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

H. Chen
D. Dhody
Huawei Technologies

July 4, 2014

The Applicability of the PCE to Computing Protection and Recovery Paths for Single Domain and Multi-Domain Networks.

draft-chen-pce-protection-applicability-05

Abstract

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

A link or node failure can significantly impact network services in large-scale networks. Therefore it is important to ensure the survivability of large scale networks which consist of various connections provided over multiple interconnected networks with varying technologies.

This document examines the applicability of the PCE architecture, protocols, and procedures for computing protection paths and restoration services, for single and multi-domain networks.

This document also explains the mechanism of Fast Re-Route (FRR) where a point of local repair (PLR) needs to find the appropriate merge point (MP) to do bypass path computation using PCE.

Status of This Memo

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

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

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

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

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction 4
  1.1. Domains 4
    1.1.1. Inter-domain LSPs 5
  1.2. Recovery 5
  1.3. Requirements Language 5
2. Terminology 5
3. Path Computation Element Architecture Considerations 7
  3.1. Online Path Computation 7
  3.2. Offline Path Computation 7
4. Protection Service Traffic Engineering 8
  4.1. Path Computation 8
  4.2. Bandwidth Reservation 8
  4.3. Disjoint Path 8
  4.4. Service Preemption 8
  4.5. Share Risk Link Groups 8
  4.6. Multi-Homing 8
    4.6.1. Ingress and Egress Protection 9
5. Packet Protection Applications 9
  5.1. Single Domain Service Protection 10
  5.2. Multi-domain Service Protection 10
  5.3. Backup Path Computation 10
  5.4. Fast Reroute (FRR) Path Computation 10
    5.4.1. Methods to find MP and calculate the optimal backup path 11
      5.4.1.1. Intra-domain node protection 12
      5.4.1.2. Boundary node protection 12
  5.5. Point-to-Multipoint Path Protection 15
6. Optical Protection Applications 16
  6.1. ASON Applicability 16
  6.2. Multi-domain Restoration 16
7. Path and Service Protection Gaps 16
8. Manageability Considerations 16
  8.1. Control of Function and Policy 16
  8.2. Information and Data Models 16
  8.3. Liveness Detection and Monitoring 16
  8.4. Verify Correct Operations 16
  8.5. Requirements On Other Protocols 16
  8.6. Impact On Network Operations 16
9. Security Considerations 16
10. IANA Considerations 16
11. Contributors 17
12. Acknowledgement 17
13. References 17
  13.1. Normative References 17
  13.2. Informative References 17

1. Introduction

Network survivability remains a major concern for network operators and service providers, particularly as expanding applications such as private and Public Cloud drive increasingly more traffic across longer ranges, to a wider number of users. A variety of well-known pre-planned protection and post-fault recovery schemes have been developed for IP, MPLS and GMPLS networks.

The Path Computation Element (PCE) [RFC4655] can be used to perform complex path computation in large single domain, multi-domain and multi-layered networks. The PCE can also be used to compute a variety of restoration and protection paths and services.

This document examines the applicability of the PCE architecture, protocols, and protocol extensions for computing protection paths and restoration services.

1.1. Domains

A domain can be defined as a separate administrative, geographic, or switching environment within the network. A domain may be further defined as a zone of routing or computational ability. Under these definitions a domain might be categorized as an Antonymous System (AS) or an Interior Gateway Protocol (IGP) area (as per [RFC4726] and [RFC4655]), or specific switching environment.

In the context of GMPLS, a particularly important example of a domain is the Automatically Switched Optical Network (ASON) subnetwork [G-8080]. In this case, computation of an end-to-end path requires the selection of nodes and links within a parent domain where some nodes may, in fact, be subnetworks. Furthermore, a domain might be an ASON routing area [G-7715]. A PCE may perform the path computation function of an ASON routing controller as described in [G-7715-2].

It is assumed that the PCE architecture should be applied to small inter-domain topologies and not to solve route computation issues across large groups of domains, I.E. the entire Internet.

Most existing protocol mechanisms for network survivability have focused on single-domain scenarios. Multi-domain scenarios are much more complex and challenging as domain topology information is typically not shared outside each specific domain.

Therefore multi-domain survivability is a key requirement for today's complex networks. It is important to develop more adaptive multi-domain recovery solutions for various failure scenarios.

1.1.1. Inter-domain LSPs

Three signaling options are defined for setting up an inter-area or inter-AS LSP [RFC4726]:

- o Contiguous LSP
- o Stitched LSP
- o Nested LSP

1.2. Recovery

Typically traffic-engineered networks such as MPLS-TE and GMPLS, use protection and recovery mechanisms based on the pre-established use of a packet or optical LSP and/or the availability of spare resources and the network topology.

1.3. Requirements Language

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 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC2119] significance.

2. Terminology

The following terminology is used in this document.

ABR: Area Border Router. Router used to connect two IGP areas (Areas in OSPF or levels in IS-IS).

ASBR: Autonomous System Border Router. Router used to connect together ASes of the same or different service providers via one or more inter-AS links.

BN: Boundary Node (BN). A boundary node is either an ABR in the context of inter-area Traffic Engineering or an ASBR in the context of inter-AS Traffic Engineering.

CPS: Confidential Path Segment. A segment of a path that contains nodes and links that the AS policy requires not to be disclosed outside the AS.

CSP: Communication Service Provide.

CSPF: Constrained Shorted Path First Algorithm.

ERO: Explicit Route Object.

FRR: Fast Re-Route.

IGP: Interior Gateway Protocol. Either of the two routing protocols, Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS).

Inter-area TE LSP: A TE LSP whose path transits through two or more IGP areas.

Inter-AS TE LSP: A TE LSP whose path transits through two or more ASs or sub-ASs (BGP confederations).

IS-IS: Intermediate System to Intermediate System.

LSP: Label Switched Path.

LSR: Label Switching Router.

MP: Merge Point. The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure.

OSPF: Open Shortest Path First.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PKS: Path Key Subobject. A subobject of an Explicit Route Object or Record Route Object that encodes a CPS so as to preserve confidentiality.

PLR: Point of Local Repair. The head-end LSR of a backup tunnel or a detour LSP.

RRO: Record Route Object.

RSVP: Resource Reservation Protocol.

SRLG: Shared Risk Link Group.

TE: Traffic Engineering.

TED: Traffic Engineering Database, which contains the topology and resource information of the domain. The TED may be fed by Interior Gateway Protocol (IGP) extensions or potentially by other means.

This document also uses the terminology defined in [RFC4655] and [RFC5440].

3. Path Computation Element Architecture Considerations

For the purpose of this document it is assumed that the path computation is the sole responsibility of the PCE as per the architecture defined in [RFC4655]. When a path is required the Path Computation Client (PCC) will send a request to the PCE. The PCE will apply the required constraints and compute a path and return a response to the PCC. In the context of this document it may be necessary for the PCE to co-operate with other PCEs in adjacent domains (as per BRPC [RFC5441]) or cooperate with the Parent PCE (as per RFC 6805).

A PCE may be used to compute end-to-end paths across single or multiple domains. Multiple PCEs may be dedicated to each area to provide sufficient path computation capacity and redundancy for each domain.

During path computation [RFC5440], a PCC request may contain backup LSP requirements in order to setup in the same time the primary and backup LSPs. This request is known as dependent path computations. A typical dependent request for a primary and backup service would request that the computation assign a set of diverse paths, so both services are disjointed from each other.

3.1. Online Path Computation

Online path computation is performed on-demand as nodes in the network determine that they need to know the paths to use for services.

3.2. Offline Path Computation

Offline path computation is performed ahead of time, before the LSP setup is requested. That means that it is requested by, or performed

as part of, a management application.

This method of computation allows the optimal placement of services and explicit control of services. A Communication Service Provide (CSP) can plan where new protection services will be placed ahead of time. Furthermore by computing paths offline specific scenarios can be considered and a global view of network resources is available.

Finally, offline path computation provides a method to compute protection paths in the event of a single, or multiple, link failures. This allows the placement of backup services in the event of catastrophic network failures.

4. Protection Service Traffic Engineering

4.1. Path Computation

This document describes how the PCE architecture defined in [RFC4655] may be utilized to compute protection and recovery paths for critical network services. In the context of this document (inter-domain) it may be necessary for the PCE to co-operate with other PCEs in adjacent domains (as per BRPC [RFC5441]) or cooperate with the Parent PCE (as per RFC 6805).

4.2. Bandwidth Reservation

4.3. Disjoint Path

Disjoint paths are required for end-to-end protection services. A backup service may be required to be fully disjoint from the primary service, link disjoint (allowing common nodes on the paths), or best-effort disjoint (allowing shared links or nodes when no other path can be found).

4.4. Service Preemption

4.5. Share Risk Link Groups

4.6. Multi-Homing

Networks constructed from multi-areas or multi-AS environments may have multiple interconnect points (multi-homing). End-to-end path computations may need to use different interconnect points to avoid single point failures disrupting primary and backup services.

Domain and path diversity may also be required when computing end-to-end paths. Domain diversity should facilitate the selection of paths that share ingress and egress domains, but do not share transit

domains. Therefore, there must be a method allowing the inclusion or exclusion of specific domains when computing end-to-end paths.

4.6.1. Ingress and Egress Protection

An end-to-end primary service carried by a primary TE LSP from a primary ingress node to a primary egress node may need to be protected against the failures in the ingress and the egress. In this case, a backup ingress and a backup egress are required, which are different from the primary ingress and the primary egress respectively. The backup ingress should be in the same domain as the primary ingress, and the backup egress should be in the same domain as the primary egress.

A source of the service traffic may be sent to both the primary ingress and the backup ingress (dual-homing). The source may not be in the same domain as the primary ingress and the backup ingress. When the primary ingress fails, the service traffic is delivered through the backup ingress.

A receiver of the service traffic may be connected to both the primary egress and the backup egress (dual-homing). The receiver may not be in the same domain as the primary egress and the backup egress. When the primary egress fails, the receiver gets the service traffic from the backup egress.

5. Packet Protection Applications

Network survivability is a key objective for CSPs, particularly as expanding revenue services (cloud and data center applications) are increasing exponentially.

Pre-fault paths are pre-computed and protection resources are reserved a priory for rapid recovery. In the event of a network failure on the primary path, the traffic is fast switched to the backup path. These pre-provisioned mechanisms are capable of ensuring protection against single link failures.

Post-fault restoration schemes are reactive and require a reactive routing procedure to set up new working paths in the event of a failure. Post fault restoration can significantly impact network services as they are typically impacted by longer restoration delays and cannot guarantee recovery of a service. However, they are much more network resource efficient and are capable of handling multi-failure situations.

5.1. Single Domain Service Protection

A variety of pre-planned protection and post-fault restoration recovery schemes are available for single domain MPLS and GMPLS networks, these include:

- o Path Recovery
- o Path Segment Recovery
- o Local Recovery (Fast Reroute)

5.2. Multi-domain Service Protection

Typically network survivability has focused on single-domain scenarios. By contrast, broader multi-domain scenarios are much more challenging as no single entity has a global view of topology information. As a result, multi-domain survivability is very important.

A PCE may be used to compute end-to-end paths across multi-domain environments using a per-domain path computation technique [RFC5152]. The so called backward recursive path computation (BRPC) mechanism [RFC5441] defines a PCE-based path computation procedure to compute inter-domain constrained LSPs.

5.3. Backup Path Computation

A PCE can be used to compute backup paths in the context of fast reroute protection of TE LSPs. In this model, all backup TE LSPs protecting a given facility are computed in a coordinated manner by a PCE. This allows complete bandwidth sharing between backup tunnels protecting independent elements, while avoiding any extensions to TE LSP signaling. Both centralized and distributed computation models are applicable. In the distributed case each LSR can be a PCE to compute the paths of backup tunnels to protect against the failure of adjacent network links or nodes.

5.4. Fast Reroute (FRR) Path Computation

As stated in [RFC4090], there are two independent methods (one-to-one backup and facility backup) of doing fast reroute (FRR). PCE can be used to compute backup path for both of the methods. Cooperating PCEs may be used to compute inter-domain backup path.

In case of one to one backup method, the destination MUST be the tail-end of the protected LSP. Whereas for facility backup, destination MUST be the address of the merge point (MP) from the

corresponding point of local repair (PLR). The problem of finding the MP using the interface addresses or node-ids present in Record Route Object (RRO) of protected path can be easily solved in the case of a single Interior Gateway Protocol (IGP) area because the PLR has the complete Traffic Engineering Database (TED). Thus, the PLR can unambiguously determine -

- o The MP address regardless of RRO IPv4 or IPv6 sub-objects (interface address or LSR ID).

- o Does a backup tunnel intersecting a protected TE LSP on MP node exist? This is the case where facility backup tunnel already exists either due to another protected TE LSP or it is pre-configured.

It is complex for a PLR to find the MP in case of boundary node protection for computing a bypass path because the PLR doesn't have the full TED visibility. When confidentiality (via path key) [RFC5520] is enabled, finding MP is very complex.

This document describes the mechanism to find MP and to setup bypass tunnel to protect a boundary node.

5.4.1. Methods to find MP and calculate the optimal backup path

The Merge Point (MP) address is required at the PLR in order to select a bypass tunnel intersecting a protected Traffic Engineering Label Switched Path (TE LSP) on a downstream LSR.

Some implementations may choose to pre-configure a bypass tunnel on PLR with destination address as MP. MP's Domain to be traversed by bypass path can be administratively configured or learned via some other means (ex Hierarchical PCE (HPCE) RFC 6805). Path Computation Client (PCC) on PLR can request its local PCE to compute bypass path from PLR to MP, excluding links and node between PLR and MP. At PLR once primary tunnel is up, a pre-configured bypass tunnel is bound to the primary tunnel, note that multiple bypass tunnels can also exist.

Most implementations may choose to create a bypass tunnel on PLR after primary tunnel is signaled with Record Route Object (RRO) being present in primary path's Resource Reservation Protocol (RSVP) Path Reserve message. MP address has to be determined (described below) to create a bypass tunnel. PCC on PLR can request its local PCE to compute bypass path from PLR to MP, excluding links and node between PLR and MP.

5.4.1.1. Intra-domain node protection

```
[R1]----[R2]----[R3]----[R4]---[R5]
          \             /
          [R6]--[R7]--[R8]

Protected LSP Path: [R1->R2->R3->R4->R5]
Bypass LSP Path:    [R2->R6->R7->R8->R4]
```

Figure 1: Node Protection for R3

In Figure 1, R2 has to build a bypass tunnel that protects against the failure of link [R2->R3] and node [R3]. R2 is PLR and R4 is MP in this case. Since, both PLR and MP belong to the same area. The problem of finding the MP using the interface addresses or node-ids can be easily solved. Thus, the PLR can unambiguously find the MP address regardless of RRO IPv4 or IPv6 sub-objects (interface address or LSR ID) and also determine whether a backup tunnel intersecting a protected TE LSP on a downstream node (MP) already exists.

TED on PLR will have the information of both R2 and R4, which can be used to find MP's TE router IP address and compute optimal backup path from R2 to R4, excluding link [R2->R3] and node [R3].

Thus, RSVP-TE can signal bypass tunnel along the computed path.

5.4.1.2. Boundary node protection

5.4.1.2.1. Area Boundary Router (ABR) node protection

```
                    |
       PCE-1        |     PCE-2
                    |
      IGP area 0 |  IGP area 1
                    |
                    |
[R1]----[R2]----[R3]----[R4]---[R5]
        \        |       /
        [R6]--[R7]--[R8]
                    |
                    |
                    |

Protected LSP Path: [R1->R2->R3->R4->R5]
Bypass LSP Path:    [R2->R6->R7->R8->R4]
```

Figure 2: Node Protection for R3 (ABR)

In Figure 2, cooperating PCE(s) (PCE-1 and PCE-2) have computed the primary LSP Path [R1->R2->R3->R4->R5] and provided to R1 (PCC).

R2 has to build a bypass tunnel that protects against the failure of link [R2->R3] and node [R3]. R2 is PLR and R4 is MP. Both PLR and MP are in different area. TED on PLR doesn't have the information of R4.

The problem of finding the MP address in a network with inter-domain TE LSP is solved by inserting a node-id sub-object [RFC4561] in the RRO object carried in the RSVP Path Reserve message. PLR can find out the MP from the RRO it has received in Path Reserve message from its downstream LSR.

But the computation of optimal backup path from R2 to R4, excluding link [R2->R3] and node [R3] is not possible with running of Constrained Shortest Path First (CSPF) algorithm locally at R2. PCE can be used to compute backup path in this case. R2 acting as PCC on PLR can request PCE-1 to compute bypass path from PLR(R2) to MP(R4), excluding link [R2->R3] and node [R3]. PCE MAY use inter-domain path computation mechanism (like HPCE (RFC 6805) etc) when the domain information of MP is unknown at PLR. Further, RSVP-TE can signal bypass tunnel along the computed path.

5.4.1.2.2. Autonomous System Border Router (ASBR) node protection

```
                          |        |
                    PCE-1 |        | PCE-2
                          |        |
                   AS 100 |        | AS 200
                          |        |
                          |        |
             [R1]----[R2]-------[R3]---------[R4]---[R5]
                          |\       |            /
                          | +-----[R6]--[R7]--[R8]
                          |        |
                          |        |

        Protected LSP Path: [R1->R2->R3->R4->R5]
        Bypass LSP Path:    [R2->R6->R7->R8->R4]
```

Figure 3: Node Protection for R3 (ASBR)

In Figure 3, Links [R2->R3] and [R2->R6] are inter-AS links. IGP extensions ([RFC5316] and [RFC5392]) describe the flooding of

inter-AS TE information for inter-AS path computation. Cooperating PCE(s) (PCE-1 and PCE-2) have computed the primary LSP Path [R1->R2->R3->R4->R5] and provided to R1 (PCC).

R2 is PLR and R4 is MP. Both PLR and MP are in different AS. TED on PLR doesn't have the information of R4.

The address of MP can be found using node-id sub-object [RFC4561] in the RRO object carried in the RSVP Path Reserve message. And Cooperating PCEs could be used to compute the inter-AS bypass path. Thus ASBR boundary node protection is similar to ABR protection.

5.4.1.2.3. Boundary node protection with Path-Key Confidentiality

[RFC5520] defines a mechanism to hide the contents of a segment of a path, called the Confidential Path Segment (CPS). The CPS may be replaced by a path-key that can be conveyed in the PCE Communication Protocol (PCEP) and signaled within in a Resource Reservation Protocol TE (RSVP-TE) explicit route object.

[RFC5553] states that, when the signaling message crosses a domain boundary, the path segment that needs to be hidden (that is, a CPS) MAY be replaced in the RRO with a PKS. Note that RRO in Resv message carries the same PKS as originally signaled in the ERO of the Path message.

5.4.1.2.3.1. Area Boundary Router (ABR) node protection

```
                          |
                 PCE-1    |     PCE-2
                          |
                IGP area 0 |  IGP area 1
                          |
                          |
          [R1]----[R2]----[R3]----[R4]---[R5]---[R9]
                  \       |       /
                  [R6]--[R7]--[R8]
                          |
                          |
                          |
```

Figure 4: Node Protection for R3 (ABR) and Path-Key

In Figure 4, when path-key is enabled, cooperating PCE(s) (PCE-1 and PCE-2) have computed the primary LSP Path [R1->R2->R3->PKS->R9] and provided to R1 (PCC).

When the ABR node (R3) replaces the CPS with PKS (as originally signaled) during the Reserve message handling, it MAY also add the immediate downstream node-id (R4) (so that the PLR (R2) can identify the MP (R4)). Further the PLR (R2) SHOULD remove the MP node-id (R4) before sending the Reserve message upstream to head end router.

Once MP is identified, the backup path computation using PCE is as described earlier. (Section 5.4.1.2.1)

5.4.1.2.3.2. Autonomous System Border Router (ASBR) node protection

```
                        |        |
                  PCE-1 |        | PCE-2
                        |        |
                 AS 100 |        | AS 200
                        |        |
                        |        |
            [R1]----[R2]-------[R3]---------[R4]---[R5]
                        |\       |            /
                        | +-----[R6]--[R7]--[R8]
                        |        |
                        |        |
```

Figure 5: Node Protection for R3 (ASBR)

The address of MP can be found using the same mechanism as explained above. Thus ASBR boundary node protection is similar to ABR protection.

5.5. Point-to-Multipoint Path Protection

A PCE utilizing the extensions outlined in [RFC6006] (Extensions to PCEP for Point-to-Multipoint Traffic Engineering Label Switched Paths), can be used to compute point-to-multipoint (P2MP) paths. A PCC requesting path computation for a primary and backup path can request that these dependent computations use diverse paths. Furthermore, the specification also defines two new options for P2MP path dependent computation requests. The first option allows the PCC to request that the PCE should compute a secondary P2MP path tree with partial path diversity for specific leaves or a specific source-to-leaf (sub-path to the primary P2MP path tree. The second option, allows the PCC to request that partial paths should be link direction diverse.

6. Optical Protection Applications

6.1. ASON Applicability

6.2. Multi-domain Restoration

7. Path and Service Protection Gaps

8. Manageability Considerations

8.1. Control of Function and Policy

TBD

8.2. Information and Data Models

TBD

8.3. Liveness Detection and Monitoring

TBD

8.4. Verify Correct Operations

TBD

8.5. Requirements On Other Protocols

TBD

8.6. Impact On Network Operations

TBD

9. Security Considerations

This document does not introduce new security issues. However, MP's node-id is carried as subobject in RRO across domain. This relaxation is required to find MP in case of BN protection. The security considerations pertaining to the [RFC3209], [RFC4090] and [RFC5440] protocols remain relevant.

10. IANA Considerations

This document makes no requests for IANA action.

11. Contributors

Venugopal Reddy Kondreddy
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: venugopalreddyk@huawei.com

12. Acknowledgement

We would like to thank Daniel King, Udayashree Palle, Sandeep Boina & Reeja Paul for their useful comments and suggestions.

13. References

13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

13.2. Informative References

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC4561] Vasseur, J., Ali, Z., and S. Sivabalan, "Definition of a Record Route Object (RRO) Node-Id Sub-Object", RFC 4561, June 2006.

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

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

[RFC5152] Vasseur, JP., Ayyangar, A., and R. Zhang, "A Per-Domain

Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.

[RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, December 2008.

[RFC5392] Chen, M., Zhang, R., and X. Duan, "OSPF Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5392, January 2009.

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

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

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

[RFC5553] Farrel, A., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, May 2009.

[RFC6006] Zhao, Q., King, D., Verhaeghe, F., Takeda, T., Ali, Z., and J. Meuric, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC 6006, September 2010.

[RFC6805] King, D. and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, November 2012.

[G-7715] ITU-T, "ITU-T Recommendation G.7715 (2002), Architecture and Requirements for the Automatically Switched Optical Network (ASON).".

[G-7715-2] ITU-T, "ITU-T Recommendation G.7715.2 (2007), ASON routing architecture and requirements for remote route query.".

[G-8080] ITU-T, "ITU-T Recommendation G.8080/Y.1304, Architecture for the automatically switched optical network (ASON).".

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
USA

EMail: huaimo.chen@huawei.com

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruv.dhody@huawei.com