

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 29 September 2025

P. Jain
MIPS
V. Moreno
Google
S. Hooda
Cisco Systems
28 March 2025

LISP Site External Connectivity
draft-ietf-lisp-site-external-connectivity-02

Abstract

This draft defines how to register/retrieve pETR mapping information in LISP when the destination is not registered/known to the local site and its mapping system (e.g. the destination is an internet or external site destination).

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 BCP 14 RFC 2119 [RFC2119] RFC 8174 [RFC8174].

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 September 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. Definition of Terms	3
3. pETR Registration	3
4. pETR Request/Subscription	3
5. pETR Notification/Publication	4
6. pETR Resolution	4
7. pETR Map-Cache Update	5
8. IANA Considerations	5
9. Security Considerations	5
10. Acknowledgements	5
11. Normative Reference	5
Authors' Addresses	6

1. Introduction

The LISP architecture and protocol [RFC9301] introduces two types of replies to a map-request sent by an ITR:

- when the EID is known or registered to the mapping system, a regular map-reply with mapping information is sent, or
- when the EID is unknown or known but not registered, a negative-map-reply (NMR) is sent.

Currently the NMR does not convey pETR RLOC-set information to specify where the ITR should send the packet.

This document describes how to use the LISP messages to register and provide pETR RLOC-set information for destinations which are EIDs not registered with the Mapping System, or simply are "not known" to be an existing EID. These destinations could be the destinations which are outside of the LISP site belonging to non-LISP domains, hence are not registered with the LISP Mapping System.

The reachability of these destinations can be provided either by configuring pETR information directly into the Mapping System, or by the registration done by certain pETRs. The pETR registration is

specifically useful when the traffic to these external destinations needs to be sent encapsulated to a preferred pETR/gateway chosen dynamically. This mechanism also helps to achieve faster convergence.

This document also specifies the structure of the map-reply containing pETR RLOC-set information.

2. Definition of Terms

Same as defined in [RFC9300] and [RFC9301].

3. pETR Registration

pETRs having external or internet connectivity MAY register the pETR with the mapping system per Instance ID (IID) of the VPN [I-D.ietf-lisp-vpn]. The pETR Map-Register procedures and record format are the same as in [RFC9301] with the following contents:

- An "EID-Prefix" as configurable "Distinguished Name" defined, agreed upon and 'distinguished' by the mechanism according to [I-D.ietf-lisp-name-encoding].
- RLOC-set for pETR information.
- Each locator in the RLOC-set MAY be encoded as per [I-D.ietf-lisp-vpn]. This enables dynamic external connectivity in VPN environments.
- Additional information MAY be encoded in vendor specific LCAF type [RFC9306] about the registering pETR such as its performance matrix, location, resource availability for the Mapping System to make preference decision.

4. pETR Request/Subscription

The Map-Request procedures and record format are the same as in [RFC9301] with the following contents:

- An "EID-Prefix" MAY be configurable "Distinguished Name" defined, agreed upon and 'distinguished' by the mechanism according to [I-D.ietf-lisp-name-encoding].
- N-bit MAY be set as per [RFC9437]
- Additional information MAY be encoded in vendor specific LCAF type [RFC9306] about the requesting ITR such as location, performance matrix to make preference decision.

5. pETR Notification/Publication

When pETR registers or updates the pETR registration with the mapping system, mapping system sends Map-Notify.

With lisp-pubsub [RFC9437], N-bit SHOULD be set in the Map-Request from ITR. Then, whenever pETR gets updated in the mapping system, mapping system sends Map-Notify messages to update subscribed ITRs as per procedures in [RFC9437].

When lisp-pubsub [RFC9437] is not available in the mapping system, then shorter TTL MAY be used for updating the map-caches at ITR as per procedures in [RFC9301].

The pETR Map-Notify procedures and record format are the same as in [RFC9301] with the following contents:

- An "EID-Prefix" as configurable "Distinguished Name" defined, agreed upon and 'distinguished' by the mechanism according to [I-D.ietf-lisp-name-encoding].
- RLOC-set for pETR information.
- Each locator in the RLOC-set MAY be encoded as per [I-D.ietf-lisp-vpn]. This enables dynamic external connectivity in VPN environments.
- TTL MAY be shorter than regular.
- Additional information MAY be encoded in vendor specific LCAF type [RFC9306] about the registering pETR such as its performance matrix, location, resource availability to communicate preference decision.

6. pETR Resolution

When the Map-Server (or ETR) determines that the requested destination is external or unknown to the mapping system, it sends a Map-Reply containing the pETR information. The Map-Reply procedures and record format are the same as described in the Map-Server processing section of [RFC9301]. This Map-Reply has the following content (as defined per regular map-reply and negative-map-reply in [RFC9301]):

- An EID-Prefix calculated as non-LISP "hole" per the procedures in [RFC9301] for negative map-reply.
- RLOC count MUST be non-zero.

- Each locator in the RLOC-set MAY be encoded as per [I-D.ietf-lisp-vpn]. This enables dynamic external connectivity in VPN environments.
- TTL MAY be shorter than regular map-reply.
- Additional information MAY be encoded in vendor specific LCAF type [RFC9306] about the mapping such as whether the mapping is based upon policy or registration.

7. pETR Map-Cache Update

On receiving pETR Map-Notify/Map-Reply from the mapping system, ITR MAY install/update map-cache and encapsulate the packets to pETR RLOCs as per [RFC9300] for the destinations not known to the mapping system as EIDs or known EIDs not registered to the mapping system. This can be implemented as follows,

- ITR SHOULD configure the known EID-blocks in its map-cache to always generate Map-Request for known EIDs. These would be more specific map-cache entries than “hole” EID-Prefix or default map-cache entries.
- On receiving pETR Map-Notify/Map-Reply, ITR MAY install/update following map-cache entries with pETR RLOCs,
 - “hole” prefix [RFC9301] map-cache entry to encapsulate the packets to pETR RLOCs
 - default map-cache entry to encapsulate the packets to pETR RLOCs.

8. IANA Considerations

No IANA considerations apply to this document.

9. Security Considerations

There are no additional security considerations except what already discussed in [RFC9301].

10. Acknowledgements

The authors would like to thank Fabio Maino, Dino Farinacci and Padma Pillay-Esnault for the suggestions and review of this document.

11. Normative Reference

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

Farinacci, D., "LISP Distinguished Name Encoding", Work in Progress, Internet-Draft, draft-ietf-lisp-name-encoding-14, 16 August 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-name-encoding-14>>.

`[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>>.

[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>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[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>>.

[RFC9306] Rodriguez-Natal, A., Ermagan, V., Smirnov, A., Ashtaputre, V., and D. Farinacci, "Vendor-Specific LISP Canonical Address Format (LCAF)", RFC 9306, DOI 10.17487/RFC9306, October 2022, <<https://www.rfc-editor.org/info/rfc9306>>.

[RFC9437] Rodriguez-Natal, A., Ermagan, V., Cabellos, A., Barkai, S., and M. Boucadair, "Publish/Subscribe Functionality for the Locator/ID Separation Protocol (LISP)", RFC 9437, DOI 10.17487/RFC9437, August 2023, <<https://www.rfc-editor.org/info/rfc9437>>.

Authors' Addresses

Prakash Jain
MIPS
San Jose, CA
United States of America
Email: prjain@mips.com

Victor Moreno
Google
Mountainview, CA
United States of America
Email: vimoreno@google.com

Sanjay Hooda
Cisco Systems
San Jose, CA
United States of America
Email: shooda@cisco.com