Workgroup: PCE Working Group
Internet-Draft:
draft-zhang-pce-resource-sharing-15
Published: 24 October 2021
Intended Status: Standards Track
Expires: 27 April 2022
Authors:
X. Zhang, Huawei Technologies
H. Zheng, Huawei Technologies
O.G. Gonzales de Dios, Telefonica
V.L. Lopez, Nokia
Y.X. Xu, CAICT

Extensions to the Path Computation Element Protocol (PCEP) to Support Resource Sharing-based Path Computation

Abstract

Resource sharing in a network means two or more Label Switched Paths (LSPs) use common pieces of resource along their paths. This can help save network resources and is useful in scenarios such as LSP recovery or when two LSPs do not need to be active at the same time. A Path Computation Element (PCE) is responsible for path computation with such requirement.

Existing extensions to the Path Computation Element Protocol (PCEP) allow one path computation request for an LSP to be associated with other (existing) LSPs through the use of the PCEP Association Object.

This document extends PCEP in order to support resource-sharing-based path computation as another use of the Association Object to enable better efficiency in the computation and in the resultant paths and network resource usage.

Status of This Memo

Table of Contents

1. Introduction

A Path Computation Element (PCE) is a way to provide path computation function, and it is especially useful in the scenarios where complex constraints and/or a demanding amount of computation resource are required [RFC4655]. The development of PCE standardization has evolved from stateless to stateful. A stateful PCE has access to the LSP database information of the networks it serves as a computation engine [RFC8231]. Unless specified, this document assumes a PCE mentioned is a stateful PCE..

Resource sharing denotes that two or more Label Switched Paths (LSPs) share common pieces of resource, (such as a common time slot of a link in an Optical Transport Network (OTN)). This is usually useful in the scenario where only one of the LSPs is active and the benefit is to save network resources. A simple example of this is dynamically calculating a recovery LSP for an existing LSP undergoing a link failure. Note that resource sharing can be worked out using a stateless PCE, but the mechanism may be complex and is out the scope of this document.

This document considers the requirement that a new LSP may request for resource sharing with one or multiple existing LSPs. Furthermore, if there is resource sharing between a new LSP and existing an LSP, the two LSPs cannot be used to carry traffic simultaneously, the new LSP will take over the traffic from the existing LSP.

In a single domain, this is a common requirement in the recovery cases especially in order to increase traffic resilience against failure while reducing the amount of network resource used for recovery purposes [RFC4428]

The current protocol supporting the communication between a PCE and a Path Computation Client (PCC), i.e. PCE Protocol (PCEP), allows for re-optimization of an existing LSP [RFC5440]. This is achieved by setting the R bit in the Request Parameter (RP) object, together with some additional information if applicable, in the Path Computation Request (PCReq) message sent from a PCC to the PCE. To support this type of resource sharing, a PCC needs to ask a PCE to compute a new path with the constraints of sharing resource with one or multiple existing LSPs. It is worth noting the "resource sharing" in this draft not only means one LSP re-using the same links of another LSP, but also the same slice of bandwidth in the network. This may occur when an LSP is required for re-routing, or online re-optimization. Current PCEP specifications do not provide such function. More specifically, this document describes the resource sharing issue during the procedure when a new LSP is required to replace an existing LSP for use together with Make-before-break (MBB) described in [RFC3209].

As mentioned in [RFC8231], the PLSP-ID provides a unique identifier for an LSP during a PCEP session between PCC and PCE. Such identification is helpful in supporting the resource sharing requirement for stateful PCEs because it greatly simplifies the operation of a PCC. Instead of the PCC determining all the resources to be shared, the PCC can request that the PCE share the resources of a specific LSP: the stateful PCE is able to determine those resource itself.

Resource sharing can also be required in an inter-layer PCEP session. This is similar to the previous requirement. However, it is more complex and therefore deserves a more detailed explanation here.

In a multi-layer network, LSPs in a lower layer are used to carry higher-layer LSPs across the lower-layer network [RFC5623]. Therefore, the resource sharing constraints in the higher layer might actually relate to resource sharing in the lower layer. Thus, it is useful to consider how this can be achieved and whether additional extensions are needed using the models defined in [RFC5623].

In the next sections, use cases are provided to show what information needs to be exchanged to fulfill these requirements. This memo then provides extensions to PCEP to enable this function.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Motivation

2.1. Single Domain Use Case

There are two potential cases that request resource to be shared: restoration and re-optimization. Figure 1 shows a single domain network with a stateful PCE, and is used as an example for the resource sharing application.

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


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

Figure 1: Figure 1: A Single Domain Example

LSP0 (existing): N1-N2-N3.

LSP1 (restoration): N1-N2-N4-N3.

LSP2 (re-optimization): N1-N5-N4-N3.

For the failure restoration, we can assume a working LSP (LSP0) exists in the network. When there is failure on the link N2-N3, it is desired to set up a restoration path for this working LSP. Suppose N1 serves as the PCC and sends a request to the stateful PCE for such an LSP. Besides the head-end and tail-end node of the working LSP, N1 may also need to check what policy should be applied for the restoration. For example, it may evaluate resource sharing and prefer to share as much resource with the working LSP as possible and specify this policy as a special object in the PCReq message. Given such policy, a probable outcome from the path computation would be LSP1, which shares the link 'N1-N2' with the existing LSP. The LSP1 will be set up by PCC via either PCInitiate of RSVP.

Re-optimization does not usually result from a specific failure in the network, but takes place on a stable network when more optimal paths may have become available. Thus switching from the existing LSP to the new LSP happens with live traffic. An example can be found in Figure 1 without failure on the link N2-N3. Instead, an online re-optimization is needed for the working LSP (LSP0) from the stateful PCE. In such cases, the best choice is to set up a backup LSP for the working LSP with totally separate routing (for example, LSP2), and move the traffic to that backup LSP. After that, the working LSP can be torn down, which will not result in any interruption during the optimization procedure. This can actually be implemented with existing PCEP mechanisms. However, if there is no such separate path, existing PCEP mechanisms will return an error. A secondary option for this case is to set up an LSP and complete re-optimization with resource sharing, even if some interruption is introduced.

In the example from Figure 1 it is assumed that the restored LSP or re-optimized LSP have the same source and destination nodes. But in some applications there is no restriction for this assumption, i.e., after an LSP is failed, it can be restored as a new LSP with different source/destination.

In the use cases above it is also assumed that the characteristics of the restored LSP or re-optimized LSP are unchanged. However, it is possible to have parameter changes during the resource sharing computation. For example, the bandwidth of the request LSP may be different from the existing LSP, while resource sharing is still preferred by the PCC. The PCE should consider the sharing request together with the policy and available resources in the network. Details can be found in Section 3.3.

Conversely to resource sharing, it may also be required to apply a disjoint constraint for the path computation. [RFC8800] discusses the solution under such a scenario, which is a companion work to this document.

2.2. Multiple Layers/Domains Use Case

As Discussed in Section 3 of [RFC5623], there are three models for inter-layer path computation. They are single PCE computation, multiple PCE with inter-PCE communication, and multiple PCE without inter-PCE communication. For the single PCE computation, the process would be similar to that of the use case in Section 2.1.

An inter-layer path computation example is shown in Figure 2. Assume an LSP (LSP1: H2-H3) has been established already, visible as H2-H3 from the view of the higher-layer PCE, and as H2-L1-L2-H3 from the global view (or from the view of the lower-layer PCE). A new request is received by H2 to establish a new LSP (LSP2: from H2 to H5), given the constraint that it can share resources with LSP1. This requirement is possible if only one of the LSPs needs to be active and resource sharing is the target.

```
                                                            -----
                          .................................| LSR |
                        .:                                 | H5  |
                      .:                                   /-----
                    .:                                    /   |
    -----    -----.:                       -----    -----/    |
   | LSR |--| LSR |.......................| LSR |--| LSR |   /
   | H1  |  | H2  |                       | H3  |  | H4  |  /
    -----    -----\                       /-----    -----  /
                   \                     /                /
                    \                   /                /
                     \                 /                /
                      \               /                /
                       \-----   -----/                /
                       | LSR |-| LSR |               /
                       | L1  | | L2  |              /
                        -----   -----\             /
                          |           \           /
                          |            \         /
                          |             \       /
                        -----            \-----/
                       | LSR |-----------| LSR |
                       | L3  |           | L4  |
                        -----             -----
```

Figure 2: Figure 2: A Two-layer Network Example

If the model of multiple PCEs with inter-PCE communication is employed, the path computation request sent by H2 to higher-layer PCE will be forwarded to lower-layer PCE since there is no resource readily available in the higher layer. So it leaves the lower-layer PCE to compute a path in the lower layer in order to support the higher layer request. In this case, the lower-layer PCE is required to compute a path between H2 and H5 under the constraint that it can share the resource with that of LSP1. At this moment the lower-layer PCE has knowledge of the mapping relationship between the higher-layer link H2-H3 and the lower layer link L1-L2, and therefore can convert the resource to be shared from higher layer to lower layer. So when the lower-layer PCE computes the path for LSP2, it can consider the resource used by L1-L2 as available with higher priority. For example, the lower-layer PCE may choose H2-L1-L2-L4-H5 as the computation result. On the other hand, if the path computation policy is to have a separate path with LSP1, the lower-layer PCE may choose H2-L1-L3-L4-H5.

During this procedure the higher-layer PCE can only use information about LSP1 (such as its five-tuple LSP information). An issue to solve is how the lower-layer PCE can resolve this information to the actual resource usage in its own layer, i.e. the lower layer. This could be solved by the edge LSR (L1) reporting this higher-lower LSP correlation to the lower-layer PCE as part of the LSP information during the LSP state synchronization process. If needed, it can be updated later when there is a change in this information. Alternatively, the lower-layer PCE can get this information from other sources, such as a network management system, where this information should be stored.

If the model of multiple PCEs without inter-PCE communication is employed, the path computation request in the lower layer will be initiated by the border LSR node, i.e., L1. The process would be similar to that of the previous scenario. A point worth noting is that the border LSR node may be able to resolve the higher layer LSP information itself, such as by mapping it to the corresponding LSP in the lower layer, in this way the lower-layer PCE does not need to perform this function. Otherwise, the mapping method mentioned above can still be used.

2.3. Bulk Path Computation Use Case

There is a potential need for resource sharing during bulk path computation, especially the processing of the "sticky resources" in [RFC7399]. It would be useful to specify the resources that can be shared among different paths, i.e., the bandwidth information.

Considering the H-PCE architecture in [RFC8751], when the parent PCE asks for a single path across a few domains, such a request may become a bulk path computation to a certain child PCE. Figure 3 shows an example of 3 domains. The parent PCE will select one of these path for establishment.

```
                           +-------+
                          /| P-PCE |\
                         / +---+---+ \
                        /       |     \
                       /        |      \
                      /         |       \
                     /          |        \
                    /           |         \
                   /            |          \
            +-----/+        +---+---+      +\------+
            |C-PCE1|        |C-PCE2 |      |C-PCE3 |
            +------+        +-------+      +-------+
             /                  |                  \
 ---------------      -----------------------       -------------
/                \   /                       \    /               \
| +---+     +---+ |  |  +---+   +---+   +---+ |  | +---+    +---+ |
| | A +-----+ B +-+--+--+ D +---+ E +---+ H +-+--+-+ J +----+ L | |
| +-\-+     +---+ |  |  +---+   +---+   +--\+ |  | +---+    +-/-+ |
|     \           |  |          /           \ |  |           /    |
|      \          |  |         /             \|  |          /     |
|        \  +---+ |  |  +---+ /               |\\|    +---+/      |
|          \+ C +-+--+--+ G +/                |  |----| K |       |
\           +---+/   \  +---+                /    \   +---+      /
 ----------------     -----------------------      --------------
```

Figure 3: Figure 3: Bulk Request example with Hierarchical PCEs

A 3-domain example is shown in Figure 3, with the hierarchical PCE architecture. In this example nodes A/B/C belong to domain 1, nodes D/E/G/H belong to domain 2, and nodes J/K/L belong to domain 3. Inter-domain links are B-D/C-G between domains 1 and 2, and H-J/H-K between domains 2 and 3. Given a path computation request from A to L, a bulk request from P-PCE would be helpful to understand whether it is possible to have different combinations on the inter-domain links. However, the resources on some specific links become 'sticky' and have to be indicated as 'sharing allowed' to avoid unnecessary resource competition. For example, both the route A-B-D-E-H-J-L and A-C-G-E-H-K-L are qualified, but these routes are competing for the resource on the link E-H and cannot be established simultaneously, so there must be one route failed to be reported to P-PCE. Given the indication of allowing resource sharing on the link E-H, both of these routes can be reported for P-PCE's decision, and there will not be any competition as the P-PCE understands that only one path needs to be set up.

3. Extensions to PCEP

3.1. Association Group and Type

According to the definition in [RFC8697], the association group is used to associate multiple LSPs into one group for further path computation considerations, such as disjointness and resource sharing, in the messages when requesting path computation. An association ID will be used to identify the resource sharing group. An association type that described disjointness has been defined in [RFC8800]. In this document, a new association type is defined as follows:

*Association type = TBD1 ("Sharing Association Type").

A sharing group should have multiple LSPs. The number of LSPs and the criteria for how LSPs share among each other are dependent on local policy.

3.2. Resource Sharing TLV

The PCEP Resource Sharing group MAY carry the following TLV. It MAY be carried within a PCReq message from the network element (or other PCCs) so as to indicate the desired resource sharing requirements to be applied by the stateful PCE during path computation.

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |         Type = [TBD2]         |            Length             |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                      Flags                            |B|S|N|L|
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                        Optional TLVs                          |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The following flags are defined:

* L (Link share) bit: when set, this flag indicates that the PCE should prioritize the links shared by existing LSPs within the sharing group for path computation. The existing LSP identifier and its available link identifiers can be contained in the optional TLVs.

* N (Node share) bit: when set, this flag indicates that the PCE should prioritize the nodes shared by existing LSPs within the sharing group for path computation. The existing LSP identifier and its available node identifiers can be contained in the optional TLVs.

* S (SRLG share) bit: when set, this flag indicates that the PCE should set the SRLG (Shared Risk Link Group) of the computed LSP to the same as existing LSPs within the sharing group for path computation. The existing LSP identifier and SRLG information can be contained in the optional TLVs.

* B (Bandwidth share) bit: when set, this flag indicates that the PCE should prioritize the bandwidth to be shared by LSPs within the sharing group for bulk path computation. The LSP identifiers can be contained in the optional TLVs.

It is worth noting that there can be multiple flags set which may conflict with each other. In this scenario, the result for path computation may not be unique, and is dependent on the implementation. The selection among multiple computation results is out of the scope of this document.

3.3. Processing Rules

To request a path allowing resource sharing with one or multiple existing LSPs, a PCC includes a Resource Sharing TLV in the Association Group Object in any kind of path computation request message, such as the PCReq, PCUpd, or PCInitiate messages specified in [RFC8231] and [RFC8281].

On receipt of a PCEP message with a Resource Sharing TLV, a stateful PCE MUST proceed as follows:

* If the Resource Sharing TLV is unknown/unsupported, the PCE will follow procedures defined in [RFC5440]. That is, the PCE sends a PCErr message with error type 26 (Association Error) and error value 6 (Association Information Mismatch), and the related path computation request is discarded.

* If the Resource Sharing TLV is extracted correctly, the PCE MUST apply the requested resource sharing requirement, i.e., try to share as much resource as possible with the LSP specified in Resource Sharing TLV.

The procedure of setting flags follows the rules defined in Section 3.1. The flags in the Resource Sharing TLV may be locally configured on the requesting nodes via external entities, such as a network management system or the entity that imposes the resource sharing requirement.

It is worth noting that the Resource Sharing TLV can be used together with other path indication objects like the IRO/XRO, with different objectives. The first difference is, the use of the Resource Sharing TLV is to set up an alternative path, instead a new path. It is also dependent on the knowledge held be the PCC, e.g., if the PCC has full knowledge of the path information and has a strong preference on the route, it may send the request message with an IRO to specify the route. On the other hand, if the PCC does not know how the path should go but just wants to set up a new LSP to replace the old one, it may use the Resource Sharing TLV instead of an IRO. The second difference is that the Resource Sharing TLV is a loose requirement. For example, if the constraint specified in an IRO/XRO in an A-Z path computation request cannot be satisfied, the reply message from PCE to PCC would be unsuccessful. However it is still possible to have a path from the A-Z. If the target node/link/SRLG/Bandwidth is set in the Resource Sharing TLV rather than an IRO, the PCE may feedback a path from A-Z that does not share the target specified in the Resource Sharing TLV.

4. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC7942].

Currently the authors are not aware of any implementations.

5. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440] and [RFC8231] apply to the PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

5.1. Control of Function and Policy

A PCE or PCC implementation MUST allow operator-configured associations and SHOULD allow setting of the resource sharing TLV (Section 3.2) as described in this document.

5.2. Information and Data Models

An implementation SHOULD allow the operator to view the resource sharing configured or created dynamically. Further implementation SHOULD allow to view resource sharing associations reported by each peer, and the current set of LSPs in the association. The PCEP YANG module [I-D.ietf-pce-pcep-yang] includes association groups information.

5.3. Liveness Detection and Monitoring

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

5.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440] and [RFC8231].

5.5. Requirements on Other Protocols

Mechanisms defined in this document do not imply any new requirements on other protocols. The configuration on local policy may be accomplished by other protocols, such as Netconf.

5.6. Impact on Network Operations

Mechanisms defined in [RFC5440] and [RFC8231] also apply to PCEP extensions defined in this document.

6. Security Considerations

Security of PCEP is discussed in [RFC5440] and [RFC6952]. The extensions in this document do not change the fundamentals of security for PCEP.

However, the introduction of the Resource Sharing TLV in the Association Group Object provides a vector that may be used to probe for information from a network. For example, a PCC that wants to discover the path of an LSP with which it is not involved can issue a request message with a Resource Sharing TLV and may be able to get back quite a lot of information about the path of the LSP through issuing multiple such requests for different endpoints and analyzing the received results. To protect against this, a PCE SHOULD be configured with access and authorization controls such that only authorized PCCs (for example, those within the network) can make computation requests, only specifically authorized PCCs can make requests for resource sharing, and such requests relating to specific LSPs are further limited to a select few PCCs. How such access controls and authorization is managed is outside the scope of this document, but it will at the least include Access Control Lists.

Furthermore, a PCC must be aware that setting up an LSP that shares resources with another LSP may be a way of attacking the other LSP, for example by depriving it of the resources it needs to operate correctly. Thus it is important that, both in PCEP and the associated signaling protocols, only authorized resource sharing is allowed.

7. IANA Considerations

7.1. Association Object Type Indicators

IANA maintains a registry called the "Path Computation Element Protocol (PCEP) Numbers" registry with a subregistry called the "Association Type Field" subregistry. IANA is requested to make an assignment from that subregistry as follows:

Object Class	Name	Object Type	Reference
TBD1	Sharing-group	Association Type	[this document]

7.2. PCEP TLV Definitions

This document defines the following TLVs to support the resource sharing scenario:

Value	Name	Reference
TBD2	Resource-sharing TLV	[this document]

IANA is requested to allocate the following bit numbers in the flag spaces of Resource-sharing TLV:

Bit	Flag name	Reference
31	Link Share	[this document]
30	Node Share	[this document]
29	SRLG Share	[this document]
28	Bandwidth Share	[this document]

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

[RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE

Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/info/rfc8281>.

[RFC8697] Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)", RFC 8697, DOI 10.17487/RFC8697, January 2020, <https://www.rfc-editor.org/info/rfc8697>.

[RFC8800] Litkowski, S., Sivabalan, S., Barth, C., and M. Negi, "Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling", RFC 8800, DOI 10.17487/RFC8800, July 2020, <https://www.rfc-editor.org/info/rfc8800>.

8.2. Informational References

[I-D.ietf-pce-pcep-yang] Dhody, D., Hardwick, J., Beeram, V. P., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-16, 22 February 2021, <https://www.ietf.org/archive/id/draft-ietf-pce-pcep-yang-16.txt>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC4428] Papadimitriou, D., Ed. and E. Mannie, Ed., "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", RFC 4428, DOI 10.17487/RFC4428, March 2006, <https://www.rfc-editor.org/info/rfc4428>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, DOI 10.17487/RFC5623, September 2009, <https://www.rfc-editor.org/info/rfc5623>.

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <https://www.rfc-editor.org/info/rfc6952>.

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://www.rfc-editor.org/info/rfc7942>.

Authors' Addresses

Xian Zhang
Huawei Technologies
China

Email: zhang.xian@huawei.com

Haomian Zheng
Huawei Technologies
H1, Xiliu Beipo Village, Songshan Lake,
Dongguan
Guangdong, 523808
China

Email: zhenghaomian@huawei.com

Oscar Gonzales de Dios
Telefonica
Spain

Email: oscar.gonzalezdedios@telefonica.com

Victor Lopez
Nokia
Spain

Email: victor.lopez@nokia.com

Yunbin Xu
CAICT
China

Email: xuyunbin@caict.ac.cn