talking about 2-3 PCEs, then we do not have this scaling concern. In particular each node needs to keep track of which PCE it has sent information to and update that information.

If a new PCE is added to the domain the node must send all its local TED information to that PCE rather than just sending status updates.

2.1.2. Nodes Send TE Info via an Intermediate System

Architecture alternative 2 is shown in Figure 2. This architecture reduces the burden on switching nodes by having the nodes send TE information to an intermediate system. This general approach is typically described in the software literature as a publish/subscribe paradigm. Here the nodes send their local TED information to an intermediate entity whose job is to insure that all PCEs receive this information. The nodes in this case being the publishers of the information and the PCEs the subscribers of the information. Publish/subscribe functionality can be found in general messaging oriented middleware such as the Java Messaging Service [JMS] and many others. A routing specific example of this approach is seen in BGP route reflectors [RFC4456].

Note that the publish/subscribe entity can be collocated with a PCE. This would then look like a master/slave type system architecture.

If a new PCE is added then the intermediate server will need to work with this new PCE to initialize its TED. Hence the publish/subscribe entity will need to also keep a copy of the entire TED and for reliability purposes a redundant server would be required. The publish/subscribe entity itself can be a PCE.

Architecture alternative 2 could be useful when there are a number of PCEs in the network and as such there is the scaling issue with each of the NEs talking to all the PCEs. The advantage of this alternative would diminish when we are dealing only with only a few PCEs.

2.1.3. Nodes Send TE Info to At Least One PCE

In this architectural alternative, shown in Figure 3, each node would be associated with at least one PCE. This implies that each

PCE will only have partial TED information directly from the nodes. It would be the responsibility of a node to get its local TED information to its associated PCE, then the PCEs within a domain would then need to share the partial TED information they learned from their associated nodes with each other so that they can create and maintain the complete TED. As we have seen in section 1.1. this is very similar to part of the functionality provided by a link state protocol, but in this case the protocol would be used between PCEs so that they can share the information they have obtained from their associated switching nodes (rather than from attached links as in a regular link state protocol).

To allow for this sharing of information PCEs would need to peer with each other. PCE discovery extensions [RFC4674] could be used to allow PCEs to find other PCEs. If a new PCE is added to the domain it would need to peer with at least one other PCE and then link state protocol procedures for TED synchronization could then be used to initialize the new PCEs TED.

A number of approaches can be used to ensure control plane resilience in this architecture. (1) Each node can be configured with a primary and a secondary PCE to send its information to; In case of failure of communications with the primary PCE the node would send its information to a secondary PCE (warm standby). (2) Each node could be configured to send its information to two different PCEs (hot standby).

2.2. Nodes Finding PCEs

In cases 1 and 3 nodes need to send TE information directly to PCEs. Path Computation Clients (PCCs) and network nodes participating in an IGP (with or without TE extensions) have a mechanism to discover a PCE and its capabilities. [RFC4674] outlines the general requirements for this mechanism and extensions have been defined to provide information so that PCCs can obtain key details about available PCEs in OSPF [RFC5088] and in IS-IS [RFC5089].

After finding candidate PCEs, a node would need to see which if any of the PCEs actually want to receive TE information directly from this node.

In architectural alternative 2 (publish/subscribe) the location of intermediate system would either need to be configured or PCE discovery could be extended so that a when a node asks a PCE if it wants to hear TE info the PCE points it to the intermediate publish/subscribe system.

2.3. Node TE Information Update Procedures

First a node must establish an association between itself and a PCE or intermediate system that will be maintaining a TED. It is the responsibility of the node to share TE information concerning its local environment, e.g., links and node properties. General and technology specific information models would specify the content of this information while the specific protocols would determine the format. Note that a node would not be sending to the PCE information it might be passed from neighbor nodes. Note that data plane neighbor information would be passed to the PCE embedded in TE link information.

There will be cases where the node would have to send to the PCE only a subset of TE link information depending on the path computation option. For instance, if the node is responsible for routing while the PCE is responsible for wavelength assignment for the route, the node would only need to send the PCE the WSON link usage information. This path computation option is referred to as separate routing (R) and wavelength assignment (WA) option in [PCE-WSON].

2.4. PCE TED Maintenance Procedures

The PCE is responsible for creating and maintaining the TED that it will use. Key functions include:

1. Establishing and authenticating communications between the PCE and sources of TED information.
2. Timely updates of the TED with information received from nodes, peers or other entities.
3. Verifying the validity of information in the TED, i.e., ensure that the network information obtained from nodes or elsewhere is relatively timely, or not stale.

3. Standardization and Protocol Considerations

In the previous section we examined a number of architectural alternatives for TED creation and maintenance on a PCE. Here we examine aspects of these alternatives that could be suitable for standardization. First there are a number of items and functions that can be independent of the particular architectural alternatives used, these include:

- o An information model for the TED

- o Basic PCE TED creation and maintenance procedures
- o Information packaging for use in TED creation, maintenance and exchange
- o NE to PCE (or Pub/Sub) communication of TED information --- interface and protocol (e.g. PCEP)
- o NEs discovering PCE (or Pub/Sub) for TED creation and maintenance purposes

By the "information model" for the TED we mean the raw information that a path computation algorithm would work with somewhat independent of how it might be packaged for TED maintenance and creation. Initial efforts along these lines have started at CCAMP for wavelength switched optical networks for non-impairment RWA [WSON-Info] and impairment aware RWA [WSON-IMP-Info].

Given a TED information model if we can agree on basic PCE TED creation and maintenance procedures we can then come up with a standardized way to package the information for use in such procedures. The analogy here is with an IGP's database maintenance procedures such as aging and the packaging of link state information into LSA (link state advertisements). LSAs form the basic chunks of an IGP's database. OSPF LSAs include an age field to assist in the ageing procedure and also has an advertising router field that aids in redistribution decisions, i.e., flooding. However the detailed TE information is encoded in LSAs via type length value (TLV) structures and it is this information that is used in path computation.

From there we could standardize the interface between a NE and a PCE for communication of TE information. This interface includes NE and PCE behaviors as well as a communications protocol.

Finally for the common behaviors we need a way for the NEs to find the PCEs or an intermediate publish/subscribe system to which they will send their TE information. As was previously pointed out this could be based on small enhancements to existing PCE discovery mechanisms.

3.1. Architecture Specific Standardization Aspects

Case 1: NEs send to all PCEs

This case has commonalities with both cases 2 and 3 and does not appear to have unique standardization aspects. As pointed out in section 2.1. we do need to consider when a new PCE comes online.

Case 2: Publish/Subscribe Server

In this case we would need to additionally standardize

1. how a new PCE coming online synchronizes with the publish/subscribe server
1. how PCEs and publish subscribe server communicate
2. Redundancy for publish subscribe server

Case 3: PCE to PCE sharing TE information learned from NEs

Here we would need the following additional mechanisms standardized:

1. The PCE to PCE interface and protocol
2. The method for PCEs to discover PCEs for the purpose of TE information sharing
3. PCE to PCE association for information sharing, in particular sharing update information.

4. Security Considerations

This draft discusses an alternative technique for PCEs to build and maintain a traffic engineering database. In this approach network nodes would directly send traffic engineering information to a PCE. It may be desirable to protect such information from disclosure to unauthorized parties in addition it may be desirable to protect such communications from interference (modification) since they can be critical to the operation of the network. In particular, this information is the same or similar to that which would be disseminated via a link state routing protocol with traffic engineering extensions.

5. IANA Considerations

This version of this document does not introduce any items for IANA to consider.

6. Conclusions

This document introduced several alternative architectures for PCEs to create and maintain a traffic engineering database (TED) via information directly or indirectly received from network elements and identified common aspects of these approaches. The TED is a critical piece of the overall PCE architecture since without it path computations cannot proceed. Though not explicitly out of scope the PCE working group does not have a work item or study item devoted to TED creation and maintenance. Such a work item can lead to enhanced interoperability and simplicity of PCE implementations. This document identified several common areas within these alternatives that could be standardized. In addition, the alternative approaches to TED creation and maintenance discussed here offloads both the network nodes and routing protocols from either some or all TED creation and maintenance duties at the same time it does not add significant new processing to a PCE that has already been participating in IGP based TED creation and maintenance.

7. Acknowledgments

TDB.

8. References

8.1. Normative References

- [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.
- [RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.
- [RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

8.2. Informative References

- [JMS] Java Message Service, Version 1.1, April 2002, Sun Microsystems.
- [PCE-Initiated] E. Crabbe, et. al., "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", draft-ietf-pce-pce-initiated-lsp, work in progress.
- [S-PCE-GMPLS] X. Zhang, et. al, "Path Computation Element (PCE) Protocol Extensions for Stateful PCE Usage in GMPLS-controlled Networks", draft-ietf-pce-pcep-stateful-pce-gmpls, work in progress.
- [BGP-LS] H. Gredler, et. al., "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution, work in progress.
- [Flexi-grid] O. Gonzalez de Dios, Ed., and R. Casellas, Ed., "Framework and Requirements for GMPLS based control of Flexi-grid DWDM networks", draft-ietf-ccamp-flexi-grid-fwk, work-in-progress.
- [PCE-WSON] Y. Lee, G. Bernstein, "PCEP Requirements for the support of Wavelength Switched Optical Networks (WSON)", work in progress, draft-lee-pce-wson-routing-wavelength-05.txt, February 2009.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, April 2006.
- [Imp-Frame] G. Bernstein, Y. Lee, D. Li, A Framework for the Control and Measurement of Wavelength Switched Optical Networks (WSON) with Impairments, Work in Progress, October 2008.
- [WSON-Frame] Y. Lee, G. Bernstein, W. Imajuku, "Framework for GMPLS and PCE Control of Wavelength Switched Optical Networks", work in progress: draft-ietf-ccamp-wavelength-switched-framework-01.txt, February 2009.
- [WSON-IMP-Info] Y. Lee, G. Bernstein, "Information Model for Impaired Optical Path Validation", work in progress: draft-bernstein-wson-impairment-info-02.txt, March 2009.

Author's Addresses

Young Lee
Huawei Technologies
5340 Legacy Drive, Building 3
Plano, TX 75023, USA

Phone: (469) 277-5838
Email: leeyoung@huawei.com

Greg Bernstein
Grotto Networking

Email: gregb@grotto-networking.com

Haomian Zheng
Huawei Technologies Co., Ltd.
F3-1-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28979835
Email: zhenghaomian@huawei.com

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

Email: dhruv.ietf@gmail.com

Contributor's Addresses

Intellectual Property Statement

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

