

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: November 2, 2022

D. Farinacci
lispers.net
May 1, 2022

LISP Distinguished Name Encoding
draft-farinacci-lisp-name-encoding-14

Abstract

This draft defines how to use the AFI=17 Distinguished Names in LISP.

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 November 2, 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 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. Definition of Terms	3
3. Distinguished Name Format	4
4. Example Use-Cases	5
5. Name Collision Considerations	5
6. Security Considerations	5
7. IANA Considerations	5
8. References	5
8.1. Normative References	5
8.2. Informative References	6
Appendix A. Acknowledgments	6
Appendix B. Document Change Log	6
B.1. Changes to draft-farinacci-lisp-name-encoding-14	6
B.2. Changes to draft-farinacci-lisp-name-encoding-13	6
B.3. Changes to draft-farinacci-lisp-name-encoding-12	7
B.4. Changes to draft-farinacci-lisp-name-encoding-11	7
B.5. Changes to draft-farinacci-lisp-name-encoding-10	7
B.6. Changes to draft-farinacci-lisp-name-encoding-09	7
B.7. Changes to draft-farinacci-lisp-name-encoding-08	7
B.8. Changes to draft-farinacci-lisp-name-encoding-07	7
B.9. Changes to draft-farinacci-lisp-name-encoding-06	7
B.10. Changes to draft-farinacci-lisp-name-encoding-05	8
B.11. Changes to draft-farinacci-lisp-name-encoding-04	8
B.12. Changes to draft-farinacci-lisp-name-encoding-03	8
B.13. Changes to draft-farinacci-lisp-name-encoding-02	8
B.14. Changes to draft-farinacci-lisp-name-encoding-01	8
B.15. Changes to draft-farinacci-lisp-name-encoding-00	8
Author's Address	8

1. Introduction

The LISP architecture and protocols [RFC6830] introduces two new numbering spaces, Endpoint Identifiers (EIDs) and Routing Locators (RLOCs) which are intended to replace most use of IP addresses on the Internet. To provide flexibility for current and future applications, these values can be encoded in LISP control messages using a general syntax that includes Address Family Identifier (AFI) [RFC1700].

The length of the value field is implicit in the type of address that follows. For AFI 17, a Distinguished Name can be encoded. A name can be a variable length field so the length cannot be determined solely from the AFI value 17. This draft defines a termination character, an 8-bit value of 0 to be used as a string terminator so the length can be determined.

LISP Distinguished Names are useful when encoded either in EID-records or RLOC-records in LISP control messages. As EIDs, they can be registered in the mapping system to find resources, services, or simply used as a self-documenting feature that accompany other address specific EIDs. As RLOCs, Distinguished Names, along with RLOC specific addresses and parameters, can be used as labels to identify equipment type, location, or any self-documenting string a registering device desires to convey.

2. Definition of Terms

Address Family Identifier (AFI): a term used to describe an address encoding in a packet. An address family currently defined for IPv4 or IPv6 addresses. See [AFI] and [RFC1700] for details on other types of information that can be AFI encoded.

3. Distinguished Name Format

An AFI=17 Distinguished Name is encoded as:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
AFI = 17										ASCII String ...																													
... ASCII String																				0																			

The string of characters are encoded in the ASCII character-set definition [RFC0020].

When Distinguished Names are encoded for EIDs, the EID-prefix length of the EIDs as they appear in EID-records for all LISP control messages is the length of the string in bits (include the null 0 byte). Where Distinguished Names are encoded anywhere else (i.e. nested in LCAF encodings), then any length field is the length of the ASCII string including the null 0 byte in units of bytes.

When Map-Requests are sent for an EID encoded in Distinguished Name format, an exact match request is performed. So the Map-Server (when configured for proxy-Map-Replying) or the ETR will return a Map-Reply with the same EID-prefix length.

4. Example Use-Cases

This section identifies three specific use-cases for the Distinguished Name format. Two are used for an EID encoding and one for a RLOC-record encoding. When storing public keys in the mapping system, as in [I-D.ietf-lisp-ecdsa-auth], a well known format for a public-key hash can be encoded as a Distinguished Name. When street location to GPS coordinate mappings exist in the mapping system, as in [I-D.farinacci-lisp-geo], the street location can be a free form ascii representation (with whitespace characters) encoded as a Distinguished Name. An RLOC that describes an xTR behind a NAT device can be identified by its router name, as in [I-D.farinacci-lisp-simple-nat], uses a Distinguished Name encoding. As well as identifying the router name (neither an EID or an RLOC) in NAT Info-Request messages uses Distinguished Name encodings.

5. Name Collision Considerations

When a Distinguished Name encoding is used to format an EID, the uniqueness and allocation concerns are no different than registering IPv4 or IPv6 EIDs to the mapping system. See [I-D.ietf-lisp-rfc6833bis] for more details. Also, the use-case documents specified in Section 4 provide allocation recommendations for their specific uses.

6. Security Considerations

There are no security considerations.

7. IANA Considerations

The code-point values in this specification are already allocated in [AFI].

8. References

8.1. Normative References

- [AFI] "Address Family Identifier (AFIs)", ADDRESS FAMILY NUMBERS <http://www.iana.org/numbers.html>, February 2007.
- [I-D.ietf-lisp-rfc6833bis] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, "Locator/ID Separation Protocol (LISP) Control-Plane", draft-ietf-lisp-rfc6833bis-30 (work in progress), November 2020.

- [RFC0020] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<https://www.rfc-editor.org/info/rfc20>>.
- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, DOI 10.17487/RFC1700, October 1994, <<https://www.rfc-editor.org/info/rfc1700>>.
- [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>>.

8.2. Informative References

- [I-D.farinacci-lisp-geo]
Farinacci, D., "LISP Geo-Coordinate Use-Cases", draft-farinacci-lisp-geo-13 (work in progress), March 2022.
- [I-D.farinacci-lisp-simple-nat]
Farinacci, D., "A Simple LISP NAT-Traversal Implementation", draft-farinacci-lisp-simple-nat-03 (work in progress), November 2021.
- [I-D.ietf-lisp-ecdsa-auth]
Farinacci, D. and E. Nordmark, "LISP Control-Plane ECDSA Authentication and Authorization", draft-ietf-lisp-ecdsa-auth-07 (work in progress), February 2022.

Appendix A. Acknowledgments

The author would like to thank the LISP WG for their review and acceptance of this draft.

Appendix B. Document Change Log

B.1. Changes to draft-farinacci-lisp-name-encoding-14

- o Submitted May 2022.
- o Update references and document expiry timer.

B.2. Changes to draft-farinacci-lisp-name-encoding-13

- o Submitted November 2021.
- o Update references and document expiry timer.

- B.3. Changes to draft-farinacci-lisp-name-encoding-12
 - o Submitted May 2021.
 - o Update references and document expiry timer.
- B.4. Changes to draft-farinacci-lisp-name-encoding-11
 - o Submitted November 2020.
 - o Made changes to reflect working group comments.
 - o Update references and document expiry timer.
- B.5. Changes to draft-farinacci-lisp-name-encoding-10
 - o Submitted August 2020.
 - o Update references and document expiry timer.
- B.6. Changes to draft-farinacci-lisp-name-encoding-09
 - o Submitted March 2020.
 - o Update references and document expiry timer.
- B.7. Changes to draft-farinacci-lisp-name-encoding-08
 - o Submitted September 2019.
 - o Update references and document expiry timer.
- B.8. Changes to draft-farinacci-lisp-name-encoding-07
 - o Submitted March 2019.
 - o Update referenes and document expiry timer.
- B.9. Changes to draft-farinacci-lisp-name-encoding-06
 - o Submitted September 2018.
 - o Update document expiry timer.

- B.10. Changes to draft-farinacci-lisp-name-encoding-05
- o Submitted March 2018.
 - o Update document expiry timer.
- B.11. Changes to draft-farinacci-lisp-name-encoding-04
- o Submitted September 2017.
 - o Update document expiry timer.
- B.12. Changes to draft-farinacci-lisp-name-encoding-03
- o Submitted March 2017.
 - o Update document expiry timer.
- B.13. Changes to draft-farinacci-lisp-name-encoding-02
- o Submitted October 2016.
 - o Add a comment that the distinguished-name encoding is restricted to ASCII character encodings only.
- B.14. Changes to draft-farinacci-lisp-name-encoding-01
- o Submitted October 2016.
 - o Update document timer.
- B.15. Changes to draft-farinacci-lisp-name-encoding-00
- o Initial draft submitted April 2016.

Author's Address

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 21 September 2022

M. Kowal
M. Portoles
Cisco Systems
A. Jain
Juniper Networks
D. Farinacci
lispers.net
20 March 2022

LISP Transport for Policy Distribution
draft-kowal-lisp-policy-distribution-02

Abstract

This document describes the use of the Locator/ID Separation Protocol (LISP) to encode and transport data models for the configuration of LISP ITRs.

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 [RFC2119].

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 September 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	2
2. Definition of Terms	2
3. Policy Distribution Use Cases	3
4. Policy Distribution: Packet Flow Description	3
4.1. Policy Distribution	4
4.2. Policy Updates	5
5. Mapping System Operations	6
6. Policy Distribution Process	6
7. Policy Distribution Encoding	6
8. IANA Considerations	7
9. Acknowledgements	7
10. References	7
10.1. Normative References	7
10.2. Informative References	8
Authors' Addresses	8

1. Introduction

When LISP ITRs are deployed with enough configuration to build a LISP overlay, they may require additional configurations such as security, QoS, and/or traffic forwarding policies. As networks continue to grow, it can be challenging to ensure these configurations are distributed to many ITRs and kept in sync. LISP network operators may wish to re-use their existing LISP architecture to distribute these configurations as opposed to configuring them by hand, using a script, or investing in a configuration management system. The configurations can be distributed via a mapping system that the network operator manages or is managed by a third-party as part of a managed service offering.

2. Definition of Terms

LISP related terms are defined as part of the LISP specification [RFC6830], notably EID, RLOC, Map-Request, Map-Reply, Map-Notify, Ingress Tunnel Router (ITR), Egress Tunnel Router (ETR), Map-Server (MS) and Map-Resolver (MR).

3. Policy Distribution Use Cases

The ITR could use the mapping system to receive configuration policies for use cases such as:

- * The RLOC interfaces of an ITR may be connected to WAN links that are policed at sub-line rate by its upstream provider. Using the mapping system, the ITR could receive and apply the QoS policies that would shape traffic to the correct rate on each ITR RLOC interface.
- * ITRs use the mapping system to receive access-list (ACL) configuration(s) that would allow them to restrict traffic from authorized sources to authorized services.
- * ITRs receive configurations that determine local forwarding policies, such as specifying ITR RLOCs to be used for egress forwarding on a per-application basis or RLOCs on different ITRs within the same LISP site to maintain application symmetry.
- * Baseline configurations for common services (e.g., DNS, SSH, Syslog) can be maintained in a mapping system and distributed across multiple ITRs.

Policy distribution is not meant to provide zero-touch provisioning for ITRs within a LISP network. At a minimum, the ITR must have a map resolver defined, IP connectivity to the map resolver, and one or more distinguished names defined for receiving specific policies from the mapping system.

4. Policy Distribution: Packet Flow Description

The following figure illustrates a reference system used to support packet flow descriptions in this section.




Figure 1: Reference system for policy distribution

The reference system contains two sites, site A and site B, with corresponding xTR-A and xTR-B providing encapsulation and decapsulation services for the overlay traffic. xTR-A uses interface-A to forward and receive encapsulated traffic through the RLOC space; and xTR-B uses interface-B for it.

For packet flow purposes the reference system assumes that a network controller provides the policies to a map-server.

When an ITR comes up, it requests it's designated policies with it's map-server. The MS may have this policy configured by the administrator via a network controller.

4.1. Policy Distribution

The following is an illustration of the sequence to distribute a policy registered by the controller with the mapping system, down to an ITR that requests its designated policies. In the example <ITR-A> represents the hostname of the ITR that learns a policy using this mechanism.

- * The Mapping-System is either configured by an operator or learns a mapping sent by a controller through a Map-Register. The Mapping System learns the mapping: EID="policy-<ITR-A>" --> RLOC= "{ "shape":{ "interface":"ifaceA", "direction":"outbound", "value":100Mbps } }". The EID is encoded as a Distinguished Name and the RLOC as a JSON string.

- * ITR-A is configured to dynamically learn policies from the Mapping System with the name "policy-ITR-A" (policy followed by its hostname).
- * ITR-A sends a Map-Request to the Mapping System with EID="policy-<ITR-A>" encoded as a Distinguished Name. The Map-Request is sent with the N-bit set.
- * The Mapping System forwards the request to the appropriate Map-Server. The Map-Server adds ITR-A to the subscription list of EID="policy-<ITR-A>" and sends back a Map-Notify with the mapping that the controller has registered.
- * When ITR-A receives the Map-Notify installs the received policy locally, to shape traffic sent over the RLOC facing interface.
- * Note that when the map-server has multiple policies associated with this ITR, it can send each one of the policies as an additional locator record (following the same JSON format) in the mapping. The locator count in the Map-Notify reflects the number of policies distributed with the mapping.

4.2. Policy Updates

Policy distribution takes advantage of the LISP pubsub model to ensure that router updates are properly distributed when policies change. In such a case, and using the same reference system as above, the information exchange is as follows:

- * The controller sends a Map-Register to the Mapping System, updating the policy mapping with: EID="policy-<ITR-A>" --> RLOC="{ "shape":{ "interface":"ifaceA", "direction":"outbound", "value":200Mbps } }".
- * When the corresponding Map-Server receives this update it checks the list of ITRs subscribed for updates of EID="policy-<ITR-A>" and finds out that ITR-A is subscribed.
- * The Map-Server sends a Map-Notify to ITR-A with the updated mapping information that has been registered.
- * When ITR-A receives and validates the Map-Notify, it updates the local policy, changing the shaping rate as specified in the new JSON description. Note that if the JSON specifies the same policy that is currently applied the notification is ignored.

5. Mapping System Operations

The mapping system that is used for distributing policy configurations can be managed by either the administrator who owns and operates their own LISP sites or a third-party administrator who offers LISP mapping system functionality as a managed service. A controller or orchestrator could be used to update and optimize policies within the mapping system based on network or ITR telemetry.

Within the mapping system, the administrator must define a distinguished name that is specific to an ITR. The distinguished name is associated with the specific policy configurations that the ITR is to receive. Each ITR is configured with the minimal requirements to perform a mapping request procedure as well as a distinguished name that can be matched upon in the mapping system.

Map-Servers should be able to receive policy registrations through the Map-Registration process. The Map-Registration must encode the policy following the specification in the policy distribution encoding section.

6. Policy Distribution Process

The ITR subscribes to its policy via the Map-Request procedure defined in section 5 of [I-D.ietf-lisp-pubsub]. The PubSub procedure is used to ensure that policies can be updated or audited after an ITR has received them. Policies are published to the ITR from the mapping system using the mapping notification procedure defined in section 6 of [I-D.ietf-lisp-pubsub].

EID-to-RLOC mappings used for policy distribution are of the type EID <Distinguished Name> to RLOC <JSON policy specification>. The EID is a distinguished name uniquely identifying a router in the system, while each RLOC record uses JSON encoding to specify the particular policy (or policies) that this router needs to implement.

7. Policy Distribution Encoding

When the ITR is configured to receive a policy using a distinguished name, the ITR sends a subscription for the EID record encoded as this Distinguished Name. When a policy has been registered with the Mapping System for this Distinguished Name, the ITR receives a publication with a list of policies as RLOC records and encoded as JSON strings (as defined in section 5.4 of [RFC8060]).

Example encoding for QoS policy that shapes traffic to 50 percent of the line-rate: EID-Record encoded as distinguished name "policy-ce-router1" RLOC-Record record encoded as JSON string

```
{"shape":{"interface":"ethernet1","direction":"outbound",  
"unit":"percent","value":50}}
```

Example encoding for setting the ITR's NTP server to 10.10.10.10: EID-Record encoded as distinguished name "policy-ce-router" RLOC-Record record encoded as JSON string

```
{"NTP-address": "10.10.10.10"}
```

Multiple ITRs can be configured to use multiple distinguished names for receiving multiple sets policies. This allows for an ITR to receive specific policies and many ITRs to receive policies that can be broadly applied. Referring to the two examples above, an ITR can be configured to use a distinguished name of "policy-ce-router1" to receive a QoS configuration that is specific to that node while also using a distinguished name of "policy-ce-router" to receive configurations that are common to each ITR in the LISP network (e.g., NTP configuration). The use of multiple distinguished names per ITR reduces the amount of configuration within the mapping system.

8. IANA Considerations

This memo includes no request to IANA.

9. Acknowledgements

Thanks to James Stankiewicz for his thorough comments and suggestions.

10. References

10.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", RFC 8060, DOI 10.17487/RFC8060, February 2017, <<https://www.rfc-editor.org/info/rfc8060>>.

10.2. Informative References

[I-D.ietf-lisp-pubsub]
Rodriguez-Natal, A., Ermagan, V., Cabellos-Aparicio, A.,
Barkai, S., and M. Boucadair, "Publish/Subscribe
Functionality for LISP", Work in Progress, Internet-Draft,
draft-ietf-lisp-pubsub-07, 8 January 2021,
<<http://www.ietf.org/internet-drafts/draft-ietf-lisp-pubsub-07.txt>>.

Authors' Addresses

Michael Kowal
Cisco Systems
111 Wood Ave. South
Iselin, NJ 08830
United States of America
Email: mikowal@cisco.com

Marc Portoles Comeras
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
United States of America
Email: mportole@cisco.com

Amit Jain
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States of America
Email: atjain@juniper.net

Dino Farinacci
lispers.net
San Jose, CA
United States of America
Email: farinacci@gmail.com

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: 6 December 2021

V. Govindan
S. Venaas
Cisco
4 June 2021

PIM Join/ Prune Attributes for LISP Environments using Underlay
Multicast
draft-vgovindan-pim-jp-extensions-lisp-01

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 6 December 2021.

Copyright Notice

Copyright (c) 2021 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 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
1.1. Requirements Language	3
2. The case for extending the Received ETR RLOC Attribute of RFC 8059	3
3. Acknowledgements	4
4. Contributors	4
5. IANA Considerations	4
6. Security Considerations	4
7. Normative References	5
Authors' Addresses	5

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