PCE
Internet-Draft
Intended status: Standards Track
Expires: January 12, 2022

S. Peng
Q. Xiong
ZTE Corporation
F. Qin
China Mobile
M. Koldychev
Cisco Systems
S. Sivabalan
Ciena Corporation

July 11, 2021

PCE TE Constraints
draft-peng-pce-te-constraints-06

Abstract

This document proposes a set of extensions for PCEP to support the TE constraints during path computation, e.g, IGP instance, virtual network, Slice-id, specific application, color template and FA-id etc.

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 https://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 12, 2022.

Copyright Notice

Table of Contents

1. Introduction

[RFC5440] describes the Path Computation Element Protocol (PCEP) which is used between a Path Computation Element (PCE) and a Path Computation Client (PCC) (or other PCE) to enable computation of Multi-protocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP). PCEP Extensions for the Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of MPLS-TE and Generalized MPLS (GMPLS) tunnels. As depicted in [RFC4655], a PCE MUST be able to compute the path of a TE LSP by operating on the TED and considering bandwidth and other constraints applicable to the TE LSP service request. The constraint parameters are provided such as metric, bandwidth, delay, affinity, etc. However these parameters can't meet the network slicing requirements.

A PCE always perform path computation based on the network topology information collected through BGP-LS [RFC7752]. BGP-LS can get multiple link-state data from multiple IGP instance, or multiple virtual topologies from a single IGP instance. It is necessary to restrict the PCE to a small topology scope during path computation for some special purpose. BGP-LS can also get application specific TE attributes for a link, it is also necessary to restrict PCE to use

TE attributes of specific application. The PCE MUST take the identifier of slicing into consideration during path computation.

This document proposes a set of extensions for PCEP to support the TE constraints during path computation, e.g, IGP instance, virtual network, Slice-id, specific application, color template and FA-id etc.

2. Conventions used in this document

2.1. Terminology

The terminology is defined as [RFC5440] and [RFC7752].

2.2. 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.

3. PCEP Extensions for TE Constraints

As defined in [RFC5440] , the LSPA object is used to specify the LSP attributes to be taken into account by the PCE during path computation such as TE constraints. This document proposes several new TLVs for the LSPA object to carry TE constraints in Network Slicing.

3.1. Source Protocol TLV

The Source Protocol TLV is optional and is defined to carry the source protocol constraint.

In a PCReq/PCRpt message, a PCC MAY insert one or more Source Protocol TLVs to indicate the source protocol that MUST be considered by the PCE. If more than one Source Protocol TLVs are carried, the PCE may perform path computation based on the sub-topology identified by the one of the source protocols. The absence of the Source Protocol TLV MUST be interpreted by the PCE as a path computation request for which no constraints need be applied to any of the source protocols.

In a PCRep/PCInit/PCUpd message, the Source Protocol TLV MAY be carried so as to provide the source protocol information for the computed path.

The format of the Source Protocol TLV is shown as Figure 1:

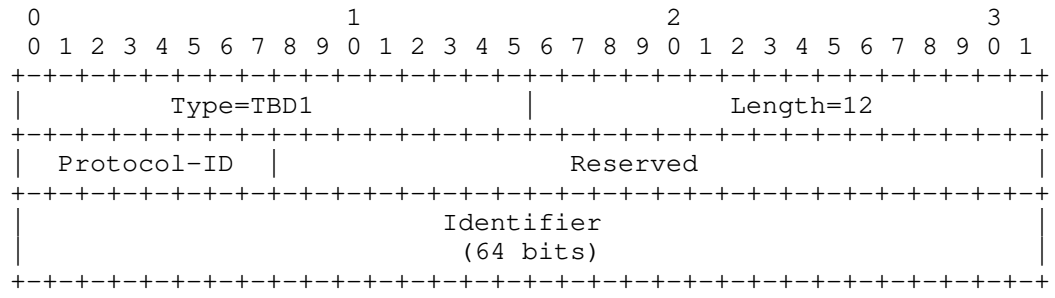

Figure 1: Source Protocol TLV

The code point for the TLV type is TBD1. The TLV length is 12 octets.

Protocol-ID (8 bits): defined in [RFC7752] section 3.2.

Reserved (24 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

Identifier (64 bits): defined in [RFC7752] section 3.2.

3.2. Multi-topology TLV

The Multi-topology TLV is optional and is defined to carry the multi-topology protocol constraint.

In a PCReq message, a PCC MAY insert one Multi-topology TLV to indicate the sub-topology of an IGP instance that MUST be considered by the PCE. The PCE will perform path computation based on the sub-topology identified by the specific Multi-Topology ID within a source protocol. The absence of the Multi-topology TLV MUST be interpreted by the PCE as a path computation request for which no constraints need be applied to any of the multi-topologies.

In a PCRep/PCInit/PCUpd message, the Multi-topology TLV MAY be carried so as to provide the Multi-topology information for the computed path.

The Multi-topology TLV MUST be carried after a Source Protocol TLV, if not it MUST be ignored.

The format of the Multi-topology TLV is shown as Figure 2:

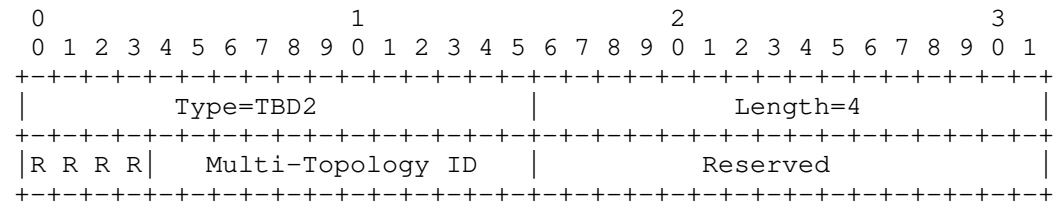

Figure 2: Multi-topology TLV

The code point for the TLV type is TBD2. The TLV length is 4 octets.

Multi-Topology ID (12 bits): Semantics of the IS-IS MT-ID are defined in Section 7.2 of [RFC5120]. Semantics of the OSPF MT-ID are defined in Section 3.7 of [RFC4915]. If the value is derived from OSPF, then the upper 9 bits MUST be set to 0. Bits R are reserved and SHOULD be set to 0 when originated and ignored on receipt.

Reserved (16 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

3.3. Slice-id TLV

PCEP message needs to carry Slice ID to let the scope of path calculation to be limited in a specific slice.

There are many control plane technologies to realize slicing. Some control plane technologies may directly maintain resources per slice granularity in the link-state database, only for the case with small slice scalability. [I-D.bestbar-teas-ns-packet] proposes a more scalable slicing scheme. The resource information in link-state database is identified by SA-ID to distinguish the logical topologies corresponding to different slice-aggregate. Within the controller, a slice-aggregate includes one or more slices mapped to it. If the number of slices is small, the resources per slice granularity can be maintained directly in the link-state database. In this case, different slice may be mapped to different slice-aggregate. If the number of slices is large, it is not recommended to maintain the slice granularity resources in the link-state database, but the aggregated SA-ID granularity.

In any case, the slice service (such as VPN service) perceives the Slice ID (not others), so it is natural for the service to include a Slice ID constraint in its TE purpose definition. For example, VPN routes may have Color attribute (refer to [I-D.ietf-idr-tunnel-encaps] and [I-D.ietf-spring-segment-routing-policy]). Color represents a

specific TE purpose, which can contain a Slice ID. Thus it is natural carry Slice ID in PCEP message.

When the controller receives the path computation request with a Slice ID constraint, it can use the resources identified by specific Slice in TED, or firstly look up the Slice ID to SA-ID mapping entry and then use the resources of specific SA-ID in TED, to calculate the path.

In a PCReq message, a PCC MAY insert one Slice-id TLV to indicate the slice based virtual network that MUST be considered by the PCE. The PCE will perform path computation based on the intra-domain or inter-domain sub-topology identified by the specific Slice-id, which is independent of routing protocols such as IGP/BGP. The absence of the Slice-id TLV MUST be interpreted by the PCE as a path computation request for which no constraints need be applied to any of slice, i.e, a default Slice-id (0) will be applied.

In a PCRep/PCInit/PCUpd message, the Slice-id TLV MAY be carried so as to provide the network slicing information for the computed path. The headend may put the Slice-id to an encapsulated data packet.

The format of the Slice-id TLV is shown as Figure 3:

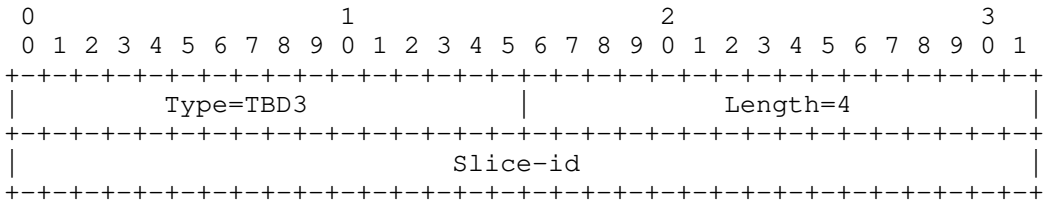

Figure 3: Slice-id TLV

The code point for the TLV type is TBD3. The TLV length is 4 octets.

Slice-id (32 bits): indicate the Slice-id information. The Slice-id is also termed as AII defined in [I-D.peng-lsr-network-slicing] to represent an IETF Network Slice that is defined in [I-D.ietf-teas-ietf-network-slice-definition].

3.4. Application Specific TLV

The Application Specific TLV is optional and is defined to carry the application specific constraints.

In a PCReq message, a PCC MAY insert one Application Specific TLV to indicate the application that MUST be considered by the PCE. The PCE will perform path computation using the specific application attributes. The absence of the Application Specific TLV MUST be interpreted by the PCE as a path computation request for which no constraints need be applied to any of the Application Specific attributes.

In a PCRep/PCInit/PCUpd message, the Application Specific TLV MAY be inserted so as to provide the Application Specific information for the computed path.

The format of the Application Specific TLV is shown as Figure 4:

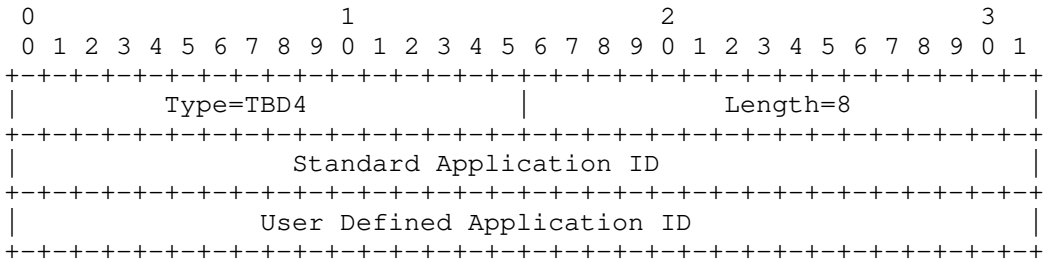

Figure 4: Application Specific TLV

The code point for the TLV type is TBD4. The TLV length is 8 octets.

Standard Application ID: Represents a bit-position value for a single STANDARD application that is defined in the IANA "IGP Parameters" registries under the "Link Attribute Applications" registry [RFC8919].

User Defined Application ID: Represents a single user defined application which is a specific implementation.

3.5. Color TLV

The Color TLV is optional and is defined to carry the color constraints.

In a PCReq message, a PCC MAY insert one Color TLV to indicate the traffic engineering purpose that is recognized by both PCE and PCC with no conflict meaning. The PCE will perform path computation based on the color template. The same color template may be also defined at PCC and the existing constraints (i.e, metric, bandwidth, delay, etc) carried in the message MUST be ignored. The absence of

the Color TLV MUST be interpreted by the PCE as a path computation request for which traditional constraints that are contained in message need be applied.

In a PCRep/PCInit/PCUpd message, the Color TLV MAY be inserted so as to provide the TE purpose information for the computed path, the PCC recognize the color value that match a local color-template. For example, the COLOR TLV can be used to identify the Color of each Candidate Path in the Composite Candidate Path as decribed in [I-D.ietf-pce-multipath]

The format of the Color TLV is shown as Figure 5:

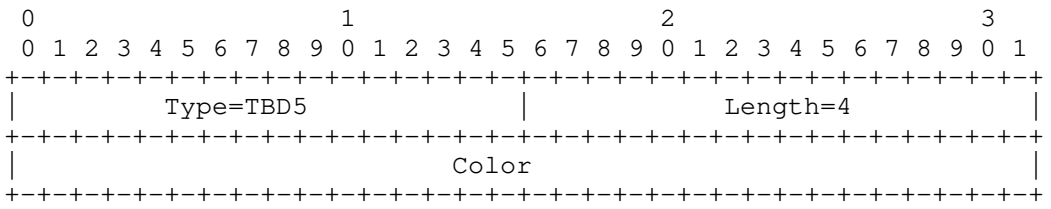

Figure 5: Color TLV

The code point for the TLV type is TBD5. The TLV length is 4 octets.

Color (32 bits): indicate a TE template, 0 is invalid value. It is consistent with the Color Extended Community defined in [I-D.ietf-idr-tunnel-encaps], and color of SR policy defined in [I-D.ietf-spring-segment-routing-policy].

Note that Color TLV defined in this document is used to represent a TE template, it can be suitable for any TE instance such as RSVP-TE, SR-TE, SR-policy. [I-D.ietf-pce-segment-routing-policy-cp] has proposed the SR policy KEY (that also includes a color information) as an association group KEY to associate many candidate paths, however it is only for association purpose but not constraint purpose for path computation.

A color template can be defined to contain existing constraints such as metric, bandwidth, delay, affinity parameters, and the sub-topology constraints above defined in this document.

3.6. FA-id TLV

FA-id defined in [I-D.ietf-lsr-flex-algo] is a short mapping of SR policy color to optimize segment stack depth for the IGP area partial of the entire SR policy. The overlay service that want to be carried over a particular SR-FA path must firstly let the SR policy supplier know that requirement. There are two possible ways to map a color to an FA-id. One is explicit mapping configuration within color template, the other is dynamicly replacing a long segment list to short FA segment by headend or controller once the constraints contained in the color-template equal to that contained in FAD.

In addition to the above mapping behavior, it is also possible to merge the constraints contained in the color-template and constraints contained in FAD. The merging behavior can be used to compute SR-TE path within a Flex-algo plane.

In a PCReq message, a PCC MAY insert one FA-id TLV to indicate the above explicit FA-id mapping or merging. For mapping case, the PCE will perform path computation based on the FA-id mapping. In detailed, The PCE will check if there are connectivity within the corresponding Flex-algo plane to the destination. If yes, the path computation result will be represented as segment list with a single prefix-SID@FA for intra-domain case, or several prefix-SID@FA for inter-domain case.

For merging case, the PCE will perform path computation based on the total constraints combinded with the ones contained in FAD identified by FA-id and other ones contained in PCReq message. The later constraints can get from color template or directly represent by a color. In this case the computed path will be limited in the specific Flex-algo plane determined by link resource Including/ Excluding rules of FAD, and at the same time the path will also meet other constraints for the TE purpose within the Flex-algo plane. The PCE can optimize the strictly path to a loosely path when a part of the strictly path is consistent with the algorithm based path, i.e, some consecutive adjacency SIDs can be replaced with a single algorithm based Prefix-SID.

In a PCRep/PCInit/PCUpd message, the FA-id TLV MAY be inserted so as to provide the FA plane information for the computed path.

In general, the FA-id TLV is only meaningful for the domain (ingress domain) that headend node belongs to. For inter-domain case, operator SHOULD ensure the FA-id configuration of different domain are same for an E2E slice, when he want to explicitly indicate FA-id in PCEP message, otherwise the PCE has to choose different FA-id for

other domain as long as the contents of FAD is consistent with the one of ingress domain.

The format of the FA-id TLV is shown as Figure 6:

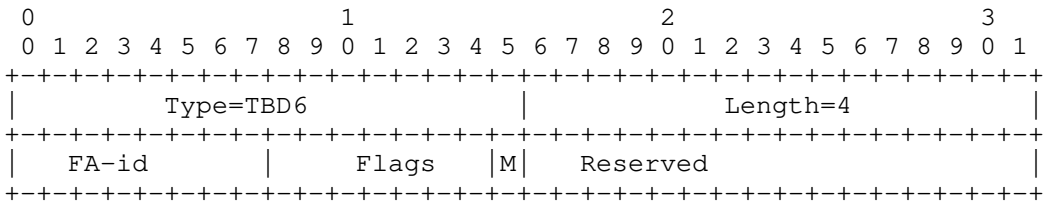

Figure 6: FA-id TLV

The code point for the TLV type is TBD6. The TLV length is 4 octets.

FA-id (8 bits): indicate an explicit FA-id mapping information.

Flags (8 bits): Currently only one flag, Flag-M, is defined.

Flag-M: Indicate mapping behavior when unset, and merging behavior when set.

4. Security Considerations

TBA

5. Acknowledgements

TBA

6. IANA Considerations

IANA is requested to make allocations from the registry, as follows:

Type	TLV	Reference
TBD1	Source Protocol TLV	[this document]
TBD2	Multi-topology TLV	[this document]
TBD3	Slice-id TLV	[this document]
TBD4	Application Specific TLV	[this document]
TBD5	Color TLV	[this document]
TBD6	FA-id TLV	[this document]

Table 1

7. Normative References

[I-D.bestbar-teas-ns-packet]
Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., and L. M. Contreras, "Realizing Network Slices in IP/MPLS Networks", draft-bestbar-teas-ns-packet-02 (work in progress), February 2021.

[I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G. V. D., Sangli, S. R., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", draft-ietf-idr-tunnel-encaps-22 (work in progress), January 2021.

[I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", draft-ietf-lsr-flex-algo-15 (work in progress), April 2021.

[I-D.ietf-pce-multipath]
Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., and S. Peng, "PCEP Extensions for Signaling Multipath Information", draft-ietf-pce-multipath-00 (work in progress), May 2021.

[I-D.ietf-pce-segment-routing-policy-cp]
Koldychev, M., Sivabalan, S., Barth, C., Peng, S., and H. Bidgoli, "PCEP extension to support Segment Routing Policy Candidate Paths", draft-ietf-pce-segment-routing-policy-cp-04 (work in progress), March 2021.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-11 (work in progress), April 2021.

[I-D.ietf-teas-ietf-network-slice-definition]
Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Definition of IETF Network Slices", draft-ietf-teas-ietf-network-slice-definition-01 (work in progress), February 2021.

[I-D.peng-lsr-network-slicing]
Peng, S., Chen, R., and G. Mirsky, "Packet Network Slicing using Segment Routing", draft-peng-lsr-network-slicing-00 (work in progress), February 2019.

[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>.

[RFC4655] Farrel, A., Vasseur, J., 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>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[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>.

[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, <https://www.rfc-editor.org/info/rfc7752>.

[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>.

[RFC8919] Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS Application-Specific Link Attributes", RFC 8919, DOI 10.17487/RFC8919, October 2020, <https://www.rfc-editor.org/info/rfc8919>.

Authors' Addresses

Shaofu Peng
ZTE Corporation
No.50 Software Avenue
Nanjing, Jiangsu 210012
China

Email: peng.shaofu@zte.com.cn

Quan Xiong
ZTE Corporation
No.6 Huashi Park Rd
Wuhan, Hubei 430223
China

Email: xiong.quan@zte.com.cn

Fengwei Qin
China Mobile
Beijing
China

Email: qinfengwei@chinamobile.com

Mike Koldychev
Cisco Systems
Canada

Email: mkoldych@cisco.com

Siva Sivabalan
Ciena Corporation
Canada

Email: ssivabal@ciena.com