

LISP Working Group
Internet-Draft
Intended status: Informational
Expires: 4 October 2025

V. Govindan
M. Hamroz
J. Gawron
Cisco
2 April 2025

LISP Multicast Deployment Experience
draft-vgovindan-lisp-multicast-deploy-01

Abstract

We present our experiences in supporting deployments of LISP Multicast using unicast and multicast underlays. This document details design considerations that can be useful for anyone interested in deploying LISP multicast services over IP networks.

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 4 October 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. Introduction	2
2. Terminology	2
3. Scope of this document	3
4. Scope not covered by this document	3
5. Industry segments/ use-cases covered	3
6. Advantages and Cost of using "PIM-over-PIM"	3
7. Underlay Deployment considerations	4
7.1. Head-End Replication	4
7.2. Native Multicast Underlay	5
8. Layer-2 BUM overlay deployment considerations	5
8.1. RP for Layer-2 BUM Multicast Underlay	6
9. Layer-3 overlay Multicast deployment considerations	6
9.1. Layer-3 Routed Any-Source Multicast (ASM) services	6
9.1.1. Layer-3 overlay ASM RP placement and redundancy	7
9.1.2. Optimisation for short-lived streams	7
9.2. Layer-3 Routed SSM services	7
10. Mobility considerations for LISP multicast	8
11. Enabling Multicast flows for multiple tenants and multiple site overlays	8
11.1. Background reading	8
11.2. LISP Uerlay deployment considerations	8
12. Extranet Multicast (Route leaking)	9
13. IANA Considerations	9
14. Security Considerations	9
15. References	9
15.1. Normative References	9
15.2. Informative References	10
Acknowledgements	11
Contributors	11
Authors' Addresses	11

1. Introduction

This document describes deployment experiences of inter domain multicast routing function in a network where Locator/ID Separation is deployed using the Locator/ID Separation Protocol (LISP) architecture as described in [I-D.ietf-lisp-rfc6831bis]

2. Terminology

All of the terminology used in this document comes from [RFC9301] [I-D.ietf-lisp-rfc6831bis], [I-D.ietf-lisp-vpn] and [I-D.moreno-lisp-uberlay].

3. Scope of this document

This document covers the following aspects:

- * Deployments based on the procedures of [I-D.ietf-lisp-rfc6831bis].
- * When using a multicast based underlay, LISP sites can provide support to forward Layer-2 Broadcast, Unknown Unicast and Multicast (BUM) packets.
- * Layer3 routed multicast services (ASM, SSM and BiDir) are provided by such LISP sites.
- * Both IPv4 and IPv6 overlays are covered by this document. Similarly, both IPv4 and IPv6 underlays are covered.

4. Scope not covered by this document

This document does not cover the following aspects:

- * L3 routed Unicast forwarding or L2 forwarding between LISP sites.
- * This document does not address services implemented using underlays such as BIER.
- * Procedures and considerations required for migrating non-LISP based networks to LISP based networks.
- * Extranet Multicast (Route Leaking).
- * Signal Free LISP Multicast [I-D.ietf-lisp-rfc8378bis].

5. Industry segments/ use-cases covered

The deployment experiences outlined in this document capture learnings from various industry segments listed below (not exhaustive:)

- * Educational Institutions (e.g. Universities with multiple campuses, school districts)
- * Public Utilities like Airports, Stadiums and Ports
- * Hospitals and Healthcare providers
- * Enterprises including Financial Institutions spread across continents and large geographical regions
- * Technology vendors and factories
- * Events like Expos, Tradeshows, Sporting events

6. Advantages and Cost of using "PIM-over-PIM"

There are both advantages and costs in using a "PIM-over-PIM" design outlined in [I-D.ietf-lisp-rfc6831bis]:

- * PIM [RFC7761] is a well-understood and deployed protocol in many types of networks (Enterprise, Global Internet etc.).

- * For the specific case of deploying PIM in the overlay in a LISP network, merely encapsulating the PIM protocol packet into a Unicast LISP packet and directing it towards the xTR that is chosen as the Upstream Multicast NH worked very well.
- * Usage of PIM Join Attributes for LISP is a very useful method for the receiver ETR to signal underlay transport attributes to the ITR [RFC8059]. The motivations for doing so are explained in the later sections of this document.
- * The PIM neighborhood was not fully established as exchange of PIM hellos were considered chatty. An alternate method of achieving this could be to use [RFC9739]
- * A simple but powerful optimization was done to use only SSM in the underlay for supporting overlay Layer-3 multicast routing.

7. Underlay Deployment considerations

7.1. Head-End Replication

A small but significant subset of deployments have been observed using the Head-End Replication (HER). This is primarily done for low-volume multicast or for deployments where there are restrictions in implementations for supporting an underlay with native multicast.

Another category of the deployments were early adopters of the technology when the software releases were limited to unicast underlay.

The primary characteristic of such networks is the presence of a limited number of LISP sites in which receivers are present. Please note that this does not necessarily mean that only a limited number of hosts receive the multicast.

Since the ASICs that form the data plane have very efficient methods to replicate multicast packets to local receivers, any deployment that has a good localization of receivers to a limited number of LISP sites can still use a unicast underlay with high efficiency.

On the positive side, there are widely deployed mechanisms for both traffic-engineering (e.g. Load balancing) and fast convergence due to link/ node failures in unicast that can be reused for overlay routed multicast when using a unicast underlay.

Another very important use-case for considering a unicast underlay is to have migration done from (say) IPv4 to IPv6.

It is also possible to perform the Head-End replication on an ITR or PITR.

Although not covered in this document, the design of HER allows for the co-existence of RLOCs using both unicast and multicast underlays.

7.2. Native Multicast Underlay

Native multicast underlay presents notable advantages over Head-End Replication, particularly in network topologies where replication occurs at multiple nodes between the ETR and the ITR. Despite advancements in modern ASICs designed for high-performance multicast packet replication at the Head-End, bandwidth consumption remains a critical factor favoring the adoption of native multicast underlay.

In native multicast mode, there is a mapping between the overlay multicast group address and the underlay multicast group. This mapping must be consistent across network devices within a LISP domain to ensure uniformity. The simplest method involves a 1:1 mapping between the overlay LISP group address and the underlay multicast group address. To maintain uniqueness in this mapping process, implementations may incorporate additional parameters, such as the source IP address and LISP instance ID, providing sufficient entropy to ensure uniqueness across LISP instances.

There exists recent work like [I-D.ietf-lisp-group-mapping] that can be leveraged here.

Using native multicast in the underlay is used by the majority of the deployments known.

- * Underlays in most deployments were homogenous e.g. IPv6 Unicast based underlay.
- * Upgrading from one underlay to another is a process that requires a lot of planning. This is not covered in this document.

8. Layer-2 BUM overlay deployment considerations

There are three deployment options that can be considered here for deployment:

- * Ingress Replication: Each L2 BUM packet is replicated by the ITR so each ETR receives a copy of the L2 packet encapsulated as unicast in the underlay.
- * Use ASM underlay: Since any xTR does not know the list of ITRs that can potentially send L2 BUM packets, it subscribes to an underlay multicast group based on the L2 LISP instance. There can be a $m:n$ mapping of 'm' LISP instances to 'n' Underlay Multicast groups with $m > n$ or $m >> n$. We have also seen many customers use $n=1$.

- * Use BIDIR underlay: Since BIDIR is a commonly supported mode we can simply reduce the multicast state in the underlay to $O(n)$ instead of $O(n \times \text{no. of ITRs})$ by choosing BIDIR over ASM. This mode is particularly popular when IPv6 underlay is used as the forwarding path resources (e.g. TCAM entries) required to support IPv6 multicast routing is double that of IPv4.

8.1. RP for Layer-2 BUM Multicast Underlay

- * When using a Multicast underlay for L2 BUM, we use ASM based underlay or a BIDIR based underlay.
- * This can be achieved by having an m LISP L2 service instances are mapped to n multicast groups where $m > n$ or $m \gg n$ since the number of LISP L2 instances are larger than the number of designated multicast groups to carry BUM traffic.
- * Since this is done flexibly, heavy users of BUM can be allocated separate underlay groups for isolation.
- * One of the most critical design element is the choice of the RP Design. We have the following options:
 - Configuring static RPs: Use of anycast IP addresses with static RP is a popular choice observed in deployments.
 - Electing RPs through mechanisms like PIM-BSR [RFC5059] has been adopted by customers as well.
 - Discovering RPs via the Mapping system is not covered in this document.

9. Layer-3 overlay Multicast deployment considerations

9.1. Layer-3 Routed Any-Source Multicast (ASM) services

LISP overlays extend ASM to networks lacking native multicast support traditionally. Native multicast in the core boosts ASM resilience and optimizes traffic distribution.

Head-end replication requires tuning to avoid ITR overload with many receivers or high traffic. LISP overlays enable ASM resilience by rerouting around underlay failures dynamically.

ASM deployments scale receiver counts flexibly without requiring underlay redesigns. Troubleshooting ASM demands monitoring both LISP overlay and underlay states concurrently.

Pre-validating underlay multicast capabilities ensures reliable ASM performance consistently. Using native multicast complicates failure diagnosis despite enhancing overall resilience.

9.1.1. Layer-3 overlay ASM RP placement and redundancy

In a Layer-3 overlay, the placement of RPs is critical for ensuring robust multicast service delivery. Unlike traditional PIM-ASM, LISP multicast relies on static Rendezvous Point (RP) configurations due to the lack of support for dynamic RP discovery mechanisms like Auto-RP or Bootstrap Router (BSR).

Discovering RPs via the mapping system is possible but not covered in this document.

RPs can be positioned both inside and outside the LISP domain. The typical configuration involves static RP setup and redundancy through the Anycast RP concept, which allows multiple RPs to share the same IP address, providing load balancing and fault tolerance. This setup requires synchronization between RPs using the MSDP to exchange information about active sources.

In some deployments, RP placements are a combination of RPs placed inside together with RPs placed outside the LISP domains. This configuration leverages advanced MSDP peering or group mesh peering to enhance multicast service resilience and efficiency.

The RP placement significantly affects the convergence between the shared and source trees. It is essential that all xTRs within a given LISP instance use a consistent address scheme with identical mapping to ensure efficient multicast routing. The RP facilitates the initial setup of the shared tree, allowing sources and receivers to establish direct multicast data flows.

9.1.2. Optimisation for short-lived streams

When working with short lived streams (e.g. PA systems for airports) it was observed that using the shared tree was optimal. The cost of switching to the shortest-path tree can be avoided in such scenarios. However such choices are better done on a case-by-case basis e.g. based on the range of group addresses.

9.2. Layer-3 Routed SSM services

SSM services over a Layer-3 LISP domain connect multicast sources and receivers via the overlay. Receivers join source trees (S,G) by signalling IGMPv3, which then is transported as PIM packets over the LISP overlay. The typical SSM services would be represented by Financial Data, IPTV and Live streaming applications.

The traffic within a LISP domain, similar to the ASM would be subject to encapsulation and depending on the multicast mode it would be either head-end replication or native (overlay to underlay multicast mapping).

The sources and receivers can be connected to the LISP site or be located outside of the LISP domain. LISP overlay provides a resiliency by rerouting traffic dynamically.

SSM services eliminate RPs and shared trees, simplifying tree management. Direct (S,G) trees enhance scalability and reduce latency for one-to-many uses.

Receivers must support IGMPv3 (or MLDv2) to specify sources, avoiding ASM fallback. Replication strategies need tuning to balance ITR load and underlay bandwidth.

10. Mobility considerations for LISP multicast

TBD

11. Enabling Multicast flows for multiple tenants and multiple site overlays

11.1. Background reading

[I-D.moreno-lisp-uberlay] provides methods to operate and interconnect multiple LISP sites using Border xTR nodes.

[I-D.ietf-lisp-vpn] describes the use of the LISP to create Virtual Private Networks (VPNs). LISP is used to provide segmentation in both the LISP data plane and control plane.

Using the procedures and construct of the above two references, we have been able to build and deploy solutions that cater to multiple tenants connected to multiple LISP sites spread across multiple site overlays.

11.2. LISP Uberlay deployment considerations

One of the primary deployment use cases involves delivering multicast services across multiple site overlays or RLOC spaces [I-D.moreno-lisp-uberlay]. There are several methods for routing multicast packets when sources and receivers are connected to LISP site overlays that are connected through VPNs. Two most common methods are given below:

- * Forwarding the traffic natively, without any encapsulation, which typically results in extending the Forwarding Context [I-D.ietf-lisp-vpn] segmentation beyond a specific LISP site overlay.
- * Implementing an overlay across the uberlay network.

Choosing between these options has significant design implications for both unicast and multicast flows:

In the native forwarding approach, traffic leaving a site overlay is decapsulated at the border xTRs and placed into the appropriate VRF corresponding to the Forwarding Context. This scenario creates considerable overhead in deploying multicast configurations across multiple site overlays, as each LISP Forwarding Context must be mapped to an individual VRF in uberlay.

Implementing an overlay over the LISP uberlay network offers advantages by extending the LISP Forwarding Context between different LISP domains. In this case, Border xTRs in each LISP domain are responsible for decapsulating and re-encapsulating traffic between different site overlays. This can be achieved by using disjoint underlay multicast groups in the different site overlays/ uberlay. PIM can be leveraged for signaling the different underlay group mappings for the site-overlays and uberlay. [RFC8059]

12. Extranet Multicast (Route leaking)

This feature is beyond the scope of this document.

13. IANA Considerations

This memo includes no request to IANA.

14. Security Considerations

This informational document does not introduce any new security considerations.

15. References

15.1. Normative References

[I-D.ietf-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-ietf-lisp-rfc6831bis-01, 29 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-rfc6831bis-01>>.

[I-D.ietf-lisp-rfc8378bis]

Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc8378bis-03, 25 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-rfc8378bis-03>>.

[RFC5059] Bhaskar, N., Gall, A., Lingard, J., and S. Venaas, "Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)", RFC 5059, DOI 10.17487/RFC5059, January 2008, <<https://www.rfc-editor.org/info/rfc5059>>.

[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>>.

[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>>.

15.2. Informative References

[I-D.ietf-lisp-group-mapping]

Govindan, V. P., Farinacci, D., Kuppusami, A., and S. Venaas, "LISP Multicast Overlay Group to Underlay RLOC Mappings", Work in Progress, Internet-Draft, draft-ietf-lisp-group-mapping-01, 11 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-group-mapping-01>>.

[I-D.ietf-lisp-vpn]

Moreno, V. and D. Farinacci, "LISP Virtual Private Networks (VPNs)", Work in Progress, Internet-Draft, draft-ietf-lisp-vpn-12, 19 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-vpn-12>>.

[I-D.moreno-lisp-uberlay]

Moreno, V., Farinacci, D., Rodriguez-Natal, A., Portoles-Comeras, M., Maino, F., and S. Hooda, "Uberlay Interconnection of Multiple LISP overlays", Work in Progress, Internet-Draft, draft-moreno-lisp-uberlay-06, 28 September 2022, <<https://datatracker.ietf.org/doc/html/draft-moreno-lisp-uberlay-06>>.

[RFC9739] Bidgoli, H., Ed., Venaas, S., Mishra, M., Zhang, Z., and M. McBride, "Protocol Independent Multicast Light (PIM Light)", RFC 9739, DOI 10.17487/RFC9739, March 2025, <<https://www.rfc-editor.org/info/rfc9739>>.

Acknowledgements

A very special and sincere thanks to Dino Farinacci for his extensive and insightful comments for improving the content and the readability of this document.

The authors would like to acknowledge Stig Venaas for his review. Many individuals also contributed to the discussions for the material of this draft including Arunkumar Nandakumar, Aswin Kuppusami, Rajavel Ganesamoorthy, Sankar S and Sundara Moorthy. All contributions are gratefully acknowledged.

Contributors

TBD

Authors' Addresses

Vengada Prasad Govindan
Cisco
Email: venggovi@cisco.com

Marcin Hamroz
Cisco
Email: mhamroz@cisco.com

Jaroslav Gawron
Cisco
Email: jagawron@cisco.com