

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2020

U. Chunduri
Y. Qu
Futurewei
R. White
Juniper Networks
J. Tantsura
Apstra Inc.
L. Contreras
Telefonica
March 8, 2020

Preferred Path Routing (PPR) in OSPF
draft-chunduri-lsr-ospf-preferred-path-routing-04

Abstract

This document specifies a Preferred Path Routing (PPR), a routing protocol mechanism to simplify the path description of data plane traffic in Segment Routing (SR) deployments with OSPFv2 and OSPFv3 protocols. PPR aims to mitigate the MTU and data plane processing issues that may result from SR packet overheads; and also supports further extensions along the paths. Preferred path routing is achieved through the addition of path descriptions to the OSPF advertised prefixes, and mapping those to a PPR data-plane identifier.

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](#) [RFC2119], [RFC8174](#) [RFC8174] when, and only when they appear in all capitals, as shown here".

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

Internet-Draft Preferred Path Routing (PPR) in OSPF

March 2020

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 September 9, 2020.

Copyright Notice

Copyright (c) 2020 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 (<https://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	3
1.1.	Acronyms	3
2.	OSPFv2 PPR TLV	4
2.1.	PPR-Flags	6
2.2.	PPR-Prefix Sub-TLV	6
2.3.	PPR-ID Sub-TLV	7
2.4.	PPR-PDE Sub-TLV	9
2.5.	PPR-Attributes Sub-TLV	11
3.	OSPFv3 PPR TLV	12
3.1.	OSPFv3 PPR-Prefix Sub-TLV	13
3.2.	OSPFv3 PPR-ID Sub-TLVs	14
3.3.	OSPFv3 PPR-PDE Sub-TLV	16
3.4.	OSPFv3 PPR-Attributes Sub-TLV	19
4.	Other Considerations	19
5.	Acknowledgements	19
6.	IANA Considerations	19
7.	Security Considerations	20
8.	References	20
8.1.	Normative References	20
8.2.	Informative References	20
	Authors' Addresses	22

1. Introduction

In a network implementing Segment Routing (SR), packets are steered through the network using Segment Identifiers (SIDs) carried in the packet header. Each SID uniquely identifies a segment as defined in [\[I-D.ietf-spring-segment-routing\]](#). SR capabilities are defined for MPLS and IPv6 data planes called SR-MPLS and SRv6 respectively.

In SR-MPLS, a segment is encoded as a label and an ordered list of segments is encoded as a stack of labels on the data packet. In SRv6, a segment is encoded as an IPv6 address, with in a new type of IPv6 hop-by-hop routing header/extension header (EH) called SRH [\[I-D.ietf-6man-segment-routing-header\]](#), where an ordered list of IPv6 addresses/segments is encoded in SRH.

Preferred path routing can be described as a) enabling route computation based on the specific path described along with the prefix as opposed to shortest path towards the prefix and b) forwarding based on the abstracted path identifier as opposed to the individual segments on the packet. This is also further described in Section 2 of [\[I-D.chunduri-lsr-isis-preferred-path-routing\]](#).

Any prefix advertised with a path description from any node in the network is called PPR. A PPR could be an SR path, an explicitly provisioned Fast Re-Route (FRR) path or a service chained path. A PPR can be signaled by any node, which receives the SR path computed by a central controller, or by statically configuring the same on a node in the network.

The issues caused by the large SID depth, and existing methods for mitigation are introduced in [\[I-D.chunduri-lsr-isis-preferred-path-routing\]](#) in [Appendix A.1](#) and [A.2](#). To mitigate these issues and also to facilitate forwarding plane extensibility, this draft proposes a new OSPFv2 PPR TLV ([Section 2](#)), OSPFv3 PPR TLV ([Section 3](#)) to use the path with a corresponding data plane identifier.

1.1. Acronyms

EL	- Entropy Label
ELI	- Entropy Label Indicator
MPLS	- Multi Protocol Label Switching
MSD	- Maximum SID Depth
MTU	- Maximum Transferrable Unit

PPR	- Preferred Path Route
SID	- Segment Identifier
SPF	- Shortest Path First
SR	- Segment Routing
SRH	- Segment Routing Header
SR-MPLS	- Segment Routing with MPLS data plane
SRv6	- Segment Routing with Ipv6 data plane with SRH
SRH	- IPv6 Segment Routing Header
TE	- Traffic Engineering

2. OSPFv2 PPR TLV

Extended Prefix Opaque LSAs defined in [[RFC7684](#)] are used for advertisements of PPRs. This section describes the encoding of PPR TLV. This TLV can be seen as having 4 logical section viz., encoding of the OSPFv2 Prefix, encoding of PPR-ID, encoding of path description with an ordered PDE Sub-TLVs and a set of optional PPR attribute Sub-TLVs, which can be used to describe one or more parameters of the path. Multiple OSPF PPR TLVs MAY be advertised in each OSPF Extended Prefix Opaque LSA, but all TLVs included in a single OSPF Extended Prefix Opaque LSA MUST have the same flooding

scope.

The PPR TLV has Type TBD (suggested value xxx), and has the following format:




Figure 1: OSPFV2 PPR TLV Format

- o Type - TBD (IANA)from OSPF Extended Prefix Opaque LSA registry.
- o Length - Total length of the value field in bytes (variable).

- o PPR-Flags - 2 Octet flags for this TLV are described below.
- o AF - Address family for the prefix. Currently, the only supported value is 0 for IPv4 unicast. The inclusion of address family in this TLV allows for future extension.
- o Reserved - 1 Octet reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.
- o PPR-Prefix - This is a variable size Sub-TLV, which represents the prefix for which path description is being attached to. This is defined in [Section 2.2](#).
- o PPR-ID - This is a variable size Sub-TLV, which represents the data plane or forwarding identifier of the PPR. This is defined in [Section 2.3](#).
- o PPR-PDEs - Variable number of ordered PDE Sub-TLVs which represents the path. This is defined in [Section 2.4](#).
- o PPR-Attributes - Variable number of PPR-Attribute Sub-TLVs which represent the path attributes. These are defined in [Section 2.5](#).

[2.1](#). PPR-Flags

Flags: 2 octet field of PPR TLV has following flags defined:

PPR TLV Flags Format

0	1	2	3	4	5	6	7	15
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+								
IA A		Reserved						
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+								

w=Where:

IA-Flag: Inter-Area flag. If set, advertisement is of inter-area type. An Area Boarder Router (ABR) that is advertising the OSPF PPR TLV between areas MUST set this bit.




Figure 2: PPR-Prefix Sub-TLV Format

- o Type - 1 (suggested value, IANA TBD) from OSPFv2 PPR TLV [Section 2](#) Sub-TLV registry.
- o Length - Total length of the value field in bytes (variable).
- o MT-ID - Multi-Topology ID (as defined in [[RFC4915](#)]).
- o Prefix Len - contains the length of the OSPF prefix being encoded in bytes.
- o Mask Length - The length of the prefix in bits. Only the most significant octets of the Prefix are encoded.
- o OSPFv2 Prefix - represents the OSPFv2 prefix at the tail-end of the advertised PPR. For the address family IPv4 unicast, the prefix itself is encoded as a 32-bit value. The default route is represented by a prefix of length 0.
- o PPR-Prefix Sub-TLVs have 2 octet type, 2 octet length and value field is defined per type.

[2.3.](#) PPR-ID Sub-TLV

This represents the actual data plane identifier in the packet and could be of any data plane as defined in PPR-ID-type field. Both OSPF Prefix and PPR-ID MUST belong to a same node in the network.

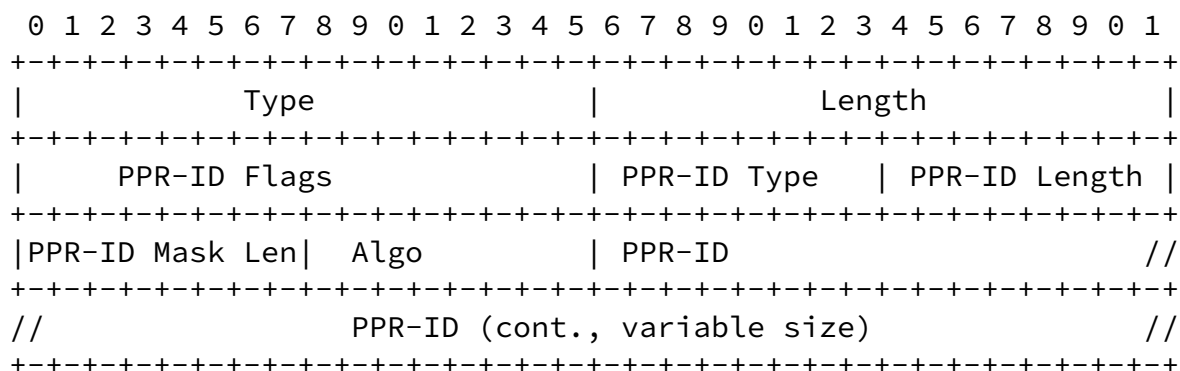


Figure 3: PPR-ID Sub-TLV Format


- o Type - 2 (suggested value, IANA TBD) from OSPFv2 PPR TLV [Section 2](#) Sub-TLV registry.
- o Length - Total length of the value field in bytes (variable).
- o PPR-ID Type - Data plane type of PPR-ID. This is a new registry (TBD IANA) for this Sub-TLV and the defined types are as follows:

Type: 1 SR-MPLS SID/Label

Type: 2 Native IPv4 Address/Prefix

- o PPR-ID Flags - 2 Octet field for PPR-ID flags:

PPR-ID Flags Format



Reserved - Reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

- o PPR-ID Type - Data plane type of PPR-ID. Values are defined in [\[I-D.chunduri-lsr-isis-preferred-path-routing\]](#). Only Type 1 and Type 2 are applicable here.
- o PPR-ID Length - Length of the PPR-ID field in octets and this depends on the PPR-ID type. See PPR-ID below for the length of this field and other considerations.

- ## 2.4. PPR-PDE Sub-TLV

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	2	3	4	5	6	7	8	9	0	1
Type										Length																															
PPR-PDE Type										PDE-ID Type										PDE-ID Len										Reserved											
PPR-PDE Flags										PDE-ID Value										//																					
//										PDE-ID Value (Contd., Variable size)										//																					
PPR-PDE										Sub-TLVs (variable)																															

- o Type - 3 (TBD IANA, suggested value) from OSPFv2 PPR TLV [Section 2](#) Sub-TLV registry.
- o Length - Total length of the value field in bytes (variable).

- o PPR-PDE Type - This is a new registry (TBD IANA) for this Sub-TLV and the defined types are as follows:

Type: 1 Topological

Type: 2 Non-Topological

- o PDE-ID Type - 1 Octet PDE-forwarding IDentifier Type. This is a new registry (TBD IANA) for this Sub-TLV and the defined types and corresponding PDE-ID Len, PDE-ID Value are as follows:

Type 0: This value MUST be set only when PPR-PDE Type is Non-Topological. PDE-ID Len specified in bytes and encoded in NBO in PDE-ID Value field which can represent a service/function. This information is provisioned on the immediate topological PDE preceding to this PDE based on the 'E' bit.

Type 1: SID/Label Sub-TLV as defined in [\[I-D.ietf-ospf-segment-routing-extensions\]](#). PDE-ID Len and PDE-ID Value fields are per [Section 2.1](#) of the referenced document.

Type 2: SR-MPLS Prefix SID. PDE-ID Len and PDE-ID Value are same as Type 1.

Type 3: SR-MPLS Adjacency SID. PDE-ID Len and PDE-ID Value are same as Type 1.

Type 4: IPv4 Node Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#).

Type 5: IPv4 P2P interface Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#).

Type 6: IPv4 LAN interface Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#). This type MUST have OSPF Neighbor ID sub-TLV in the PDE.

- o PDE-ID Len - 1 Octet. Length of PDE-ID field.
- o Reserved - 1 Octet reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.
- o PPR-PDE Flags - 2 Octet flags for this TLV are described below:

PPR-PDE Flags Format

```

0 1 2 3 4 5 6 7... 15
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|D|E|   Reserved   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


L: Loose Bit: This bit indicates the type of next "Topological PDE-ID" in the path description. If set, the next PDE is Loose. If this flag is unset, the next Topological PDE is Strict Type.

D: Destination Bit: By default this bit MUST be unset. This bit MUST be set only for PPR-PDE Type is Topological and this PDE represents the PDE-ID corresponding to the PPR-Prefix [Section 2.2](#).

E: Egress Bit. By default this bit MUST be unset. This bit MUST be set only for PPR-PDE Type is 2 i.e., Non-Topological and the service needs to be applied on the egress side of the topological PDE preceding this PDE.

Reserved: Reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

- o PPR-PDE Sub-TLVs have 2 octet type, 2 octet length and value field is defined per type.
- o PPR-PDE Sub-TLV: Type 4 (IANA TBD), Length Total length of value field in bytes, Value: The Router ID of the neighbor for which the LAN interface is advertised. This Sub-TLV MUST NOT be present, if the PPR-PDE Type is not equal to 1 i.e., Topological PDE and PDE-



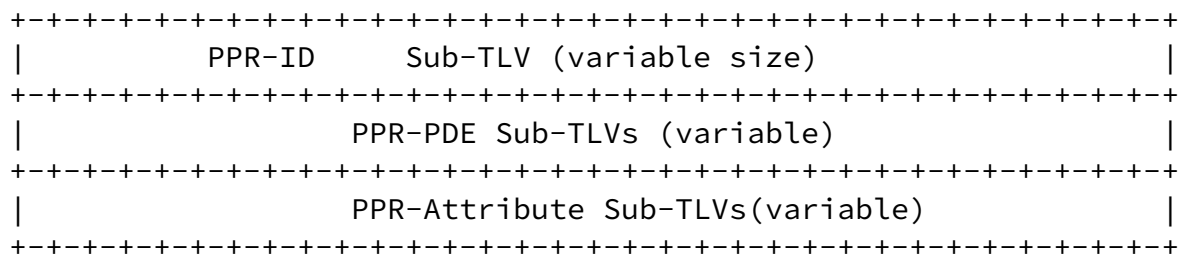


Figure 5: OSPFv3 PPR TLV Format

- o Type - TBD (IANA)from OSPF Extended Prefix Opaque LSA registry.
- o Length - Total length of the value field in bytes (variable).
- o PPR-Flags - 2 Octet flags for this TLV are described below.
- o AF: Address family for the prefix.


AF: 0 - IPv4 unicast

AF: 1 - IPv6 unicast

- o Reserved - 1 Octet reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

Flags: 2 octet field. The following flags are defined:

OSPFv3 PPR TLV Flags Format



IA-Flag: Inter-Area flag. If set, advertisement is of inter-area type. An ABR that is advertising the OSPF PPR TLV between areas MUST set this bit.

[\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#)

A: The originator of the PPR TLV MUST set the A bit in order to signal that the prefixes and PPR-IDs advertised in the PPR TLV are directly connected to the originators. If this bit is not set, this allows any other node in the network advertise this TLV on behalf of the originating node of the "OSPF Prefix". If PPR TLV is propagated to other areas the A-flag MUST be cleared. In case if the originating node of the prefix has to be disambiguated for any reason including, if it is a Multi Homed Prefix (MHP) or propagated to a different OSPF area, then PPR-Attribute Sub-TLV Source Router ID SHOULD be included.

Reserved - reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

PPR path description for each OSPF area is computed and given to one of the nodes in that area for dissemination. Similarly path information when crossing the area boundaries MUST be relevant to the destination area. If there is no path information available for the destination area, PPR TLV MUST NOT be leaked regardless of the IA bit status.

[3.1.](#) OSPFv3 PPR-Prefix Sub-TLV

The structure of OSPFv3 PPR-Prefix, for which path description is attached to is as follows:





Figure 6: OSPFv3 PPR-Prefix Sub-TLV Format

- o Type - 1 (suggested value, IANA TBD) from OSPFv3 PPR TLV [Section 3](#) Sub-TLV registry.
- o Length - Total length of the value field in bytes (variable).
- o Prefix Len - contains the length of the prefix in bits. Only the most significant octets of the Prefix are encoded.
- o Mask Length - The length of the prefix in bits. Only the most significant octets of the Prefix are encoded.
- o OSPFv3 Prefix - represents the OSPFv3 prefix at the tail-end of the advertised PPR. For the address family IPv4 unicast, the prefix itself is encoded as a 32-bit value. The default route is represented by a prefix of length 0. For the address family (AF in OSPFv3 PPR TLV) in IPv6 unicast, the prefix, encoded as an even multiple of 32-bit words, padded with zeroed bits as necessary. This encoding consumes $((\text{PrefixLength} + 31) / 32)$ 32-bit words.
- o PPR-Prefix Sub-TLVs have 2 octet type, 2 octet length and value field is defined per type.

3.2. OSPFv3 PPR-ID Sub-TLVs

This represents the actual data plane identifier in the packet and could be of any data plane as defined in PPR-ID-type field. Both OSPF Prefix and PPR-ID MUST belong to a same node in the network.




```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      PPR-ID Flags          | PPR-ID Type   | PPR-ID Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|PPR-ID Mask Len|  Algo          | PPR-ID              //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                               PPR-ID (cont, variable size)          //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 7: OSPFv3 PPR-ID Sub-TLV Format

- o Type - 2 (suggested value, IANA TBD) from OSPFv3 PPR TLV [Section 3](#) Sub-TLV registry.
- o Length - Total length of the value field in bytes (variable).
- o PPR-ID Type - Data plane type of PPR-ID. This is a new registry (TBD IANA) for this Sub-TLV and the defined types are as follows:

Type: 1 SR-MPLS SID/Label

Type: 2 Native IPv4 Address/Prefix

Type: 3 Native IPv6 Address/Prefix

Type: 4 IPv6 SID in SRv6 with SRH

- o PPR-ID Flags - 2 Octet field for PPR-ID flags:

PPR-ID Flags Format

```

      0 1 2 3 4 5 6 7          15
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|A| Reserved                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Reserved - Reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

- o PPR-ID Length - Length of the PPR-ID field in octets and this depends on the PPR-ID type. See PPR-ID below for the length of this field and other considerations.

- o PPR-ID Mask Len - It is applicable for only for PPR-ID Type 2, 3 and 4. For Type 1 this value MUST be set to zero. It contains the length of the PPR-ID Prefix in bits. Only the most significant octets of the Prefix are encoded. This is needed, if PPR-ID followed is an IPv4/IPv6 Prefix instead of 4/16 octet Address respectively.
- o Algo - 1 octet value represents the SPF algorithm. Algorithm registry is as defined in [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).
- o PPR-ID - This is the Preferred Path forwarding identifier that would be on the data packet. The value of this field is variable and it depends on the PPR-ID Type - for Type 1, this is encoded as SR-MPLS SID/Label. For Type 2 this is encoded as 4 byte IPv4 address. For Type 3 this is encoded as 16 byte IPv6 address. For Type 2 and Type 3 encoding is similar to OSPF Prefix as specified in [Section 2.2](#). For Type 4, this is encoded as 16 byte IPv6 SID.

3.3. OSPFv3 PPR-PDE Sub-TLV

This is a new Sub-TLV type in PPR TLV [Section 3](#) and is called as PPR Path Description Element (PDE). PPR-PDEs are used to describe the path in the form of set of contiguous and ordered Sub-TLVs, where first Sub-TLV represents (the top of the stack in MPLS data plane or) first node/segment of the path. These set of ordered Sub-TLVs can have both topological SIDs and non-topological SIDs (e.g., service segments).

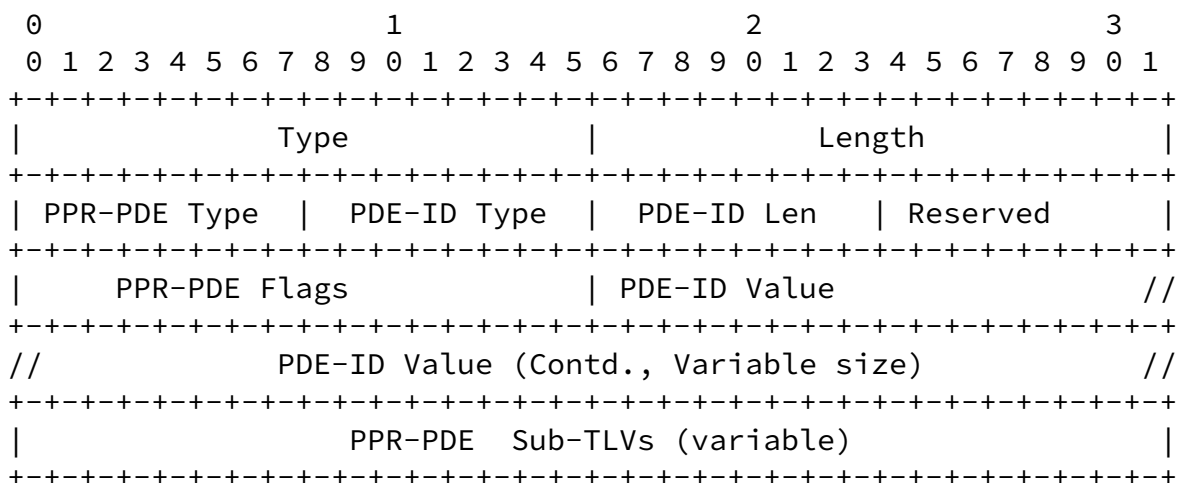


Figure 8: OSPFv3 PPR-PDE Sub-TLV Format

- o Type - 3 (suggested value, IANA TBD) from OSPFv3 PPR TLV [Section 3](#) Sub-TLV registry.

- o Length - Total length of the value field in bytes (variable).
- o PPR-PDE Type - This is a new registry (TBD IANA) for this Sub-TLV and the defined types are as follows:
 - Type: 1 Topological
 - Type: 2 Non-Topological
- o PDE-ID Type - 1 Octet PDE-forwarding IDentifier Type. This is a new registry (TBD IANA) for this Sub-TLV and the defined types and corresponding PDE-ID Len, PDE-ID Value are as follows:
 - Type 0: This value MUST be set only when PPR-PDE Type is Non-Topological. PDE-ID Len specified in bytes and encoded in NBO in PDE-ID Value field which can represent a service/function. This information is provisioned on the immediate topological PDE preceding to this PDE based on the 'E' bit.
 - Type 1: SID/Label Sub-TLV as defined in [\[I-D.ietf-ospf-segment-routing-extensions\]](#). PDE-ID Len and PDE-ID Value fields are per [Section 2.1](#) of the referenced document.
 - Type 2: SR-MPLS Prefix SID. PDE-ID Len and PDE-ID Value are same as Type 1.
 - Type 3: SR-MPLS Adjacency SID. PDE-ID Len and PDE-ID Value are same as Type 1.
 - Type 4: IPv4 Node Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#).
 - Type 5: IPv4 P2P interface Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#).
 - Type 6: IPv4 LAN interface Address. PDE-ID Len is 4 bytes and PDE-ID Value is 4 bytes IPv4 address encoded similar to IPv4 Prefix described in [Section 2.2](#). This type MUST have OSPF

Neighbor ID Sub-TLV in the PDE.

Type 7: IPv6 Node Address. PDE-ID Len is 16 bytes and PDE-ID Value is 16 bytes IPv6 address encoded similar to IPv6 Prefix described in [Section 2.2](#).

Type 8: IPv6 P2P interface Address. PDE-ID Len is 16 bytes and PDE-ID Value is 16 bytes IPv6 address encoded similar to IPv6 Prefix described in [Section 2.2](#).

Type 9: IPv6 LAN interface Address. PDE-ID Len is 16 bytes and PDE-ID Value is 16 bytes IPv6 address encoded similar to IPv6 Prefix described in [Section 2.2](#). This type MUST have OSPF Neighbor ID Sub-TLV in the PDE.

Type 10: SRv6 Node SID as defined in [\[I-D.li-ospf-ospfv3-srv6-extensions\]](#). PDE-ID Len and PDE-ID Value are as defined in SRv6 SID.

Type 11: SRv6 Adjacency-SID. PDE-ID Len and PDE-ID Value are as defined in Type 6.

- o PDE-ID Len - 1 Octet. Length of PDE-ID field.
- o Reserved - 1 Octet reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.
- o PPR-PDE Flags - 2 Octet flags for this TLV are described below:

PPR-PDE Flags Format

```

0 1 2 3 4 5 6 7... 15
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|L|D|E|      Reserved      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

L: Loose Bit. This bit indicates the type of next "Topological PDE-ID" in the path description and overrides the L bit in

[Section 3.2](#). If set, the next PDE is Loose. If this flag is unset, the next Topological PDE is Strict Type.

D: Destination Bit. By default this bit MUST be unset. This bit MUST be set only for PPR-PDE Type is Topological and this PDE represents the PDE-ID corresponding to the PPR-Prefix [Section 3.1](#).

E: Egress Bit. By default this bit MUST be unset. This bit MUST be set only for PPR-PDE Type is 2 i.e., Non-Topological and the service needs to be applied on the egress side of the topological PDE preceding this PDE.

Reserved - Reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

- o PPR-PDE Sub-TLVs have 2 octet type, 2 octet length and value field is defined per type.
- o PPR-PDE Sub-TLV: Type 4 (IANA TBD), Length Total length of value field in bytes, Value: The Router ID of the neighbor for which the LAN interface is advertised. This Sub-TLV MUST NOT be present, if the PPR-PDE Type is not equal to 1 i.e., Topological PDE and PDE-ID Type 6/9.

[3.4](#). OSPFv3 PPR-Attributes Sub-TLV

PPR-Attribute Sub-TLVs describe the attributes of the path. The following Sub-TLVs draw from a new registry for Sub-TLV numbers; this registry is to be created by IANA, and administered using the first come first serve process:

- o Type 1 (suggested value, IANA TBD): PPR-Metric Sub-TLV. Length 4 bytes, and Value is metric of this path represented through the PPR-ID. Different nodes can advertise the same PPR-ID for the same Prefix with a different set of PPR-PDE Sub-TLVs and the receiving node MUST consider the lowest metric value.

[4](#). Other Considerations

Please refer to [I-D.chunduri-isis-preferred-path-routing] [section 4](#), 5, 6 and 7.

[5.](#) Acknowledgements

Thanks to Richard Li, Alex Clemm, Padma Pillay-Esnault, Toerless Eckert, Kiran Makhijani and Lin Han for initial discussions on this topic. Thanks to Kevin Smith and Stephen Johnson for various deployment scenarios applicability from ETSI WGs perspective. Authors also acknowledge Alexander Vainshtein for detailed discussions and suggestions on this topic.

Earlier versions of [draft-ietf-ospf-segment-routing-extensions](#) have a mechanism to advertise EROs through Binding SID.

[6.](#) IANA Considerations

This document requests the following new TLV in IANA OSPFv2 and OSPFv3 TLV code-point registry as specified in [Section 2](#) [Section 3](#) respectively .

TLV #	Name
-----	-----
TBD	PPR TLV

This document also requests IANA to create new registries for PPR TLV Flags field, PPR Flags, and PPR Sub-TLVs in PPR TLV as described in [Section 2](#) and [Section 3](#).

[7.](#) Security Considerations

Existing security extensions as described in [[RFC2328](#)] and [[RFC7684](#)] apply to the extensions specified in this document. While OSPF is under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the OSPF routing domain. In these deployments, stronger authentication mechanisms such as those specified in [[RFC7474](#)] SHOULD be used.

Advertisement of the additional information defined in this document

introduces no new security concerns in OSPF protocol. However as this extension is related to SR-MPLS and SRH data planes as defined in [[I-D.ietf-spring-segment-routing](#)], those particular data plane security considerations does apply here.

[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>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [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>>.

[8.2.](#) Informative References

- [I-D.chunduri-lsr-isis-preferred-path-routing]
Chunduri, U., Li, R., White, R., Tantsura, J., Contreras, L., and Y. Qu, "Preferred Path Routing (PPR) in IS-IS", [draft-chunduri-lsr-isis-preferred-path-routing-04](#) (work in progress), July 2019.

Chunduri, et al.	Expires September 9, 2020	[Page 20]
------------------	---------------------------	-----------

Internet-Draft	Preferred Path Routing (PPR) in OSPF	March 2020
----------------	--------------------------------------	------------

- [I-D.ietf-6man-segment-routing-header]
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-26](#) (work in progress), October 2019.

- [I-D.ietf-ospf-ospfv3-lsa-extend]
Lindem, A., Roy, A., Goethals, D., Vallem, V., and F. Baker, "OSPFv3 LSA Extendibility", [draft-ietf-ospf-ospfv3-lsa-extend-23](#) (work in progress), January 2018.

- [I-D.ietf-ospf-ospfv3-segment-routing-extensions]

Psenak, P. and S. Previdi, "OSPFv3 Extensions for Segment Routing", [draft-ietf-ospf-ospfv3-segment-routing-extensions-23](#) (work in progress), January 2019.

[I-D.ietf-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-27](#) (work in progress), December 2018.

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-15](#) (work in progress), January 2018.

[I-D.li-ospf-ospfv3-srv6-extensions]

Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", [draft-li-ospf-ospfv3-srv6-extensions-07](#) (work in progress), November 2019.

[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>>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", [RFC 7474](#), DOI 10.17487/RFC7474, April 2015, <<https://www.rfc-editor.org/info/rfc7474>>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.

Authors' Addresses

Uma Chunduri
Futurewei
2330 Central Expressway

Santa Clara, CA 95050
USA

Email: umac.ietf@gmail.com

Yingzhen Qu
Futurewei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: yingzhen.qu@futurewei.com

Russ White
Juniper Networks
Oak Island, NC 28465
USA

Email: russ@riw.us

Jeff Tantsura
Apstra Inc.
333 Middlefield Road
Menlo Park, CA 94025
USA

Email: jefftant.ietf@gmail.com

Luis M. Contreras
Telefonica
Sur-3 building, 3rd floor
Madrid 28050
Spain

Email: luismiguel.contrerasmurillo@telefonica.com