

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 13 November 2025

V. Govindan, Ed.
Cisco
D. Farinacci
lispers.net
A. Kuppusami
S. Venaas
Cisco
12 May 2025

LISP Multicast Overlay Group to Underlay RLOC Mappings
draft-ietf-lisp-group-mapping-02

Abstract

This draft augments LISP [RFC9300] multicast functionality described in [I-D.farinacci-lisp-rfc6831bis] and [I-D.farinacci-lisp-rfc8378bis] to support the mapping of overlay group addresses to underlay RLOC addresses. This draft defines a many-to-1, 1-to-many, and many-to-many relationship between multicast EIDs and the Replication List Entries (RLEs) RLOC records they map to. The mechanisms in this draft allow a multicast LISP overlay to run over a mixed underlay of unicast and/or multicast packet forwarding functionality.

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 13 November 2025.

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. Requirements Language	2
2. Introduction	2
3. Definition of Terms	3
4. Source Site Procedures	4
5. Hashed Based Underlay Group Selection	5
6. Mapping System Lookup Based Underlay Group Selection	5
7. Receiver Site Procedures	6
8. Interworking with non-Overlay Receivers	7
9. IANA Considerations	7
10. Security Considerations	7
11. Normative References	7
Appendix A. Acknowledgements	9
Appendix B. Document Change Log	9
B.1. Changes to draft-ietf-lisp-group-mapping-02	9
B.2. Changes to draft-ietf-lisp-group-mapping-01	9
B.3. Changes to draft-ietf-lisp-group-mapping-00	9
B.4. Changes to draft-vda-lisp-group-mapping-03	9
B.5. Changes to draft-vda-lisp-group-mapping-02	9
B.6. Changes to draft-vda-lisp-group-mapping-01	9
B.7. Changes to draft-vda-lisp-group-mapping-00	10
B.8. Changes to draft-vda-lisp-underlay-multicast-trees-00	10
Authors' Addresses	10

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 [RFC2119].

2. Introduction

When a multicast capable underlay connects multiple LISP sites, we can take advantage of the multicast capabilities and perform replication more efficiently than using head-end replication. This draft addresses the problem of selecting the underlay multicast group(s) to transport a given overlay multicast flow. There are 4 different scenarios possible:

1. A 1:1 mapping of an overlay multicast flow, either a (Source, Group) or (*, Group) to an underlay group.
2. A many:1 mapping where many overlay multicast flows share the same underlay group.
3. A 1:many mapping where a single overlay group is transported over different underlay groups. Note in this case, the "many" means mapping one EID to multiple RLOCs in a RLE-set where the set can be a mix of underlay groups and underlay unicast addresses. It is really important that this algorithm compute unicast RLOCs as well as multicast RLOCs.
4. Given the "many" mapping case relationship above, a 4th relationship of m:n many-to-many mappings exist and will be supported.

There are two methods being proposed to derive an underlay group from an overlay group:

- * Computation based underlay group assignment using a hash function. Here all ETRs will use the same multicast underlay group, when they are all on the same native multicast forwarding underlay. The ETRs hash the bits of the group address G for an (S-EID, G) flow. The low-order 24-bits of the hash is used to produce an IPv4 underlay group address. The low-order 120-bits of the hash is used to produce an IPv6 underlay group address.
- * Lookup based group assignment where a mapping system lookup is used to coordinate the assignment of the underlay group among ETRs.

The scope of this draft covers underlays based on IPv4 and IPv6 only. It does not cover other transport mechanisms like BIER, multicast MPLS, or layer-2 underlays.

3. Definition of Terms

Note all terminology used in this document is based on the formal definitions from [I-D.farinacci-lisp-rfc6831bis] and [I-D.farinacci-lisp-rfc8378bis]. The definitions below extend these formal definitions to address the introduction of group overlay to group/unicast underlay mappings.

(S-EID,G): refers to multicast state in multicast source and

receiver sites where S-EID is the IP address of the multicast source host (its EID). An S-EID can appear in an IGMPv3 report, an MSDP SA message or a PIM Join/Prune message that travels inside of a site. The notation (S-EID,G-EID) is also used to signify that this is an overlay source and overlay group entry.

(S-RLOC G): refers to multicast state in the underlay core where S-RLOC is a source address of the ITR that is performing encapsulation. The notation (S-RLOC,G-RLOC) is also used to denote G-RLOC as the underlay group address. The (S-RLOC,G-RLOC) is mapped from the (S-EID,G-EID) entry by doing a mapping database lookup for (S-EID,G-EID).

(S-RLOC,U-RLOC): refers to multicast state in the underlay core where S-RLOC is a source address of the ITR that is performing encapsulation. In this case, the underlay RLOC is a unicast address and denoted as U-RLOC. A U-RLOC is used for encapsulation when it is determined that the ETR has no multicast connectivity and cannot be reached via a G-RLOC. So the ITR encapsulates to the unicast address U-RLOC of the ETR.

RLE: is an Replication List Entry encoding used for replicating multicast packets which are encapsulated. The list can be made up of multiple G-RLOCs and U-RLOCs. The RLE appears in control messages in a single RLOC-record.

Multicast Mapping: an entry in the mapping system that maps a multicast (S-EID,G-EID) or (*,G-EID) entry to an RLOC-set. An example would be EID: (S-EID,G-EID) -> RLE-set: [G-RLOCa, G-RLOCb, U-RLOCc U-RLOCd] where the S-EID is the source of a multicast packet sending to a G-EID destination. The RLOC-record is an RLE entry made up of 4 replications, the first to G-RLOCa, the second to G-RLOCb, the third to U-RLOCc, and finally to U-RLOCd. It is at the discretion of the encapsulator (an ITR or RTR) what source RLOC address is used as the source address in the outer header.

4. Source Site Procedures

The Source Site Procedures from [I-D.farinacci-lisp-rfc8378bis] are followed for overlay nodes and [I-D.farinacci-lisp-rfc6831bis] for non-overlay nodes. There are no modifications to these procedures other than the source multicast site -ITR should be capable of replicating to multiple G-RLOC and U-RLOC addresses from its map-cache.

5. Hashed Based Underlay Group Selection

When an ETR receives an IGMP/MLD report it needs to decide which G-RLOC to join in the underlay. A hash-based algorithm can be used so all ETRs that process the report can join the same G-RLOC. Therefore when an (S-EID,G-EID) is reported, the hash will be over the pair of 32-bits for IPv4 IGMP reports and over the pair of 128-bits for IPv6 MLD reports. When (*,G-EID) is being reported, the hash is over just G-EID. The hash function used will be sha256 [RFC6234].

The hashed based approach creates perfect replication since it results in a 1-to-1 mapping, but at the expense of more underlay state.

6. Mapping System Lookup Based Underlay Group Selection

There will be scenarios where native multicast underlay provider will want to control what group addresses are used in the network. Therefore, a hashed based algorithm may be at conflict. In this case G-RLOCs need to be registered to the mapping system which are part of the provider supported group address block.

The proposed procedure is for the provider to register the G-RLOC as an RLOC record for a distinguished-name EID encoded from procedures in [I-D.ietf-lisp-name-encoding]. The EID will be registered in a well defined and configured instance-id with name "group-<group-address>".

For example, say there exists a G-EID of 224.1.1.1 and a provider G-RLOC of 225.1.1.1. When a receiver joins G-EID 224.1.1.1, the receiving ETR lookups up EID "group-224.1.1.1" which returns an G-RLOC of 225.1.1.1. The ETR uses 225.1.1.1 as the G-RLOC to register the (S-EID,G-EID) to the mapping system.

Using this method, the provider can create 1-to-many, many-to-1, and many-to-many relationships between G-EID and G-RLOC. The provider could also have EID "group-224.1.1.1" map to a (S-RLOC,G-RLOC) so an SSM based distribution tree can be joined in the underlay, by either PIM or IGMP/MLD. This example illustrates how to do a many-to-1 mapping by having multiple distinguish-name entries (encoded with different G-EID addresses) map to a single G-RLOC.

In a many-to-many scenario, the ETR could be configured with a G-EID set of prefixes so a power-of-2 range of G-EIDs would be looked up that returns a single G-RLOC. For example, say there exists 224.1.0.0/16 and 224.2.0.0/16 G-EID prefix ranges and G-RLOCs 225.1.1.1 and 225.2.2.2. The provider allows only 2 underlay groups

to be used but the overlay has large ranges of G-EIDs. So the provider registers to the mapping system "group-224.1.0.0-16" with G-RLOC 225.1.1.1 and "group-224.2.0.0-16" with G-RLOC 225.2.2.2. When an ETR receives an IGMP report for 224.1.1.1, it registers G-EID 224.1.1.1 with G-RLOC 225.1.1.1. It can even scale better, by registering G-EID 224.1.0.0/16 with G-RLOC 225.1.1.1 so subsequent IGMP reports in the 224.1.0.0/16 range would not need to be registered. But this does come at expense of receiving multicast packets for a G-EID when there are no receivers.

When using multicast and G-EID prefixes, there is a tradeoff between state in the underlay and using unnecessary bandwidth.

7. Receiver Site Procedures

The Receiver Site Procedures from [I-D.farinacci-lisp-rfc8378bis] are followed. This draft adds an additional step to the procedure on how to select the G-RLOC address for registration.

From the two approaches to obtain a G-RLOC address discussed in the previous sections, the ETR can start and complete the (S-EID,G-EID) registration process:

1. When ETR receives either an IGMP or MLD report, be it (S,G) or (*,G), the entry is used as the multicast EID to Map-Register to the mapping system. Just like [I-D.farinacci-lisp-rfc6831bis] and [I-D.farinacci-lisp-rfc8378bis] specify.
2. If the ETR is connected to a native multicast underlay, it will register (S-EID,G-EID) or (*,G-EID) with a G-RLOC based on one of G-RLOC selection options. All ETRs joining the same underlay group for a given overlay group MUST use the same G-RLOC selection option. For early versions of this design, the network operator configures the option consistently in all ETRs. The ITRs do not need such configuration since they use what the mapping system returns.
3. If the ETR is configured to know there are ITRs that are not on the same native multicast underlay, it MAY also register its U-RLOC. The ITR will determine which RLOC to use by the testing reachability via RLOC-probing according to [RFC9301].
4. Other ETRs registering the same (S-EID,G-EID) or (*,G-EID) entries, will choose the same underlay group G-RLOC so when the ITR replicates to G-RLOC, all ETRs receive the packet so they can deliver to multicast receivers.

5. After the registration is performed by ETRs, they need to tell the underlay to deliver multicast packets to them for G-RLOC. They do that by either sending PIM-Join messages into the underlay PIM domain or they send IGMP/MLD reports. This is left to the implementation to decide.

If ETRs underlay join on more than one interface, they may receive duplicate packets. Care must be taken to not IGMP/MLD join on multiple interfaces or duplicates will occur. If PIM joining occurs on different interfaces, RPF failures will occur to stop the duplicates from being delivered to receivers but at the expense of using unnecessary underlay bandwidth.

Packet duplicates can also occur when ETRs register both G-RLOC and their own U-RLOC addresses to the mapping system. ETRs cannot deregister one or the other when they see duplicates because different ITRs use the same mapping, so some will be on the multicast underlay and some will not. In the case, when they are on a multicast underlay, duplicates will occur across the RLOCs. The ETRs must monitor this and not answer RLOC probes sent to the U-RLOC so the ITR suppresses sending to it.

8. Interworking with non-Overlay Receivers

If there are multicast receivers that IGMP/MLD join a G-EID but are not attached to a native multicast underlay, they cannot receive multicast packets. They cannot receive multicast packets from overlay attached sources because the ITR has no ETR to encapsulate to. Hence, they are not on the overlay. However, if they are attached to a native multicast underlay, they can receive multicast packets.

There are numerous multicast connectivity combinations which are documented in detail in [I-D.farinacci-lisp-rfc6831bis]. Those procedures should be followed to deliver multicast packets from overlay attached sources to underlay only attached receivers. As well as non-overlay attach sources to overlay attached receivers.

9. IANA Considerations

There are no requests for IANA.

10. Security Considerations

There are no security considerations at this time.

11. Normative References

[I-D.farinacci-lisp-rfc6831bis]

Farinacci, D., Meyer, D., Zwiebel, J., Venaas, S., and V. P. Govindan, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", Work in Progress, Internet-Draft, draft-farinacci-lisp-rfc6831bis-02, 29 May 2024, <<https://datatracker.ietf.org/doc/html/draft-farinacci-lisp-rfc6831bis-02>>.

[I-D.farinacci-lisp-rfc8378bis]

Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", Work in Progress, Internet-Draft, draft-farinacci-lisp-rfc8378bis-00, 3 May 2024, <<https://datatracker.ietf.org/doc/html/draft-farinacci-lisp-rfc8378bis-00>>.

[I-D.ietf-lisp-name-encoding]

Farinacci, D. and L. Iannone, "LISP Distinguished Name Encoding", Work in Progress, Internet-Draft, draft-ietf-lisp-name-encoding-17, 9 December 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-name-encoding-17>>.

[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>>.

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.

[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>>.

[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>>.

[RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <<https://www.rfc-editor.org/info/rfc9301>>.

Appendix A. Acknowledgements

The authors would like to thank the LISP WG for their careful review and commentary. A special thank you goes to Stig Venaas.

Appendix B. Document Change Log

B.1. Changes to draft-ietf-lisp-group-mapping-02

- * Submitted May 2025.
- * Update document timer and references.

B.2. Changes to draft-ietf-lisp-group-mapping-01

- * Submitted November 2024.
- * Fixed typos found by Raphael Lienard's.

B.3. Changes to draft-ietf-lisp-group-mapping-00

- * Submitted September 2024.
- * Renamed draft-vda-lisp-group-mapping to draft-ietf-lisp-group-mapping now that the draft has been accepted as a working group document.

B.4. Changes to draft-vda-lisp-group-mapping-03

- * Submitted May 2024.
- * Added references to 6831bis and 8378bis as we put those RFCs to draft standards track.

B.5. Changes to draft-vda-lisp-group-mapping-02

- * Submitted April 2024.
- * Update document timer and references.

B.6. Changes to draft-vda-lisp-group-mapping-01

- * Submitted October 2023.
- * Update document timer and references.

B.7. Changes to draft-vda-lisp-group-mapping-00

- * Submitted April 2023.
- * Completed adding details to compliment [I-D.farinacci-lisp-rfc6831bis] and [I-D.farinacci-lisp-rfc8378bis].
- * Changed name to draft-vdas-lisp-group-mapping-00 from draft-vda-lisp-underlay-multicast-trees-00.

B.8. Changes to draft-vda-lisp-underlay-multicast-trees-00

- * Initial posting December 2020.

Authors' Addresses

Vengada Prasad Govindan (editor)
Cisco
Email: venggovi@cisco.com

Dino Farinacci
lisppers.net
Email: farinacci@gmail.com

Aswin Kuppusami
Cisco
Email: aswkuppu@cisco.com

Stig Venaas
Cisco
Email: stig@cisco.com