

Workgroup: Internet Engineering Task Force

Internet-Draft:

draft-ietf-pim-jp-jp-extensions-lisp-01

Published: 19 April 2022

Intended Status: Experimental

Expires: 21 October 2022

Authors: V. Govindan S. Venaas

Cisco

Cisco

PIM Join/ Prune Attributes for LISP Environments using Underlay Multicast

Abstract

This document specifies an extension to PIM Receiver RLOC Join/ Prune attribute that supports the construction of multicast distribution trees where the root 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 ITR (Ingress Tunnel Router).

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 21 October 2022.

Copyright Notice

Copyright (c) 2022 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](#)
- [3. Acknowledgements](#)
- [4. Contributors](#)
- [5. IANA Considerations](#)
- [6. Security Considerations](#)
- [7. Normative References](#)
- [Authors' Addresses](#)

1. Introduction

The construction of multicast distribution trees where the root and receivers are located in different LISP sites [[RFC6830](#)] is defined in [[RFC6831](#)].

[[RFC6831](#)] specifies that (root-EID, G) data packets are to be LISP-encapsulated into (root-RLOC, G) multicast packets. [[RFC8059](#)] defines PIM J/P attribute extensions to construct multicast distribution trees. This document extends the Receiver ETR RLOC PIM J/P 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 xTR nodes and avoid unnecessary replication or traffic hairpinning.

Since the Receiver RLOC Attribute defined in [[RFC8059](#)] only addresses the Ingress Replication case, an extension of the scope of that PIM J/P attribute is defined by this draft to include scenarios where the underlay uses Multicast transport. The scope extension proposed here complies with the base specification [[RFC5384](#)].

This document uses terminology defined in [[RFC6830](#)], 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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. The case for extending the Received ETR RLOC Attribute of RFC 8059

When LISP based Multicast trees can be built using IP Multicast in the underlay, the mapping between the overlay group address and the underlay group address becomes a very crucial engineering decision:

Flexible mapping of overlay to underlay group ranges:

Three different types of overlay to underlay group mappings are possible: Many to one mapping: Many (root-EID, G) flows originating from a RLOC can be mapped to the same underlay (root-RLOC, G-u) flow. One to many mapping: Conversely the 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 different (root-RLOC, G-u) flow. The overlay can use ASM while the underlay can use SSM ranges.

Multicast Address Range constraints:

It is possible that under certain circumstances, different subsets of xTRs subscribing to the same overlay multicast stream would be constrained to use different underlay multicast mapping ranges. This definitely involves a trade-off between replication and the flexibility in assigning address ranges and could be required in certain situations further below.

Inter-site PxTR:

When multiple LISP sites are connected through a LISP based transit, the site border node interconnects the site-facing interfaces and the external LISP based core. Under such circumstances, there could be different ranges of multicast group addresses used for building the (S-RLOC, G) trees inside the LISP site and the external LISP based core. This is desired for various reasons:

Hardware resource restrictions:

Platform limitations could force engineering decisions to be made on restricting multicast address ranges in the underlay.

Other Use-cases:

TBD

Editorial Note: Comments from Stig: There should be some text indicating that the group address used should ideally only be used for LISP encapsulation (if ASM), and perhaps that it is preferable to use an SSM group. Also, that the group obviously must be a group that the underlay supports/allows. I think it is also worth noting that ideally, different ETRs should request the same group.

3. Acknowledgements

The authors would like to thank Dino Farinacci and Victor Moreno for their valuable comments.

4. Contributors

Sankaralingam
Cisco

Email: sankt@cisco.com

Amit Kumar
Cisco

Email: kumaram3@cisco.com

5. IANA Considerations

No new requests to IANA

6. Security Considerations

There is perhaps a new attack vector where an attacker can send a bunch of joins with different group addresses. It may interfere with other multicast traffic if those group addresses overlap. Also, it may take up a lot of resources if replication for thousands of groups are requested. However PIM authentication (?) should come to the rescue here. TBD Since explicit tracking would be done, perhaps it is worth enforcing that for each ETR RLOC (the RLOC used as the source of the overlay join), there could be a configurable number of maximum permissible group(s). TBD

Ed Note: To be addressed - Comments from Stig: Regarding security considerations and PIM authentication. The only solution we have here is to use IP-Sec to sign the J/P messages. I don't know if anyone has tried to use IPSec between LISP RLOCs. Are there any LISP security mechanisms that would help here for authenticating LISP encapsulated messages between xTRs?

7. 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>>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.

[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>>.

[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>>.

[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>>.

Authors' Addresses

Vengada Prasad Govindan
Cisco

Email: venggovi@cisco.com

Stig Venaas
Cisco

Email: svenaas@cisco.com