Workgroup: Internet Engineering Task Force
Internet-Draft:
draft-ietf-pim-jp-jp-extensions-lisp-05
Updates: 8059 (if approved)
Published: 26 February 2024
Intended Status: Experimental
Expires: 29 August 2024
Authors: V. Govindan
Cisco
S. Venaas
Cisco

PIM Join/ Prune Attributes for LISP Environments using Underlay Multicast

Abstract

This document specifies an update to the 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 Ingress Tunnel Router (ITR).

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 29 August 2024.

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

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 (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 [RFC9300], 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.

Multicast Address Range constraints:
It is possible that under certain circumstances, differnt 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.

3. Updates to RFC 8059

3.1. Scope

There are no changes proposed to the syntax and semantics of the Transport attributed defined in RFC8059 [RFC8059].

The scope of the updates proposed to RFC 8059 [RFC8059] are 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 RFC 8059 [RFC8059] is updated as follows:

Receiver RLOC:
The Receiver RLOC field of the Receiver RLOC Attribute MAY contain a multicast IP address. This 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. The source IP 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, beyond the scope of this document.

* It identifies the upstream LISP site where the underlay multicast tree tree needs to be rooted.

* It constructs the PIM Join/ Prune message as specified in RFC 8059 [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 oif-list 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. Acknowledgements

The authors would like to thank Dino Farinacci, Victor Moreno, Alvaro Retana and Aswin Kuppusami for their valuable comments. The

authors also thank Sankaralingam T and Amit Kumar for their contributions to the document.

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 could be used here. Since explicit tracking could be done, an implementation may consider knobs to ensure 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).

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>.

[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>.

[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>.

Authors' Addresses

Vengada Prasad Govindan
Cisco

Email: venggovi@cisco.com

Stig Venaas
Cisco

Email: svenaas@cisco.com