

Internet Engineering Task Force
Internet-Draft
Updates: 8059, 9798 (if approved)
Intended status: Standards Track
Expires: 30 December 2025

V. Govindan
S. Venaas
Cisco
28 June 2025

PIM Join Attributes for Locator/ID Separation Protocol (LISP)
Environments
draft-vs-pim-rfc8059-9798bis-00

Abstract

This document defines two PIM Join/Prune attributes that support the construction of multicast distribution trees where the root and receivers are located in different Locator/ID Separation Protocol (LISP) sites. These attributes allow the receiver site to select between unicast and multicast underlying transport, to convey the RLOC (Routing Locator) address of the receiver ETR (Egress Tunnel Router) to the control plane of the root ITR (Ingress Tunnel Router) and to signal the underlay multicast group to the control plane of the root ITR. This document updates RFC 8059 and RFC 9798.

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 30 December 2025.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. PIM Join/Prune Attributes	4
3. The Transport Attribute	4
3.1. Transport Attribute Format	4
3.2. Using the Transport Attribute	5
4. Receiver ETR RLOC Attribute	5
4.1. Using the Received ETR RLOC Attribute for Unicast Underlays	5
4.2. Using the Received ETR RLOC Attribute for Multicast Underlays	6
4.2.1. Multicast Address Range constraints:	6
4.3. Receiver RLOC Attribute format	7
4.4. Using the Receiver RLOC Attribute	8
5. Acknowledgements	9
6. IANA Considerations	9
7. Security Considerations	9
8. Contributors	10
9. Normative References	10
Authors' Addresses	12

1. Introduction

The construction of multicast distribution trees where the root and receivers are located in different LISP sites [RFC9300] is defined in [I-D.ietf-lisp-rfc6831bis]. Creation of (root-EID,G) state in the root site requires that unicast LISP-encapsulated Join/Prune messages be sent from an ETR on the receiver site to an ITR on the root site. The term "EID" is short for "Endpoint ID". [RFC9300] specifies that (root-EID,G) data packets are to be LISP- encapsulated into (root-RLOC,G) multicast packets. However, a wide deployment of multicast connectivity between LISP sites is unlikely to happen any time soon. In fact, some implementations are initially focusing on unicast transport with head-end replication between root and receiver sites.

The unicast LISP-encapsulated Join/Prune message specifies the (root-EID,G) state that needs to be established in the root site, but conveys nothing about the receiver's capability or desire to use multicast as the underlying transport. This document specifies a Join/Prune attribute that allows the receiver ETR to select the desired transport.

The term "transport" in this document is intentionally somewhat vague. Currently, it is used just to indicate whether multicast or head-end replication is used; this means that the outer destination address is either a unicast or multicast address. Future documents may specify how other types of delivery, encapsulation, or underlay are used.

Knowledge of the receiver ETR's RLOC address is essential to the control plane of the root ITR. The RLOC address determines the downstream destination for unicast head-end replication and identifies the receiver ETR that needs to be notified should the root ITR of the distribution tree move to another site. The root ITR can change when the source EID is roaming to another LISP site.

Service providers may implement unicast reverse path forwarding (uRPF) policies requiring that the outer source address of the LISP-encapsulated Join/Prune message be the address of the receiver ETR's core-facing interface used to physically transmit the message. However, due to policy and load-balancing considerations, the outer source address may not be the RLOC on which the receiver site wishes to receive a particular flow. This document specifies a Join/Prune attribute that conveys the appropriate receiver ETR's RLOC address to the control plane of the root ITR.

This document uses terminology defined in [RFC9300] and [RFC9301], such as EID, RLOC, ITR, and ETR.

The construction of multicast distribution trees where the root and receivers are located in different LISP sites [RFC9300] is defined in [I-D.ietf-lisp-rfc6831bis].

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. PIM Join/Prune Attributes

PIM Join/Prune attributes are defined in [RFC5384] by introducing a new Encoded-Source type that, in addition to the Join/Prune source, can carry multiple Type-Length-Value (TLV) attributes. These attributes apply to the individual Join/Prune sources on which they are stored.

The attributes defined in this document conform to the format of the encoding type defined in [RFC5384]. The attributes would typically be the same for all the sources in the Join/Prune message. Hence, we RECOMMEND using the hierarchical Join/Prune attribute scheme defined in [RFC7887]. This hierarchical system allows attributes to be conveyed in the Upstream Neighbor Address field, thus enabling the efficient application of a single attribute instance to all the sources in the Join/Prune message.

LISP Tunnel Routers (xTRs) do not exchange PIM Hello Messages, and hence no Hello option is defined to negotiate support for these attributes. Systems that support unicast head-end replication are assumed to support these attributes.

3. The Transport Attribute

It is essential that a mechanism be provided by which the desired transport can be conveyed by receiver sites. Root sites with multicast connectivity will want to leverage multicast replication. However, not all receiver sites can be expected to have multicast connectivity. It is thus desirable that root sites be prepared to support (root-EID,G) state with a mixture of multicast and unicast output state. This document specifies a Join/Prune attribute that allows the receiver to select the desired underlying transport.

3.1. Transport Attribute Format




Figure 1

- * F-bit: The Transitive bit. Specifies whether the attribute is transitive or non-transitive. MUST be set to zero. This attribute is ALWAYS non-transitive.

- * End-of-Attributes bit. Specifies whether this attribute is the last. Set to zero if there are more attributes. Set to 1 if this is the last attribute.
- * The Transport Attribute type is TBD.
- * Length: The length of the Transport Attribute value. MUST be set to 1.
- * Transport: The type of transport being requested. Set to 0 for multicast. Set to 1 for unicast. The values from 2 to 255 may be assigned in the future.

3.2. Using the Transport Attribute

Hierarchical Join/Prune attribute instances [RFC7887] SHOULD be used when the same Transport Attribute is to be applied to all the sources within the Join/Prune message or all the sources within a group set. The root ITR MUST accept Transport Attributes in the Upstream Neighbor Encoded-Unicast address, Encoded-Group addresses, and Encoded-Source addresses.

There MUST NOT be more than one Transport Attribute within the same encoded address. If an encoded address has more than one instance of the attribute, the root ITR MUST discard all affected Join/Prune sources. The root ITR MUST also discard all affected Join/Prune sources if the Transport Attribute value is unknown.

4. Receiver ETR RLOC Attribute

4.1. Using the Received ETR RLOC Attribute for Unicast Underlays

When a receiver ETR requests unicast head-end replication for a given (root-EID,G) entry, the PIM control plane of the root ITR must maintain an outgoing interface list ("oif-list") entry for the receiver ETR and its corresponding RLOC address. This allows the root ITR to perform unicast LISP-encapsulation of multicast data packets to each and every receiver ETR that has requested unicast head-end replication.

The PIM control plane of the root ITR could potentially determine the RLOC address of the receiver ETR from the outer source address field of the LISP-encapsulated Join/Prune message. However, receiver ETRs are subject to uRPF checks by the network providers on each core-facing interface. The outer source address must therefore be the RLOC of the core-facing interface used to physically transmit the LISP-encapsulated Join/Prune message. Due to policy and load-balancing considerations, that may not be the RLOC on which the receiver site wishes to receive a particular flow. This document specifies a Join/Prune attribute that conveys the appropriate receiver RLOC address to the PIM control plane of the root ITR.

To support root-EID mobility, receiver ETRs must also be tracked by the LISP control plane of the root ITR, regardless of the underlying transport. When the root-EID moves to a new root ITR in a different LISP site, the receiver ETRs do not know the root-EID has moved and therefore do not know the RLOC of the new root ITR. This is true for both unicast and multicast transport modes. The new root ITR does not have any receiver ETR state. Therefore, it is the responsibility of the old root ITR to inform the receiver ETRs that the root-EID has moved. When the old root ITR detects that the root-EID has moved, it sends a LISP Solicit-Map-Request (SMR) message to each receiver ETR. The receiver ETRs do a mapping database lookup to retrieve the RLOC of the new root ITR. The old root ITR detects that the root-EID has moved when it receives a Map-Notify from the Map-Server. The transmission of the Map-Notify is triggered when the new root ITR registers the root-EID [EID-MOBILITY]. When a receiver ETR determines that the root ITR has changed, it will send a LISP-encapsulated PIM prune message to the old root xTR and a LISP-encapsulated PIM join message to the new root xTR.

4.2. Using the Received ETR RLOC Attribute for Multicast Underlays

When LISP based Multicast trees are constructed using IP Multicast in the underlay, the mapping between the overlay group address and the underlay group address becomes a crucial engineering decision:

Three distinct types of overlay to underlay group mappings are possible:

Many to one mapping:

Many (root-EID, G) flows originating from an RLOC can be mapped to a single underlay multicast (root-RLOC, G-u) flow.

One to many mapping:

Conversely a single same overlay flow can be mapped to two or more flows, e.g., (root-RLOC, G-u1) and (root-RLOC, G-u2) to cater to the requirements of downstream xTR nodes.

One to one mapping:

Every (root-EID, G) flow is mapped to a unique (root-RLOC, G-u) flow.

4.2.1. Multicast Address Range constraints:

Under certain conditions, different subsets of xTRs subscribing to the same overlay multicast stream may be constrained to use distinct underlay multicast mapping ranges.

This introduces a trade-off between replication overhead and the flexibility of address range assignment, which may be necessary in specific use-cases like Proxy Tunnel Routers or when using nodes with limited hardware resources as explained below:

Inter-site Proxy Tunnel Routers (PxTR):

When multiple LISP sites are interconnected through a LISP-based transit, the site border node (PxTR) connects the site-facing interfaces with the external LISP core. In such cases, different ranges of multicast group addresses may be used for constructing (S-RLOC, G) trees within the LISP site and in the external LISP core. This distinction is desirable for various operational reasons.

Hardware resource restrictions:

Platform limitations may necessitate engineering decisions to restrict multicast address ranges in the underlay due to hardware resource constraints.

4.3. Receiver RLOC Attribute format




Figure 2

- * F bit: The Transitive bit. Specifies whether this attribute is transitive or non-transitive. MUST be set to zero. This attribute is ALWAYS non-transitive.
- * E bit: End-of-Attributes bit. Specifies whether this attribute is the last. Set to zero if there are more attributes. Set to 1 if this is the last attribute.
- * Type: The Receiver RLOC Attribute type is TBD[suggested: 6].
- * Length: The length in octets of the attribute value. MUST be set to the length in octets of the receiver RLOC address plus 1 octet to account for the Address Family field.
- * Addr Family: The PIM Address Family of the receiver RLOC as defined in [RFC7761].

- * Receiver RLOC: The RLOC address on which the receiver ETR wishes to receive the encapsulated flow. A unicast IP Receiver RLOC address is used for unicast-encapsulated flows. Alternately, a multicast IP Receiver RLOC address is used for for multicast-encapsulated flows. A multicast IP address MUST be used only when the underlay network of the LISP core supports IP Multicast transport.

When the ITR needs to track the list of ETRs from which the PIM joins are received, the ITR MUST use the source IP address field of the incoming PIM Join/Prune message. The source IP address of the PIM Join/Prune MUST be an ETR RLOC IP address.

4.4. Using the Receiver RLOC Attribute

Hierarchical Join/Prune attribute instances [RFC7887] SHOULD be used when the same Receiver RLOC Attribute is to be applied to all the sources within the message or all the sources within a group set. The root ITR MUST accept Transport Attributes in the Upstream Neighbor Encoded-Unicast address, Encoded-Group addresses, and Encoded-Source addresses.

There MUST NOT be more than one Receiver RLOC Attribute within the same encoded address. If an encoded address has more than one instance of the attribute, the root ITR MUST discard all affected Join/Prune sources. The root ITR MUST also discard all affected Join/Prune sources if the address family is unknown or the address length is incorrect for the specified address family.

When the ETR determines to use the multicast underlay:

- * It chooses an underlay multicast group that it can join. This is a matter of local decision, beyond the scope of this document.
- * It identifies the upstream LISP site where the underlay multicast tree needs to be rooted.
- * It constructs the PIM Join/Prune message as specified in RFC 8059 [RFC8059] and RFC 9798 [RFC9798]. Only the Receiver RLOC attribute is encoded as above.

When the ITR receives a PIM Join/Prune message:

- * It allocates a new entry in the OutgoingInterfaceList [I-D.ietf-lisp-rfc6831bis] for every unique underlay multicast mapping.

- * The ITR MAY apply local policy to perform any kind of rate-limiting on the number of copies it needs to make in the underlay. Such actions are beyond the scope of this document.

5. Acknowledgements

TBD

6. IANA Considerations

Following the guidelines of [RFC8126], IANA is asked to assign values as follows:

Two new PIM Join/Prune attribute types have been requested: value 5 for the Transport Attribute and value 6 for the Receiver RLOC Attribute [RFC8059] and [RFC9798].

The "PIM Join/Prune Transport Types" registry has been created for the Join/Prune Transport attribute. The registration policy is IETF Review [RFC5226], and the values are in the range 0-255. This document requests value 0 for multicast and value 1 for unicast [RFC8059] and [RFC9798].

7. Security Considerations

Security of Join/Prune attributes is only guaranteed by the security of the PIM packet. The attributes specified herein do not enhance or diminish the privacy or authenticity of a Join/Prune message. A site that legitimately or maliciously sends and delivers a Join/Prune message to another site will equally be able to append these and any other attributes it wishes. See [RFC5384] for general security considerations for Join/Prune attributes.

An attack vector arises where an attacker sends numerous PIM Join messages with different group addresses. This could interfere with legitimate multicast traffic if the group addresses overlap. Additionally, resource exhaustion may occur if replication is requested for a large number of groups, potentially resulting in significant resource consumption.

To mitigate these risks, PIM authentication mechanisms RFC 5796 [RFC5796] could be employed to validate join requests. Furthermore, implementations may consider explicit tracking mechanisms to manage joins more effectively. Configurable controls could be introduced, allowing for a maximum permissible number of groups for each ETR RLOC used as the source of overlay joins. These controls would limit the impact of such attacks and ensure that resource allocation is managed appropriately.

8. Contributors

Jesus Arango
TBD
Email: tbd@tbd.com

Isidor Kouvelas
Arista
Email: kouvelas@arista.com

Dino Farinacci
lispers.net
Email: farinacci@gmail.com

9. Normative References

- [I-D.ietf-lisp-rfc6831bis]
Farinacci, D., Meyer, D., Zwiebel, J., Venaas, S., and V. P. Govindan, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6831bis-01, 29 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-rfc6831bis-01>>.
- [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>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.

- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, DOI 10.17487/RFC5384, November 2008, <<https://www.rfc-editor.org/info/rfc5384>>.
- [RFC5796] Atwood, W., Islam, S., and M. Siami, "Authentication and Confidentiality in Protocol Independent Multicast Sparse Mode (PIM-SM) Link-Local Messages", RFC 5796, DOI 10.17487/RFC5796, March 2010, <<https://www.rfc-editor.org/info/rfc5796>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7887] Venaas, S., Arango, J., and I. Kouvelas, "Hierarchical Join/Prune Attributes", RFC 7887, DOI 10.17487/RFC7887, June 2016, <<https://www.rfc-editor.org/info/rfc7887>>.
- [RFC8059] Arango, J., Venaas, S., Kouvelas, I., and D. Farinacci, "PIM Join Attributes for Locator/ID Separation Protocol (LISP) Environments", RFC 8059, DOI 10.17487/RFC8059, January 2017, <<https://www.rfc-editor.org/info/rfc8059>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <<https://www.rfc-editor.org/info/rfc9300>>.
- [RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <<https://www.rfc-editor.org/info/rfc9301>>.

[RFC9798] Govindan, V. and S. Venaas, "PIM Join/Prune Attributes for Locator/ID Separation Protocol (LISP) Environments Using Underlay Multicast", RFC 9798, DOI 10.17487/RFC9798, June 2025, <<https://www.rfc-editor.org/info/rfc9798>>.

Authors' Addresses

Vengada Prasad Govindan
Cisco
Email: venggovi@cisco.com

Stig Venaas
Cisco
Email: svenaas@cisco.com