

Internet Engineering Task Force (IETF)
Request for Comments: 9798
Updates: 8059
Category: Experimental
ISSN: 2070-1721

V. Govindan
S. Venaas
Cisco Systems, Inc.
June 2025

PIM Join/Prune Attributes for Locator/ID Separation Protocol (LISP) Environments Using Underlay Multicast

Abstract

This document specifies an update to the Receiver RLOC (Routing Locator) field of the PIM Join/Prune attribute that supports the construction of multicast distribution trees where the source and receivers are located in different Locator/ID Separation Protocol (LISP) sites and are connected using underlay IP multicast. This attribute allows the receiver site to signal the underlay multicast group to the control plane of the root Ingress Tunnel Router (ITR). This document updates RFC 8059.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9798>.

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
 - 1.1. Requirements Language
2. The Case for Extending the Received ETR RLOC Attribute of RFC 8059
 - 2.1. Flexible Mapping of Overlay to Underlay Group Ranges
 - 2.2. Multicast Address Range Constraints

- 3. Updates to RFC 8059
 - 3.1. Scope
 - 3.2. Receiver ETR RLOC Attribute
 - 3.3. Using the Receiver RLOC Attribute
- 4. IANA Considerations
- 5. Security Considerations
- 6. Normative References
- Acknowledgements
- Authors' Addresses

1. Introduction

The construction of multicast distribution trees where the root and receivers are located in different LISP sites [RFC9300] is defined in [RFC6831].

[RFC6831] specifies that (EID, G) data packets are to be LISP-encapsulated into (RLOC, G) multicast packets. In this document, we use the term root-EID or root-RLOC to refer to the source of the multicast tree rooted at the EID or RLOC. [RFC8059] defines PIM Join/Prune attribute extensions to construct multicast distribution trees. Please refer to Section 3 of [RFC6831] for the definition of the terms Endpoint ID (EID) and Routing Locator (RLOC). This document extends the Receiver ETR RLOC PIM Join/Prune attribute [RFC8059] to facilitate the construction of underlay multicast trees for (root-RLOC, G).

Specifically, the assignment of the underlay multicast group needs to be done in consonance with the downstream Tunnel Router (xTR) nodes needed to avoid unnecessary replication or traffic hairpinning.

Since the Receiver RLOC Attribute defined in [RFC8059] only addresses the Ingress Replication case, this document extends the scope of that PIM Join/Prune attribute to include scenarios where the underlay uses multicast transport. The scope extension complies with the base specification [RFC5384].

This document uses terminology defined in [RFC6831], such as EID, RLOC, ITR and ETR.

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. The Case for Extending the Received ETR RLOC Attribute of RFC 8059

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.

2.1. Flexible Mapping of Overlay to Underlay Group Ranges

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.

2.2. 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 (i.e., 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.

3. Updates to RFC 8059

3.1. Scope

There are no changes to the syntax or semantics of the Transport Attribute defined in [RFC8059].

The scope of the updates to [RFC8059] is limited to the case where the "Transport" field of the Transport Attribute is set to zero (multicast) only.

3.2. Receiver ETR RLOC Attribute

The definition of the "Receiver RLOC" field of the Receiver ETR RLOC attribute (see Section 5.1 of [RFC8059]) is updated as follows:

OLD:

| Receiver RLOC: The RLOC address on which the receiver ETR wishes
| to receive the unicast-encapsulated flow.

NEW:

| 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 multicast-
| encapsulated flows. A multicast IP address MUST be used only
| when the underlay network of the LISP core supports IP
| multicast transport.

The definitions of the other fields of the Receiver ETR RLOC Attribute remain unchanged.

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.

3.3. Using the Receiver RLOC Attribute

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, which is 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 [RFC8059]. 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 outgoing interface list [RFC6831] 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.

4. IANA Considerations

This document has no IANA actions.

5. Security Considerations

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 [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.

6. 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>>.
- [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>>.

- [RFC6831] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", RFC 6831, DOI 10.17487/RFC6831, January 2013, <<https://www.rfc-editor.org/info/rfc6831>>.
- [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>>.
- [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>>.

Acknowledgements

The authors would like to thank Dino Farinacci, Victor Moreno, Alvaro Retana, Aswin Kuppusami, Joe Clarke, and Peter Yee for their valuable comments. The authors also thank Sankaralingam T and Amit Kumar for their contributions to the document. The authors thank Gunter Van de Velde for his valuable comments.

Authors' Addresses

Vengada Prasad Govindan
Cisco Systems, Inc.
Email: venggovi@cisco.com

Stig Venaas
Cisco Systems, Inc.
Tasman Drive
San Jose, CA 95134
United States of America
Email: stig@cisco.com