

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

J. Arango
S. Venaas
Cisco Systems
I. Kouvelas
Arista Networks Inc.
D. Farinacci
lispers.net
January 2017

PIM Join Attributes
for Locator/ID Separation Protocol (LISP) Environments

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 and 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).

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 a candidate 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 <http://www.rfc-editor.org/info/rfc8059>.

Copyright Notice

Copyright (c) 2017 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 (<http://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 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

1. Introduction	2
2. Requirements Notation	3
3. PIM Join/Prune Attributes	3
4. The Transport Attribute	4
4.1. Transport Attribute Format	4
4.2. Using the Transport Attribute	5
5. Receiver ETR RLOC Attribute	5
5.1. Receiver RLOC Attribute Format	6
5.2. Using the Receiver RLOC Attribute	6
6. Security Considerations	7
7. IANA Considerations	7
8. References	8
8.1. Normative References	8
8.2. Informative References	8
Authors' Addresses	9

1. Introduction

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

[RFC6831] 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 [RFC6830], such as EID, RLOC, ITR, and ETR.

2. Requirements Notation

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 [RFC2119].

3. 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.

4. 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.

4.1. Transport Attribute Format



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.

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 Transport Attribute type is 5.

Length: The length of the Transport Attribute value. MUST be set to 1.

Transport: The type of transport being requested. Set to zero for multicast. Set to 1 for unicast. The values from 2 to 255 may be assigned in the future.

4.2. 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.

5. Receiver ETR RLOC Attribute


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.

5.1. Receiver RLOC Attribute Format



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 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 receiver the unicast-encapsulated flow.

5.2. 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.

6. 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.

7. IANA Considerations

Two new PIM Join/Prune attribute types have been assigned: value 5 for the Transport Attribute and value 6 for the Receiver RLOC Attribute.

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 assigns value 0 for multicast and value 1 for unicast.

8. References

8.1. 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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc7887>>.

8.2. Informative References

- [EID-MOBILITY] Portoles-Comeras, M., Ashtaputre, V., Moreno, V., Maino, F., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", Work in Progress, draft-portoles-lisp-eid-mobility-01, October 2016.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

Authors' Addresses

Jesus Arango
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
United States of America

Email: jearango@cisco.com

Stig Venaas
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
United States of America

Email: stig@cisco.com

Isidor Kouvelas
Arista Networks Inc.
5453 Great America Parkway
Santa Clara, CA 95054
United States of America

Email: kouvelas@arista.com

Dino Farinacci
lispers.net
San Jose, CA
United States of America

Email: farinacci@gmail.com

