

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2018

J. Dong
M. Chen
D. Dhody
Huawei Technologies
J. Tantsura
Individual
K. Kumaki
KDDI Corporation
T. Murai
Furukawa Network Solution Corp.
July 3, 2017

BGP Extensions for Path Computation Element (PCE) Discovery
draft-dong-pce-discovery-proto-bgp-07

Abstract

In networks where a Path Computation Element (PCE) is used for path computation, it is desirable for the Path Computation Clients (PCCs) to discover dynamically and automatically a set of PCEs along with certain information relevant for PCE selection. RFC 5088 and RFC 5089 define the PCE discovery mechanisms based on Interior Gateway Protocols (IGP). This document defines extensions to BGP for the advertisement of PCE Discovery information. The BGP based PCE discovery mechanism is complementary to the existing IGP based mechanisms.

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 RFC 2119 [RFC2119].

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 4, 2018.

Copyright Notice

Copyright (c) 2017 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	2
2. Carrying PCE Discovery Information in BGP	3
2.1. PCE NLRI	3
2.1.1. PCE Descriptors	4
2.2. PCE Attribute TLVs	5
2.2.1. PCE Domain TLV	6
2.2.2. Neighbor PCE Domain TLV	6
3. Operational Considerations	7
4. IANA Considerations	7
5. Security Considerations	7
6. Contributors	7
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Authors' Addresses	9

1. Introduction

In networks where a Path Computation Element (PCE) is used for path computation, it is desirable for the Path Computation Clients (PCCs) to discover dynamically and automatically a set of PCEs along with certain information relevant for PCE selection. [RFC5088] and [RFC5089] define the PCE discovery mechanisms based on Interior Gateway Protocols (IGP). When PCCs are LSRs participating in the IGP (OSPF or IS-IS), and PCEs are either LSRs or servers also participating in the IGP, an effective mechanism for PCE discovery within an IGP routing domain consists of utilizing IGP advertisements.

[RFC4674] presents a set of requirements for a PCE discovery mechanism. This includes the discovery by a PCC of a set of one or more PCEs which may potentially be in some other domains. This is a desirable function in the case of inter-domain path computation. For example, Backward Recursive Path Computation (BRPC) [RFC5441] can be used by cooperating PCEs to compute an inter-AS path, in which case the discovery of PCE as well as the domain information is useful.

BGP has been extended for north-bound distribution of routing and TE information to PCE [RFC7752] and [I-D.ietf-idr-te-pm-bgp]. Similarly this document extends BGP to also carry the PCE discovery information.

This document defines extensions to BGP to allow a PCE to advertise its location, along with some information useful to a PCC for the PCE selection, so as to satisfy dynamic PCE discovery requirements set forth in [RFC4674].

This specification contains two parts: definition of a new BGP-LS NLRI [RFC7752] that describes PCE information and definition of PCE Attribute TLVs as part of BGP-LS attributes.

2. Carrying PCE Discovery Information in BGP

2.1. PCE NLRI

The PCE discovery information is advertised in BGP UPDATE messages using the MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760]. The "Link- State NLRI" defined in [RFC7752] is extended to carry the PCE information. BGP speakers that wish to exchange PCE discovery information MUST use the BGP Multiprotocol Extensions Capability Code (1) to advertise the corresponding (AFI, SAFI) pair, as specified in [RFC4760].

The format of "Link-State NLRI" is defined in [RFC7752]. A new "NLRI Type" is defined for PCE Information as following:

- o Type = TBD1: PCE NLRI

The format of PCE NLRI is shown in the following figure:

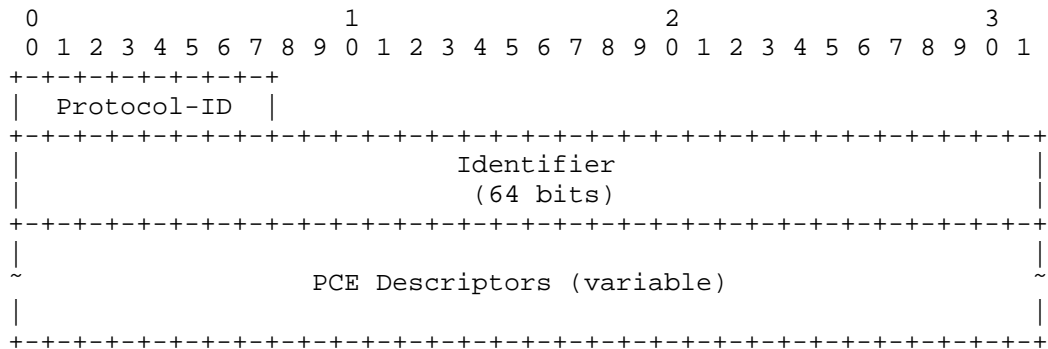


Figure 1. PCE NLRI

The 'Protocol-ID' field is defined in [RFC7752], to be set to the appropriate value that indicates the source of the PCE information. If BGP speaker and PCE are co-located, the Protocol-ID SHOULD be set to "Direct". If PCE information to advertise is configured at the BGP speaker, the Protocol-ID SHOULD be set to "Static configuration".

As defined in [RFC7752], the 64-Bit 'Identifier' field is used to identify the "routing universe" where the PCE belongs.

2.1.1.1. PCE Descriptors

The PCE Descriptor field is a set of Type/Length/Value (TLV) triplets. The format of each TLV is as per Section 3.1 of [RFC7752]. The PCE Descriptor TLVs uniquely identify a PCE. The following PCE descriptor are defined -

Codepoint	Descriptor TLV	Length
TBD2	IPv4 PCE Address	4
TBD3	IPv6 PCE Address	16

Table 1: PCE Descriptors

The PCE address TLVs specifies an IP address that can be used to reach the PCE. The PCE-ADDRESS Sub-TLV defined in [RFC5088] and [RFC5089] is used in the OSPF and IS-IS respectively. The format of the PCE address TLV are -

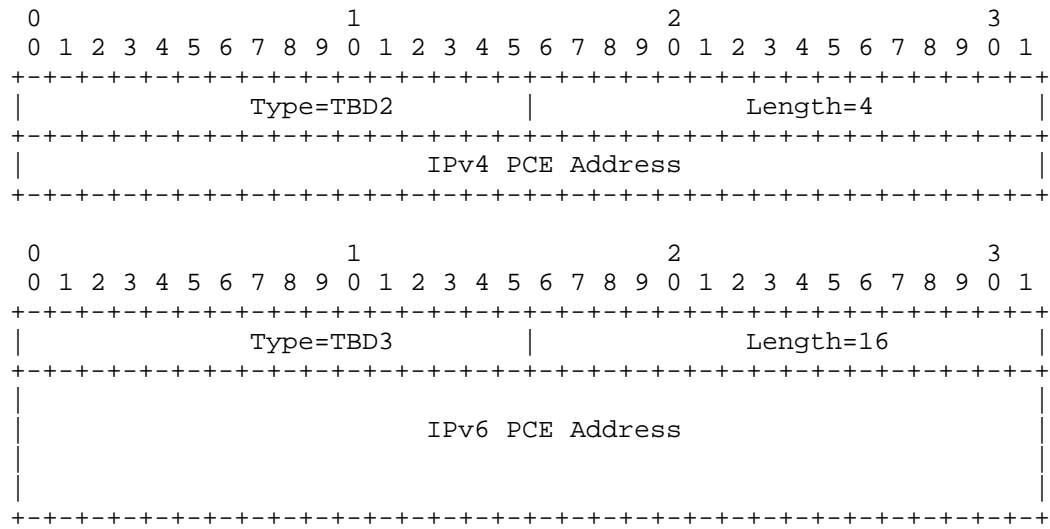


Figure 2. PCE Address TLVs

When the PCE has both an IPv4 and IPv6 address, both the TLVs MAY be included.

2.2. PCE Attribute TLVs

PCE Attribute TLVs are TLVs that may be encoded in the BGP-LS attribute [RFC7752] with a PCE NLRI. The format of each TLV is as per Section 3.1 of [RFC7752]. The format and semantics of the Value fields in some PCE Attribute TLVs correspond to the format and semantics of the Value fields in IS-IS PCED Sub-TLV, defined in [RFC5089]. Other PCE Attribute TLVs are defined in this document.

The following PCE Attribute TLVs are valid in the BGP-LS attribute with a PCE NLRI:

TLV Code Point	Description	IS-IS TLV /Sub-TLV	Reference (RFC/Section)
TBD4	Path Scope	5/2	[RFC5089]/4.2
TBD5	PCE Domain	-	-
TBD6	Neighbor PCE Domain	-	-
TBD7	PCE Capability	5/5	[RFC5089]/4.5

Table 2: PCE Attribute TLVs

The format and semantics of Path Scope and PCE capability is as per [RFC5089]. The Path Scope TLV is mandatory.

2.2.1. PCE Domain TLV

The PCE Domain TLV specifies a PCE-Domain (IGP area and/or AS) where the PCE has topology visibility and through which the PCE can compute paths.

```

      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=TBD5                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Length                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Domain Sub-TLVs (variable)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The length of this TLV is variable. The value contains one or more domain sub-TLVs as listed below -

Sub-TLV Code Point	Description	Length
512	Autonomous System	4
514	OSPF Area-ID	4
1027	IS-IS Area Identifier	Variable

Multiple sub-TLVs MAY be included, when the PCE has visibility into multiple PCE-Domains.

2.2.2. Neighbor PCE Domain TLV

The Neighbor PCE Domain TLV specifies a neighbor PCE-Domain (IGP area and/or AS) toward which a PCE can compute paths.

```

      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=TBD6                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Length                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Domain Sub-TLVs (variable)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The length of this TLV is variable. The value contains one or more domain sub-TLVs as listed above. Multiple sub-TLVs MAY be included, when the PCE can compute paths towards several neighbor PCE-Domains.

3. Operational Considerations

Existing BGP-LS operational procedures apply to the advertisement of PCE information as per [RFC7752]. This information is treated as pure application level data which has no immediate impact on forwarding states. The PCE information SHOULD be advertised only to the domains where such information is allowed to be used. This can be achieved by policy control on the ASBRs.

The PCE information is considered relatively stable and does not change frequently, thus this information will not bring significant impact on the amount of BGP updates in the network.

4. IANA Considerations

IANA needs to assign a new NLRI Type for 'PCE NLRI' from the "BGP-LS NLRI-Types" registry.

IANA needs to assign new TLV code point as per Table 1 and 2 from the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

[Editor's Note - Check if name of the registry should be changes with following instructions - Further IANA is requested to rename the registry as "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, PCE Descriptor, and Attribute TLVs".]

5. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

Existing BGP-LS security considerations as per [RFC7752] continue to apply.

6. Contributors

The following individuals gave significant contributions to this document:

Takuya Miyasaka
KDDI Corporation
ta-miyasaka@kddi.com

7. Acknowledgements

The authors would like to thank Zhenbin Li, Hannes Gredler, Jan Medved, Adrian Farrel, Julien Meuric and Jonathan Hardwick for the valuable discussion and comments.

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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<http://www.rfc-editor.org/info/rfc4760>>.
- [RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, DOI 10.17487/RFC5088, January 2008, <<http://www.rfc-editor.org/info/rfc5088>>.
- [RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, DOI 10.17487/RFC5089, January 2008, <<http://www.rfc-editor.org/info/rfc5089>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<http://www.rfc-editor.org/info/rfc7752>>.

8.2. Informative References

- [I-D.ietf-idr-te-pm-bgp]
Previdi, S., Wu, Q., Gredler, H., Ray, S.,
jeffrant@gmail.com, j., Filsfils, C., and L. Ginsberg,
"BGP-LS Advertisement of IGP Traffic Engineering
Performance Metric Extensions", draft-ietf-idr-te-pm-
bgp-06 (work in progress), June 2017.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis",
RFC 4272, DOI 10.17487/RFC4272, January 2006,
<<http://www.rfc-editor.org/info/rfc4272>>.
- [RFC4674] Le Roux, J., Ed., "Requirements for Path Computation
Element (PCE) Discovery", RFC 4674, DOI 10.17487/RFC4674,
October 2006, <<http://www.rfc-editor.org/info/rfc4674>>.
- [RFC5441] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux,
"A Backward-Recursive PCE-Based Computation (BRPC)
Procedure to Compute Shortest Constrained Inter-Domain
Traffic Engineering Label Switched Paths", RFC 5441,
DOI 10.17487/RFC5441, April 2009,
<<http://www.rfc-editor.org/info/rfc5441>>.
- [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,
<<http://www.rfc-editor.org/info/rfc6952>>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: mach.chen@huawei.com

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

Email: dhruv.ietf@gmail.com

Jeff Tantsura
Individual
US

Email: jefftant.ietf@gmail.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower, Iidabashi, Chiyoda-ku
Tokyo 102-8460
Japan

Email: ke-kumaki@kddi.com

Tomoki Murai
Furukawa Network Solution Corp.
5-1-9, Higashi-Yawata, Hiratsuka
Kanagawa 254-0016
Japan

Email: murai@fnsc.co.jp