

LISP Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: July 13, 2019

A. Rodriguez-Natal  
Cisco Systems, Inc.  
V. Ermagan  
Google  
F. Maino  
D. Dukes  
P. Camarillo  
C. Filsfils  
Cisco Systems, Inc.  
January 9, 2019

LISP Control Plane for SRv6 Endpoint Mobility  
draft-rodrigueznatal-lisp-srv6-01

## Abstract

This document describes the use of the LISP Control Plane to support endpoint mobility and Location/ID separation in Segment Routing v6 (SRv6) deployments.

## 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 July 13, 2019.

## Copyright Notice

Copyright (c) 2019 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

Internet-Draft

LISP-SRV6

January 2019

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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Overview . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Operation . . . . .	<a href="#">4</a>
<a href="#">3.1.</a>	Provisioning . . . . .	<a href="#">4</a>
<a href="#">3.2.</a>	Endpoint Registration . . . . .	<a href="#">5</a>
<a href="#">3.3.</a>	Endpoint Resolution . . . . .	<a href="#">6</a>
<a href="#">3.4.</a>	SR Policy Resolution . . . . .	<a href="#">6</a>
<a href="#">3.4.1.</a>	Decentralized . . . . .	<a href="#">7</a>
<a href="#">3.4.2.</a>	Centralized . . . . .	<a href="#">7</a>
<a href="#">3.5.</a>	Traffic Encapsulation . . . . .	<a href="#">8</a>
<a href="#">3.6.</a>	Endpoint Mobility . . . . .	<a href="#">9</a>
<a href="#">4.</a>	Segment Routing LCAF (SR LCAF) . . . . .	<a href="#">9</a>
<a href="#">4.1.</a>	SR-MPLS . . . . .	<a href="#">10</a>
<a href="#">4.2.</a>	SRv6 . . . . .	<a href="#">11</a>
<a href="#">4.3.</a>	Traffic Steering Tag . . . . .	<a href="#">12</a>
<a href="#">5.</a>	References . . . . .	<a href="#">12</a>
<a href="#">5.1.</a>	Normative References . . . . .	<a href="#">12</a>
<a href="#">5.2.</a>	Informative References . . . . .	<a href="#">13</a>
	Authors' Addresses . . . . .	<a href="#">13</a>

## [1.](#) Introduction

This document defines how to use the LISP Control Plane [[I-D.ietf-lisp-rfc6833bis](#)] to support endpoint mobility and Location/ID separation in those IPv6 deployments that are using SRv6 Network Programming [[I-D.filsfils-spring-srv6-network-programming](#)]. The LISP control plane is used to lookup "where" an endpoint is located, while SRv6 specifies "how" to program the SRv6 network infrastructure to transport traffic to that location.

To enable this, the egress Provider Edge (PE) nodes register via LISP control plane their local SRv6 Segment IDs (SIDs) to be used to reach endpoints attached to them. Ingress PE nodes retrieve via LISP control plane this mapping information, and use SRv6 network programming to encapsulate and steer the traffic towards the destination egress PE node.

The LISP control plane provides on-demand endpoint-to-SID mapping, as well as support for anchorless dynamic endpoint mobility.

## [2.](#) Overview

The ingress PE nodes receive traffic from endpoints in their local networks, encapsulate it in IPv6 packets following SRv6 policies and forward it to the SRv6 domain. Similarly, egress PE nodes receive SRv6 traffic from the SRv6 domain, remove the SRv6 encapsulation and forward the traffic to one of their locally attached networks according to the decapsulation SID received in the packets. Ingress PE nodes and egress PE nodes can be co-located in the same node, in this document we use "PE node" to refer to those collocated nodes.

This specification leverages the LISP Mapping System to store mappings of endpoints to decapsulation SIDs. Endpoints are neither LISP nor SRv6 aware and can roam freely across different PE nodes. The mappings in the LISP Mapping System can be used by PE nodes to know to which PE node an endpoint is currently connected. The LISP Mapping System is composed of LISP Map-Resolvers and Map-Servers but, for convenience, this document refers to the Mapping System as a single logical entity.

To use the LISP Mapping System, in this specification the PE nodes also implement the control-plane logic of LISP Ingress/Egress Tunnel Routers (xTRs). As a LISP xTR, a PE node keeps a local database of all the endpoints that are attached to its local network(s). It also keeps a list of local decapsulation function SIDs that map to each of its local network(s). A PE node registers into the LISP Mapping System its local endpoint-to-SID mappings. These mappings also contain the traffic steering tag associated with the endpoint. In addition to its local mappings, a PE node also keeps a local map-cache of remote endpoint-to-SID mappings that it uses to find the SID to use to encapsulate traffic towards a remote endpoint.

This specification also assumes an SR Path Computation Element (PCE) [[RFC4655](#)] that can compute and provide SRv6 paths from a given ingress PE node to a given egress PE node. The paths are based on the ingress and egress PE nodes and on the traffic steering tag that

is associated with the destination endpoint. Figure 1 shows an example diagram to be used as a reference for the architecture described in this document.

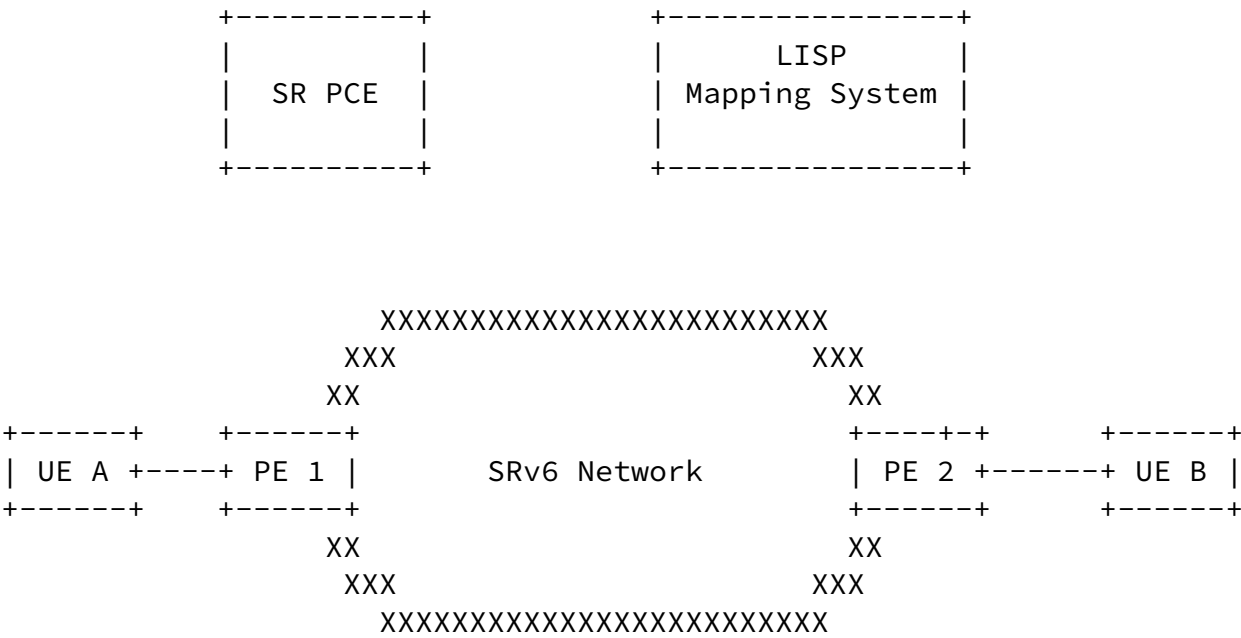


Figure 1: Reference Architecture

3. Operation

3.1. Provisioning

Each PE node is connected to the SRv6 domain and to one or more local networks. These local networks may represent different tenants/VPNs. Each tenant network is globally identified via a unique Instance

Identifier (IID). Locally, these tenant networks are identified at each PE node by the local IP table assigned to them. While the IID for a tenant network is global across the domain, the local IP table assigned to it in a given PE node is local to that PE node. Each PE node needs to be provisioned with the one-to-one mapping of global IID to local IP table per each tenant network it is serving. The provisioning of the IID to the local IP table is out of the scope of this document. PE nodes use this IID to local IP table information to know which IID they need to use when requesting mappings for traffic coming from a particular tenant network (i.e. from a particular local IP table).

Per each local IP table, each PE node has a local SRv6 "End.DT46" function (see [[I-D.filsfils-spring-srv6-network-programming](#)]) that decapsulates SRv6 traffic and forwards the inner traffic for lookup into that particular IP table. The PE node concatenates one of its local SRv6 locators with each of these decapsulation functions to create SIDs that point to each of the different tenant networks it is serving. These SIDs are kept onto the "My local SIDs table" of the

PE node and they are used to decapsulate incoming SRv6 traffic into the appropriate local IP table.

Beyond SRv6 state, each PE node has to be provisioned with the address of at least one Map-Resolver it will use to retrieve remote endpoint-to-SID mappings from the Mapping System. Similarly, it has to be provisioned with the address of the Map-Server to which it is going to register their local endpoint-to-SID mappings.

### [3.2.](#) Endpoint Registration

Endpoints attach to the local networks served by PE nodes. Upon attachment of an endpoint, the PE node will correlate the local IP table that corresponds to the local network where the endpoint attached to a global IID. It does so by using its local information of global IID to local IP table. The local IP table also correlates with the local SID that remote PE nodes need to use to send SRv6 encapsulated traffic towards the endpoint. The local SID directly translates into which local IP table the PE node should use to lookup for the endpoint when receiving traffic for it.

It should be noted that the endpoint does not need to attach to the

PE node directly (i.e. it can attach to another network device downlink) as long as the PE node is notified (e.g. via off-band orchestration) of the following:

- o Which endpoint has been attached (i.e. Endpoint EID)
- o Which tenant network (if any) the endpoint belongs to (i.e. Endpoint IID)

For this version of the specification, it is assumed that the endpoint only attaches to a single PE node and that all traffic steering will be handled by SRv6. Therefore, for this version of the specification, a PE only register one SID per endpoint (or group of endpoints) into the Mapping System. Future versions of this specification will describe how to an endpoint can be served by more than one PE node and/or by more than one SID per PE node.

For SRv6 operation, endpoints need to be associated with a particular tag to be used for traffic steering policies. This means that the traffic addressed towards the endpoint may need to be steered through a particular path in the SRv6 domain. This draft assumes that a PE node knows the tag associated to endpoints attached to the local networks it is serving. How a PE node knows which tag corresponds to a particular endpoint is out of the scope of this document. In addition to the endpoint tag, to compute the path through the SRv6 network the loopback address of the egress PE node is also used.

Therefore, to make all this information available to the LISP-SRv6 deployment, the PE node will register into the Mapping System the mapping of endpoint address and IID to endpoint's tag, PE node local SID and PE node loopback. To do so, the egress PE node will send a Map-Register as specified in [[I-D.ietf-lisp-rfc6833bis](#)] with the appropriate IID and endpoint address as EID. The endpoint's tag, the PE node local SID and the PE node loopback are encoded using the SR LCAF described in [Section 4](#) as an RLOC in the Map-Register.

### [3.3.](#) Endpoint Resolution

When a PE node needs to encapsulate traffic from a local endpoint towards a remote endpoint served by a remote PE node, it looks up in its map-cache to find the appropriate destination SID (and SR policy) to use to reach the remote endpoint. This lookup takes into

consideration the local IP table serving the local endpoint (i.e the IID of the mapping). If no entry is found for that remote endpoint, the PE node has to retrieve it from the Mapping System. To do so, it follows the procedures described in [\[I-D.ietf-lisp-rfc6833bis\]](#) and sends a Map-Request towards the Mapping System. In the Map-Request the PE node encodes its loopback address as ITR RLOC and as destination EID the address of the remote endpoint for which a destination SID is needed. It uses the IID associated with the local IP table serving the local endpoint that triggered the request.

What the Mapping System replies via a Map-Reply depends on how the SR policy is resolved. The Mapping System can reply with either the destination SID only (along with egress PE loopback address and endpoint tag) or with the complete SID list to be applied in the SRv6 underlay. This two options are discussed in detail in [Section 3.4](#). Note that the Map-Reply may return a prefix as returned EID, which means all the endpoints within the prefix can be reached through the same SID. Optionally, the PE node can request to also be subscribed to updates in the endpoint(s) mapping following [\[I-D.ietf-lisp-pubsub\]](#).

Alternatively, it is also possible for a PE node to subscribe in advance for endpoint(s) mappings for a set (or sets) of endpoint EIDs. That way the map-cache will be pre-populated for those destination endpoint(s), avoiding the need for an on-demand Map-Request/Map-Reply. This optimization is at the cost of keeping more state in the PE node and Mapping System.

### [3.4.](#) SR Policy Resolution

Besides retrieving the destination SID for a remote endpoint, a PE node also needs to find a suitable SR policy to send the traffic towards the destination SID through the SRv6 underlay. The

architecture discussed in this document offers different options to make the SR policy available at the PE node.

#### [3.4.1.](#) Decentralized

+-----+	+-----+
	LISP





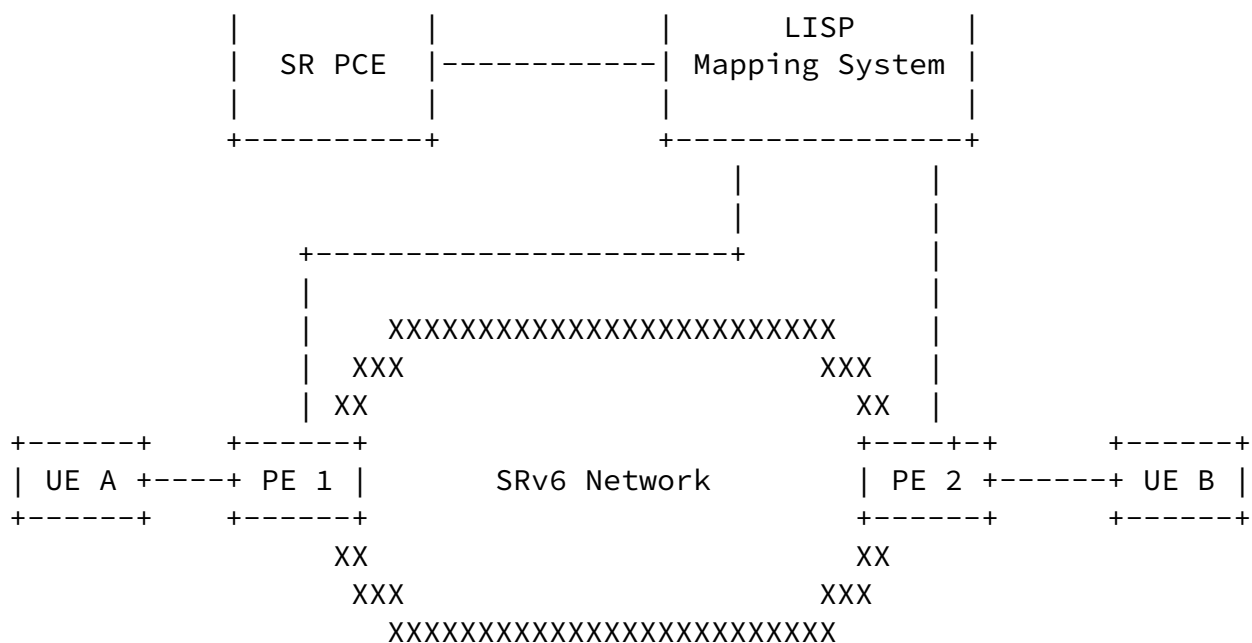


Figure 3: Centralized SR Policy Resolution

With the centralized approach the SR policies are resolved at the Mapping System when the mapping is requested by the PE node. Upon receiving a Map-Request for an endpoint from a PE node, the Mapping System queries the SR PCE to find the SR policy using [RFC5440]. To query the SR PCE, the Mapping System uses the ingress and egress PE nodes loopback addresses and the traffic steering tag of the requested endpoint. The Mapping System obtains the loopback address of the ingress PE node from the ITR-RLOC field of the Map-Request. In this case, the Mapping System returns not only the destination SID of the remote endpoint, but also the rest of the SIDs that the traffic has to go through from the ingress PE node to the egress PE node. The SR policy SID list along with the destination egress PE node decapsulation SID are encoded as an Explicit Locator Path (ELP) [RFC8060] in the Map-Reply returned. The ingress PE node can directly use this ELP to build the SRv6 packets and does not need to query the PCE to obtain the SR policy.

### 3.5. Traffic Encapsulation

Once the destination SID and SR policy for a given endpoint EID are known by a PE node, the PE node will use this information to build the Segment Routing Header (SRH), if needed, and encapsulate the traffic through the SRv6 domain towards the egress PE node. Note that if there is no SR policy for a particular endpoint, no SRH is needed and the packets can be encapsulated using a vanilla IPv6 header with the destination SID as destination address. This follows

common SRv6 operation as specified in [\[I-D.filsfils-spring-srv6-network-programming\]](#).

When the SRv6 traffic reaches the destination, the destination SID points to the local IP table where the decapsulated traffic has to be delivered. Once decapsulated, the traffic will be routed towards the intended endpoint via lookup in that local IP table.

### [3.6.](#) Endpoint Mobility

LISP-SRV6 deployment relies on the LISP mechanisms defined in [\[I-D.ietf-lisp-eid-mobility\]](#) to support mobility of endpoints. Following that specification, when a PE node registers an endpoint mapping, the previous PE node that had registered a mapping for the same endpoint will be notified. This serves to (1) notify the former egress PE node for the endpoint that the endpoint has moved and (2) to let former PE node know the new egress PE node for the endpoint. Once the former PE node receives the notification, it (1) stops registrations for the endpoint, (2) re-encapsulates any traffic received for the endpoint towards the new egress PE node, and (3) sends a Solicit-Map-Request message to any ingress PE node from which it receives traffic for the endpoint to let them know that they should trigger a Map-Request to update their map-caches.

In addition, if the Mapping System supports the Publish/Subscribe mechanisms described in [\[I-D.ietf-lisp-pubsub\]](#), a PE node can ask the Mapping System to be notified of changes in the mapping for a particular destination endpoint when it requests the mapping. This way a PE node subscribed to a particular endpoint will receive a mapping update with the new destination SID for the endpoint whenever the endpoint moves to a new PE node.

## [4.](#) Segment Routing LCAF (SR LCAF)

The following Segment Routing LISP Canonical Address Format (SR LCAF) [\[RFC8060\]](#) is introduced to support the operation of LISP-SR deployments.

Internet-Draft

LISP-SRv6

January 2019




Figure 4: Segment Routing LCAF

The SR LCAF defines the SR-Type field to indicate the type of Segment Routing and the specific format of the SR LCAF. The following values are defined:

0: Reserved


1: SR-MPLS

2: SRv6

For definitions of the rest of the LCAF fields please refer to [\[RFC8060\]](#).

#### [4.1.](#) SR-MPLS

When the SR-Type is SR-MPLS (SR-Type = 1) the SR LCAF has the following format:



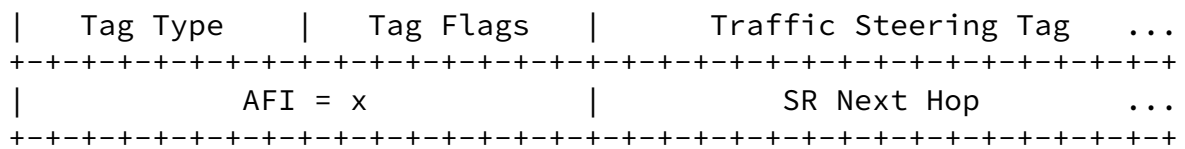


Figure 5: SR-MPLS Segment Routing LCAF

The SR-MPLS SR LCAF defines the following fields:

MPLS label: MPLS VPN label associated with the EID.

Tag Type: Type of traffic steering tag associated with the EID. Details on this traffic steering tag and different Tag Types are discussed in [Section 4.3](#).

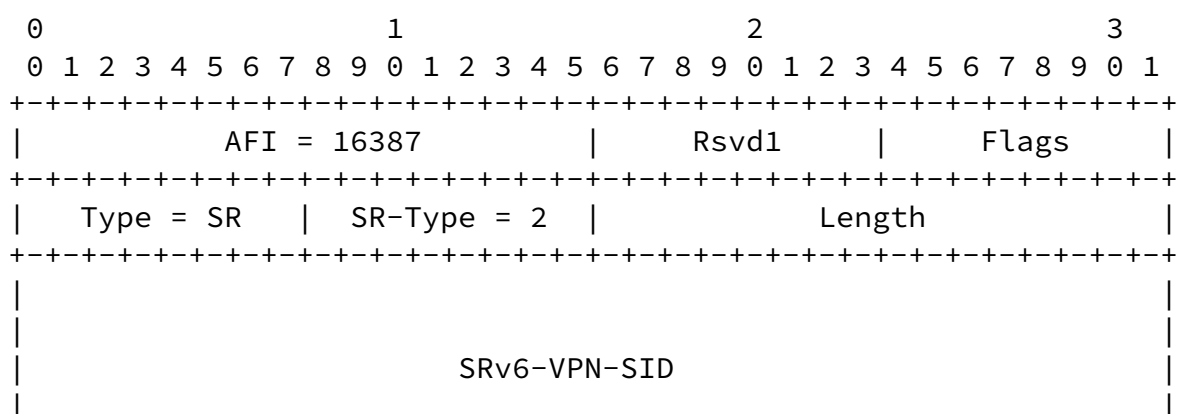
Tag Flags: Flags for the particular type of traffic steering tag associated with the EID.

Traffic Steering Tag: Tag associated with the EID that is used to compute SR paths.

SR Next Hop: Loopback address of the egress PE node associated with the EID.

## 4.2. SRv6

When the SR-Type is SR-SRv6 (SR-Type = 2) the SR LCAF has the following format:



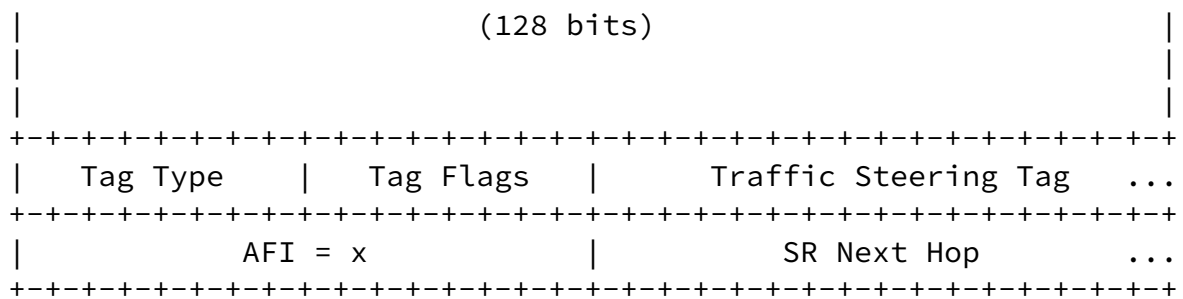


Figure 6: SRv6 Segment Routing LCAF

The SRv6 SR LCAF defines the following fields:

SRv6-VPN-SID: The SRv6 VPN SID associated with the EID.

See [Section 4.1](#) for definitions of the rest of the fields of the SRv6 SR LCAF.

### 4.3. Traffic Steering Tag

The SR LCAF supports different traffic steering tags to be associated with the EID and be used in the operation of the SR deployment. The following values are defined for the Tag Type field:

0: No tag. When the Tag Type is 0 the Tag Flags are 0 and the Tag field has length of 0 octets.

1: Color. When the Tag Type is 0 the Tag Flags and Tag field have the following format.

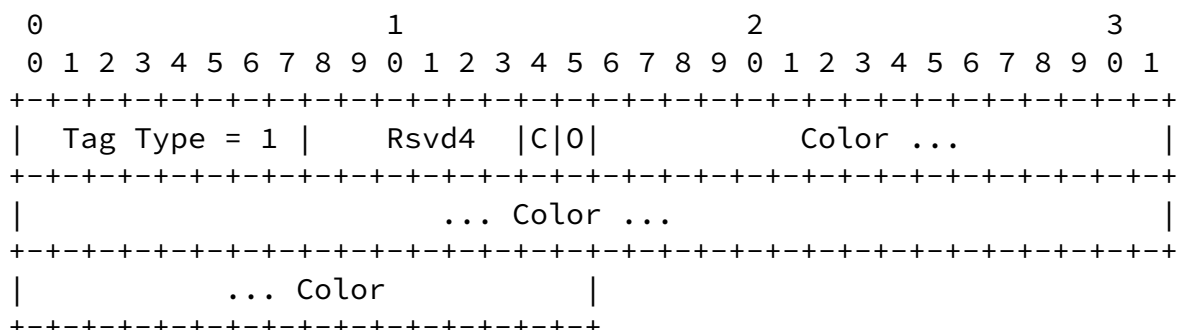


Figure 7: Color Tag

2: Slice. The tag format and flags for this Tag Type are to be defined in future versions of this document.

## 5. References

### 5.1. Normative References

- [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>>.
- [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>>.

- [RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", [RFC 8060](#), DOI 10.17487/RFC8060, February 2017, <<https://www.rfc-editor.org/info/rfc8060>>.

### 5.2. Informative References

- [I-D.filsfils-spring-srv6-network-programming]  
Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-filsfils-spring-srv6-network-programming-06](#) (work in progress), October 2018.
- [I-D.ietf-lisp-eid-mobility]  
Portoles-Comeras, M., Ashtaputre, V., Moreno, V., Maino, F., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", [draft-ietf-lisp-eid-mobility-03](#) (work in progress), November 2018.

[I-D.ietf-lisp-pubsub]

Rodriguez-Natal, A., Ermagan, V., Leong, J., Maino, F., Cabellos-Aparicio, A., Barkai, S., Farinacci, D., Boucadair, M., Jacquenet, C., and S. Secci, "Publish/Subscribe Functionality for LISP", [draft-ietf-lisp-pubsub-02](#) (work in progress), November 2018.

[I-D.ietf-lisp-rfc6833bis]

Fuller, V., Farinacci, D., and A. Cabellos-Aparicio, "Locator/ID Separation Protocol (LISP) Control-Plane", [draft-ietf-lisp-rfc6833bis-23](#) (work in progress), December 2018.

#### Authors' Addresses

Alberto Rodriguez-Natal  
Cisco Systems, Inc.  
United States of America

Email: natal@cisco.com

Vina Ermagan  
Google  
United States of America

Email: ermagan@gmail.com

Fabio Maino  
Cisco Systems, Inc.  
United States of America

Email: fmaino@cisco.com

Darren Dukes  
Cisco Systems, Inc.  
Canada

Email: ddukes@cisco.com

Pablo Camarillo  
Cisco Systems, Inc.  
Spain

Email: pcamaril@cisco.com

Clarence Filsfils  
Cisco Systems, Inc.  
Belgium

Email: cf@cisco.com