Network Working Group
Internet-Draft
Intended status: Informational

Fatai Zhang
Xian Zhang
Young Lee
Huawei
Ramon Casellas
CTTC
Oscar Gonzalez de Dios
Telefonica

Expires: September 05, 2012
March 05, 2012

Applicability of Stateful Path Computation Element (PCE)

draft-zhang-pce-stateful-pce-app-00.txt

Status of this Memo

Abstract

The Path Computation Element (PCE) provides a solution for Traffic Engineering (TE) based path calculation in large, multi-domain, multi-region, or multi-layer networks. Depending on whether a PCE keeps information about LSPs and reserved resource usage in the network or not, it can be categorized as either stateful or stateless.

This memo describes general considerations for stateful PCE(s) and examines its applicability through a number of typical scenarios. It shows how stateful PCE(s) can be applied to facilitate these applications.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Table of Contents

1. Introduction

[RFC 4655] defines the architecture for a Path Computation Element (PCE)-based model for the computation of Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering Label Switched Paths (TE LSPs). To perform such a constrained computation, a PCE stores the network topology (i.e., TE links and nodes) and resource information (i.e., TE attributes) in its TE Database (TED). To request path computation services to a PCE, [RFC 5440] defines the PCE Communication Protocol (PCEP) for communications between a Path Computation Client (PCC) and a PCE, or between two PCEs. A PCC can initiate a path computation request to a PCE through a Path Computation Request (PCReq) message, and then the PCE will return the computed path to the requesting PCC in response to a previously received PCReq message through a PCEP Path Computation Reply (PCRep) message.

Per [RFC 4655], a PCE can be either stateful or stateless. Compared to a stateless PCE, a stateful PCE stores not only the network states, but also the set of computed paths and reserved resources in use in the network. In other words, the ''state'' in a stateful PCE is determined not only by the TED but also by the set of active LSPs and their corresponding reserved resources. Such augmented state allows the PCE to compute constrained paths while considering individual LSPs and their interaction. Note that [RFC4655] further specifies that the TED contains link state and bandwidth availability as distributed by the IGPs or collected via other methods. Even if such information can provide increased granularity and more detail, it is not state information in the PCE context and so a model that uses it is still described as a stateless PCE.

As described in section 6.8 of [RFC 4655], there are many applications which can benefit from stateful PCE(s), e.g.:

o Minimum perturbation: stateful PCE(s) can minimize the number of existing TE LSPs that are affected and preempted by a higher-priority TE LSP request in a crowded network.

o Virtual Network Topology (VNT) maintenance: the information of existing LSPs in the higher layer is used as an input for setting up/tearing down the LSPs in the lower layer (i.e., VNT modification).

Besides these scenarios, there are some additional scenarios that should be investigated further. For instance, in impairment-aware Wavelength Switched Optical Networks (WSON) [WSON-Impairment], stateful PCEs could be used to perform Impairment-Aware Routing and Wavelength Assignment (IA-RWA) procedures. In this case, PCE(s) need to know the detailed information of the existing LSPs so that the new LSP(s) will not impact them. Such PCE(s) would maintain the existing LSPs states (e.g., route, wavelength and speed) to perform impairment aware RWA procedures simpler and with less protocol overhead.

[RFC 4655] also discusses potential scalability and synchronization issues in order to implement stateful PCE(s). The main problem pointed out by [RFC 4655] is that a PCE would be constrained if all the states of TE LSPs of a network are to be maintained by a PCE. Moreover, such state, when there are multiple PCEs, needs to be properly synchronized. These issues are especially relevant in packet networks, such as MPLS-TE networks, given a potentially large number of LSPs. Nonetheless, it is expected that in transport networks, such as OTN networks, the number of the LSPs will be much smaller, which makes stateful PCEs more applicable. Finally, with the increasing power and memory of the hardware platforms that a PCE may run, the number of LSPs that can be managed by a PCE is significantly large. Hence, there is lesser scaling issue for a PCE to store all the LSPs states, especially for a transport network.

This document presents general considerations for stateful PCE(s) and several examples of its application scenarios. It exhibits the utility of stateful PCE(s) in effective support of these applications to obtain better performance.

2. General Considerations

2.1. Architectural Considerations

Several PCE architectures are described in Section 5 of [RFC4655]. A stateful PCE needs to maintain a large amount of data and potentially incur in a very high amount of control plane overhead. Moreover, there might be high computational demands on stateful PCE entities to effectively support the applications listed in Section 3. Therefore, the composite PCE architecture is NOT RECOMMENDED to support stateful PCEs. It does not exclude the possibility that multiple PCEs with different capabilities are included in the network. For example, both stateless and stateful PCEs can co-exist to be in charge of path computation of different types.

2.2. LSP State Synchronization

As suggested by the definition, a stateful PCE maintains two databases for path computation. The first one is the Traffic Engineering Database (TED) which includes the topology and resource in the network. TED can be obtained through participating in routing distribution of TE information or other means as explained in Section 6.7 of [RFC4655].

The other database is the LSP state Database (LSP-DB), in which a PCE stores attributes of all existing LSPs in the network, such as payload signal, switching types and bandwidth/resource usage etc. A stateful PCE should gather the LSP information either from the network management system (NMS) or from the nodes in the network. For a NMS-based PCE, if the PCE is not collocated with the NMS, a standard communication protocol might be needed for LSP state synchronization; otherwise, proprietary APIs can be used. If a PCE rely on network nodes for state synchronization, the strategies may vary depending on the network scenarios in which the PCE is applied to (i.e., single domain, multiple domain or multi-layer networks.) as well as the adoption of PCE computation model.

2.2.1. Single Domain

In a single domain network, LSP state information is maintained locally by the nodes initiating LSP(s). Therefore, PCE(s) should gather the LSP state information either passively or actively from the nodes in the network they have visibility. With a centralized stateful PCE computation model, it is straightforward that all nodes in the domain could communicate with the PCE for its LSP-DB synchronization. As for distributed stateful PCE computation model (i.e., there are multiple stateful PCEs in the network), there are several alternatives for synchronization:

o Every node can update the PCE LSP-DBs by sending the LSP state information to each of the PCEs in the network separately.

o Another feasible strategy is to choose one of the PCEs (i.e., a designated PCE) for synchronization with all the nodes in the network and also update the LSP-DBs of all the other PCE(s).

o A mixed of these two methods listed above can also be considered in which more than one PCEs (e.g., two PCEs) are chosen to interact directly with nodes in the network for state synchronization while other PCEs are updated via these PCEs.

2.2.2. Multi-domain

In a multi-domain network with a centralized PCE model, the LSP state synchronization is similar to that of a single domain scenario. If there is a stateful PCE responsible for performing path computation within each domain, the LSPs (segments) traversing the domain/layer should be synchronized to the PCE.

As described in [RFC4726], there are four methods to set up a LSP traversing multiple domains: LSP nesting, contiguous LSP, LSP stitching and hybrid methods, respectively. Hence, the ingress nodes of a LSP traversing a domain may exist in another domain (e.g., a contiguous LSP spanning across multiple domains). In this case, the border node of a domain (i.e., an intermediate node of a LSP), could be responsible for synchronizing the LSP segment in the domain to the PCE.

```
        +---------------------+---------------------+
        |      +----+         |       +----+        |
        |      |PCE1|         |       |PCE2|        |
        |      +----+         |       +----+        |
        |      Domain 1       |       Domain 2      |
        |  +--+   +--+   +--+ | +--+   +--+   +--+  |
        |  |N1+---+N2+---+N3+---+N7+---+N8+---+N9|  |
        |  +-++   +--+   +-++ | +-++   +--+   +-++  |
        |    |             |  |   |             |   |
        |    |             |  |   |             |   |
        |  +-++   +--+   +-++ | +-+-+        +--++  |
        |  |N4+---+N5+---+N6+---+N10+--------+N11|  |
        |  +--+   +--+   +--+ | +---+        +---+  |
        +---------------------+---------------------+
```

Figure 1: Multi-domain Scenario

Figure 1 shows an example of multi-domain scenario. Suppose a contiguous LSP traverses N1-N2-N3-N7-N8-N9. Then in domain 1, the ingress node of the LSP (i.e., N1) SHOULD synchronize the state of the LSP segment N1-N2-N3 to PCE1. In domain 2, the border node (i.e., N7) SHOULD synchronize the state of the LSP segment N7-N8-N9 to PCE2.

This approach requires that N7 has a PCEP adjacency with its PCE (PCE2) even if no path computation expansions are required. N7 needs to check whether its RSVP-TE upstream node belongs to another domain and notify the PCE when the LSP is released. Note that synchronization may require detailed information of the LSP (e.g., a full record route, the actual reserved resources) which may only be available during Resv message processing.

Alternatively, inter-PCE communication strategy can be adopted for LSP-DB synchronization. For instance, in Figure 1, upon the notification of the setup of LSP N1-N2-N3-N7-N8-N9, PCE1 can establish a PCEP adjacency to inform PCE2 to update its LSP-DB. This method SHOULD be preferred only when PCE1 has sufficient and valid information of the across-domain LSP, such as explicit LSP information. Otherwise, the method in which the border node(s) are in charge of LSP state update is more appropriate. For example, Backward Recursive Path Computation (BRPC) [RFC5441] in conjunction with path-key-based mechanism [RFC5520] can be adopted for inter-domain path computation. If this is the case with the example in Figure 1, PCE1 only acquires a loose LSP path (e.g., N1-N2-N3-N7-KEY1, where KEY1 can be interpreted only by PCE2). Since it depends on the local policy that how long a Path-Key should be stored, KEY1 might not be valid anymore when it is used by PCE1 for PCE2 LSP-DB update notification. In this case, N7 will need to request PCE2 to unlock the Path-Key in order to complete the signaling process. Therefore, it is possible to use N7 instead for updating PCE2 LSP-DB.

Note that a timely synchronization of PCEs and these two databases is a prerequisite to maintaining a good performance of a stateful PCE.

2.2.3. Multi-layer

In multi-layer scenarios, one node/domain may have multiple switching capabilities. For instance, Optical Transport Network (OTN) nodes may have both of electrical (e.g., ODU1, ODU2, ODU3) and optical switch capabilities. ODU LSPs and wavelength LSPs may be established in an OTN network.

In such networks, a PCE may have the capability of performing single layer path computation or multi-layer path computation. If a stateful PCE has single layer path computation capability, the nodes should be aware of information pertaining to which layer should be synchronized to a specific PCE. Otherwise, the state of the LSPs in all layers should be synchronized to the single stateful PCE.

2.3. PCE Survivability/Reliability

Since a PCE supports a centralized path computation model, its survivability should be carefully considered to ensure its proper operation. If a multiple stateful PCE model is used and these PCEs have a consistent view of the network, they can act as a hot backup for each other. Otherwise, other backup strategies SHOULD be present if only one PCE is deployed in the network to avoid a single point of failure.

2.4. Delegation and Policy

Stateful PCE(s) are still subject to policies when performing path computation based on TED and LSP-DB as well as in what concerns LSP-DB organization and maintenance.

For LSP-DB maintenance, a basic function of stateful PCEs that SHOULD be supported is the ability to keep LSP state information in the network within which they have visibility. OPTIONALLY, a stateful PCE can also extend its ability to support modification of LSP state information. This can be realized by obtaining the temporal LSP state control through negotiation with LSRs (i.e., LSP delegation). Please note that LSP state delegation should comply with the policy imposed by LSP state owner (i.e., LSRs) as well as the policy imposed upon PCE(s).

3. Application Scenarios

In this section, several examples exploiting the capabilities of stateful PCE(s) are presented, although the application of stateful PCE(s) is not limited to them. In general, stateful PCE(s) can be deployed for applications where LSP state as well as traffic engineering information in the network are necessary inputs to achieve one or multiple of the following goals:

o Improving the performance such as reducing network blocking probability, achieving load balancing, improve network resources utilization or increasing the route computation success rate;

o Reducing the complexity of the relevant procedure(s) associated with the application(s);

o Lowering resource consumption;

As discussed in [PSU-WSON] and [LCA-Stateless], some of the objectives can be achieved through limited LSP awareness in stateless PCE by exploiting objects defined in existing protocols, such as the SVEC object defined in [RFC5440] and/or XRO object defined in [RFC5521]. These methods are considered as transitional solutions because of two reasons. Firstly, these methods only have local/partial/temporal LSP related information and thus have limited utility in terms of achieving the goals, particularly for objectives set at a network level. Secondly, it might incur a substantial amount of overhead since it requires frequent message exchanges among PCC and PCE entities.

3.1. Impairment-Aware Routing and Wavelength Assignment (IA-RWA)

In WSON networks [RFC6163], a wavelength-switched LSP traverses one or multiple fiber links. The bit rates of the client signals carried by the wavelength LSPs may be the same or different. Hence, a fiber link may transmit a number of wavelength LSPs with equal or mixed bit rate signals. For example, a fiber link may multiplex the wavelengths with only 10G signals, mixed 10G and 40G signals, or mixed 40G and 100G signals.

IA-RWA in WSONs refers to the RWA process (i.e., lightpath computation) that takes into account the optical layer/transmission imperfections by considering as additional (i.e., physical layer) constraints. To be more specific, linear and non-linear effects associated with the optical network elements should be incorporated into the route and wavelength assignment procedure. For example, the physical imperfection can result in the interference of two adjacent lightpaths. Thus, a guard band should be reserved between them to alleviate these effects. The width of the guard band between two adjacent wavelengths depends on their characteristics, such as modulation formats and bit rates. Two adjacent wavelengths with different characteristics (e.g., different bit rates) may need a wider guard band and with same characteristics may need a narrower guard band. For example, 50GHz spacing may be acceptable for two adjacent wavelengths with 40G signals. But for two adjacent wavelengths with different bit rates (e.g., 10G and 40G), a larger spacing such as 300GHz spacing may be needed. Hence, the characteristics (states) of the existing wavelength LSPs SHOULD be considered for a new RWA request in WSON.

In summary, when stateful PCE(s) are used to perform the IA-RWA procedure, it needs to know the characteristics of the existing wavelength LSPs. The impairment information relating to existing and to-be-established LSPs can be obtained by nodes in WSON networks via external configuration or other means such as monitoring or estimation based on a vendor-specific impair model. However, WSON related routing protocols, i.e., [GEN-OSPF] and [WSON-OSPF], only advertise limited information (i.e., availability) of the existing wavelengths, without defining the supported client bit rates. It will incur substantial amount of control plane overhead if routing protocols are extended to support dissemination of the new information relevant for the IA-RWA process. In this scenario, stateful PCE(s) would be a more appropriate mechanism to solve this problem. Stateful PCE(s) can exploit impairment information of LSPs stored in LSP-DB to provide accurate RWA calculation.

3.2. Defragmentation in Flexible Grid Networks

Traditionally, in Dense Wavelength Division Multiplexing (DWDM) networks, the frequency and channel spacing for a single wavelength allocated to an optical connection is fixed, in terms of a fixed channel spacing grid. With the development of mixed-rate transmission and the increase in the speed of optical signal, the issue of poor optical spectrum usage needs to be addressed. Flexible grid is proposed to solve this problem [G.FLEXIGRID]. In Flexible grid networks, LSPs with different slot widths (such as 12.5G, 25G etc.) can co-exist so as to accommodate the services with different bandwidth requests.

Yet another problem arises in this type of DWDM networks. Since in flexible grid networks LSPs are dynamically allocated and released over time, the optical spectrum resource becomes fragmented. The overall available spectrum resource on a link might be sufficient for a new LSP request. But if the available spectra are not continuous, the request would be rejected. In order to perform frequency defragmentation procedure, stateful PCE(s) COULD be used, since existing TE LSPs information (i.e., slot width and spectrum location information associated with TE LSPs) is required to accurately assess spectrum resources on the LSPs, and perform de-fragmentation while ensuring a minimal disruption of the network, e.g., based on active LSP priorities.

3.3. Recovery

3.3.1. Protection

For protection purposes, a PCC may send a request to a PCE for computing a set of paths for a given LSP. Alternatively, the PCC can send multiple requests to the PCE, asking for working and backup LSPs separately. In either way, the resources bound to backup paths can be shared by different LSPs to improve the overall network efficiency. If resource sharing is supported for LSP protection, the information relating to existing LSPs is required to avoid allocation of shared protection resources to two LSPs that might fail together and cause protection contention issues. If such information is required on each network node, extensions to existing signaling or routing protocols are needed in order to carry the necessary information for avoiding allocating shared protection resources for two non-disjoint working LSPs. However, stateful PCE(s) can easily accommodate this need using the information stored in its LSP-DB, without requiring extensions to existing routing protocols.

```
                 +----+
                 |PCE |
                 +----+

            +-------+         +-------+         +-------+
            |  N1  +----------+  N2  +----------+  N3   |
            +--+---+          +---+--+          +---+---+
               |                  |                 |
               |        +---------+                 |
               |        |                           |
               |     +--+---+          +-------+    |
               +-----+  N5  +----------+  N4  +-----+
                     +-------+         +-------+
```

Figure 2: Example Network

For example, in the network depicted in Figure 2, suppose there exists LSP1 (N1->N5) with backup route following N1->N2->N5. A request arrives asking for a working and backup path pair to be computed for a request from N2 to N5. If the PCE decides N2->N1->N5 to be the best working route, then the backup path should not use the same protection resource with LSP1 since the new LSP shares part of its resource with LSP1 (i.e., these two LSPs are in the same shared risk group). Alternatively, there is no such constraint if N2->N3->N4->N5 is chosen to be the right candidate for undertaking the request.

3.3.2. Restoration

In case of a link failure, such as fiber cut, multiple LSPs may fail at the same time. Thus, the source nodes of the affected LSPs will be informed of the failure by the nodes detecting the failure. These source nodes will send requests to a PCE for rerouting. In order to reuse the resource taken by an existing LSP, the source node can send a PCReq message including the XRO object with F bit set, together with RRO object, as specified in [RFC5521].

If a stateless PCE is exploited, it might respond to the rerouting requests separately if they arrive at different times. Thus, it might result in sub-optimal resource usage. Even worse, it might unnecessarily block some of the rerouting requests due to insufficient resources for later-arrived rerouting messages. If a stateful PCE is used to fulfill this task, it can re-compute the affected LSPs concurrently while reusing part of the existing LSPs resources when it is informed of the failed link identifier provided by the first request. This is made possible since the stateful PCE can check what other LSPs are affected by the failed link and their

route information by inspecting its LSP-DB. As a result, a better performance, such as better resource usage, minimal probability of blocking upcoming new rerouting requests sent as a result of the link failure, can be achieved.

In order to further reduce the amount of LSP rerouting messages flow in the network, the notification can be performed at the node(s) which detect the link failure. For example, suppose there are two LSPs in the network as shown in Figure 2: (i) LSP1: N1->N5->N4->N3; (ii) LSP2: N2->N5->N4. They traverse the failed link between N5-N4. When N4 detects the failure, it can send a notification message to a stateful PCE. Note that the stateful PCE stores the path information of the LSPs that are affected by the link failure, so it does not need to acquire this information from N4. Moreover, it can make use of the bandwidth resources occupied by the affected LSPs when performing path recalculation. After N4 receives the new paths from the PCE, it notifies the ingress nodes of the LSPs, i.e., N1 and N2, and specifies the new paths which should be used as the rerouting paths. To support this, it would require extensions to existing signaling protocol.

3.4. SRLG Diversity

A common requirement is to maintain SRLG disjointness between LSPs. This can be achieved at provisioning time, if the routes of all the LSPs are requested together, using a synchronized computation of the different LSPs with SRLG disjointness constraint. If the LSPs need to be provisioned at different times, (more general, the routes are requested at different times, e.g. in the case of a restoration), the PCC can specify, as constraints to the path computation a set of Shared Risk Link Groups (SRLGs) using the Explicit route Object [RFC 5521]. However, for the latter to be effective, it is needed that the entity that request the route to the PCE maintains updated SRLG information of all the LSPs to which it must maintain the disjointness.

Using a stateful PCE allows the maintenance of the updated SRLG information of the established LSPs in the PCE. Having such information in the PCE facilitates the PCC to specify, as constraint to the path computation, the SRLG disjointess of a set of already established LSPs.

3.5. Maintenance of Virtual Network Topology (VNT)

In Multi-Layer Networks (MLN), a Virtual Network Topology (VNT) [RFC5212] consists of a set of one or more TE LSPs in the lower layer to provide TE links to the upper layer. In [RFC5623], the PCE-

based architecture is proposed to support path computation in MLN networks in order to achieve inter-layer TE.

The establishment/teardown of a TE link in VNT needs to take into consideration the state of existing LSPs and/or new LSP request(s) in the higher layer. Without detailed LSP information, this process would be inefficient or even infeasible, requiring the cooperation of a NMS or a VNT manager (VNTM). Therefore, a stateful PCE seems more suitable to make the decision of when and how to modify the VNT either to accommodate new LSP requests or to re-optimize resource use across layers irrespective of PCE models. As described in Section 2.2, path computation for a VNT change can be performed by the PCE if a single PCE model is adopted. On the other hand, if a per-layer PCE model is more appropriate, coordination between PCEs is required.

3.6. Global Concurrent Optimization (GCO)

GCO is introduced in [RFC5557] to calculate multiple paths concurrently so as to improve network resource efficiency. By taking into consideration the network topology as well as existing TE LSPs information, GCO can (re)optimize the entire network simultaneously. Alternatively, GCO can be applied to (re)optimize one or a subset of existing TE LSPs or plan for forthcoming LSP(s) with specific objectives. GCO can also support off-line one-time optimization (i.e., planning) given a traffic matrix and network topology. Due to its complexity and potentially high computational demand, it is recommended to be performed in a centralized way (e.g., based on a management-based PCE).

In case of a stateless PCE, in order to optimize network resource usage dynamically through online planning, PCC (e.g., NMS) should send a request to PCE together with detailed path/bandwidth information of the LSPs that need to be concurrently optimized. This would require a PCC (e.g., NMS) to determine when and which LSPs should be optimized. Given all of the existing LSP state information kept at a stateful PCE, it allows automation of this process without the PCC (e.g. NMS) to supply the existing LSP state information. Moreover, since a stateful PCE can maintain the information regarding to all LSPs that are currently under signaling, it makes the optimization procedures be performed more intelligently and effectively.

3.7. Point-to-Multipoint (P2MP) Application

Route computation for P2MP application involves selection of branching points together with calculating multiple sub-LSPs with certain objective(s) such as minimizing the overall cost of the P2MP

tree. Moreover, egress nodes addition and removal in a P2MP tree necessitates (re)optimization. Besides these, there are also some constraints and policies that make the P2MP tree computation hard, requiring high computation power. Therefore, PCE is proposed to support P2MP application [RFC5671].

If a stateless PCE is used for P2MP calculation or optimization under constraints such as load balancing or path disjointedness, then a large amount of sub-LSP information might need to be exchanged between the PCE and the requesting entities. Moreover, if the requesting entity cannot provide complete information of sub-LSPs pertaining to the P2MP tree, then the performance of stateless PCE will be sub-optimal. On the contrary, a stateful PCE can support the P2MP tree computation/optimization with reduced overhead and improved efficiency.

3.8. Time-based Scheduling

Time-based scheduling allows network operators to reserve resources in advance upon request from the customers to transmit large bulk of data with specified starting time and duration, such as in support of scheduled data transmission between data centers.

Traditionally, this can be supported by NMS operation through path pre-establishment and activation on the agreed starting time. However, this does not provide efficient network usage since the established paths exclude the possibility of being used by other services even when they are not used for undertaking any service. It can also be accomplished through GMPLS protocol extensions by carrying the related request information (e.g., starting time and duration) across the network. Nevertheless, this method inevitably increases the complexity of signaling and routing process.

A stateful PCE can support this application with better efficiency since it can alleviate the burden of processing on network elements as well as enable the flexibility of resources usage by only excluding the time slot(s) reserved for time-based scheduling requests. In order to support this application, a stateful PCE should also maintain a database that stores all the reserved information with time reference. This can be achieved either by maintaining a separate database or incorporated into LSP-DB. The details of organizing time-based scheduling related information as well as its impact on LSP-DB is subject to network provider's policy and administrative consideration and thus outside of the scope of this document.

4. Manageability Considerations

TBD.

5. Security Considerations

TBD.

6. References

6.1. Normative References

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

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

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

[RFC6163] Lee, Y., Bernstein, G., "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, April, 2011.

[RFC5521] Oki, E., Farrel, A., "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", RFC5521, April 2009.

6.2. Informative References

[WSON-Impairment] Lee, Y., Bernstein, G., Li, D., Martinelli, G., "A Framework for the Control of Wavelength Switched Optical Network (WSON) with Impairments", draft-ietf-ccamp-wson-impairments, work in progress.

[RFC4726] Farrel, A., Vasseur, J.-P., Ayyangar, A., "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.

[RFC5520] Bradford, R., Vasseur, JP., Farrel, A., "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.

[RFC5441] Vasseur, J.-P., Zhang, R., Bitar, N., Le Roux, JL., "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, April 2009.

[PSU-WSON] Giorgetti, A, Cugini, G, et al, "Path state-based update of PCE traffic engineering database in wavelength switched optical networks", IEEE Com. Let., June 2010.

[LCA-Stateless] Gonzalez de Dios, O., et al, "Benefits of limited context awareness in stateless PCE", Optical Fiber Communication Conference, March 2011.

[WSON-OSPF] Lee, Y., Bernstein, G., "GMPLS OSPF Enhancement for Signal and Network Element Compatibility for Wavelength Switched Optical Networks", draft-ietf-ccamp-wson-signal-compatibility-ospf-07, October 2011.

[GEN-OSPF] Zhang, Fatai, Lee, Y., Han, Jianrui, Bernstein, G., Xu, Yunbin, "OSPF-TE Extensions for General Network Element Constraints", draft-ietf-ccamp-gmpls-general-constraints-ospf-te-02, September 2011.

[G.FLEXIGRID] Draft revised G.694.1 version 1.3, Unpublished ITU-T Study Group 15, Question 6.

[RFC5212] Shiomoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., Brungard, D., "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, July 2008.

[RFC5557] Lee, Y., Le Roux, JL., King, D., Oki E., "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", RFC 5557, July, 2009.

[RFC5671] Yasukawa, S., Farrel, A., "Applicability of the Path Computation Element (PCE) to Point-to-Multipoint (P2MP) MPLS and GMPLS Traffic Engineering (TE)", October, 2009.

[RFC5623] Oki, E., Takeda, T., Le Roux, JL., Farrel, A., "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC5623, September 2009.

7. Contributors' Address

Dhruv Dhody
Huawei Technology
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruvd@huawei.com

Xiaobing Zi
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28973229
Email: zixiaobing@huawei.com

Authors' Addresses

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972913
Email: zhang.xian@huawei.com

Young Lee
Huawei
1700 Alma Drive, Suite 100

Plano, TX 75075
US

Phone: +1 972 509 5599 x2240
Fax: +1 469 229 5397
EMail: ylee@huawei.com

Ramon Casellas
CTTC - Centre Tecnologic de Telecomunicacions de Catalunya
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
Spain

Phone:
Email: ramon.casellas@cttc.es

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain

Phone: +34 913374013
Email: ogondio@tid.es

Intellectual Property