

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: March 28, 2019

D. Farinacci
lispers.net
E. Nordmark
Zededa
September 24, 2018

LISP Control-Plane ECDSA Authentication and Authorization
draft-ietf-lisp-ecdsa-auth-00

Abstract

This draft describes how LISP control-plane messages can be individually authenticated and authorized without a priori shared-key configuration. Public-key cryptography is used with no new PKI infrastructure required.

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 March 28, 2019.

Copyright Notice

Copyright (c) 2018 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	4
3. Overview	5
4. Public-Key Hash	6
5. Hash-EID Mapping Entry	7
6. Hash-EID Structure	7
7. Keys and Signatures	8
8. Signed Map-Register Encoding	8
9. Signed Map-Request Encoding	9
10. Signed Map-Notify Encoding	10
11. Other Uses	10
12. EID Authorization	11
13. Security Considerations	13
14. IANA Considerations	13
15. References	13
15.1. Normative References	13
15.2. Informative References	15
Appendix A. Acknowledgments	15
Appendix B. Document Change Log	16
B.1. Changes to draft-ietf-lisp-ecdsa-auth-00	16
B.2. Changes to draft-farinacci-lisp-ecdsa-auth-03	16
B.3. Changes to draft-farinacci-lisp-ecdsa-auth-02	16
B.4. Changes to draft-farinacci-lisp-ecdsa-auth-01	16
B.5. Changes to draft-farinacci-lisp-ecdsa-auth-00	16
Authors' Addresses	17

1. Introduction

The LISP architecture and protocols [RFC6830] introduces two new numbering spaces, Endpoint Identifiers (EIDs) and Routing Locators (RLOCs) which provide an architecture to build overlays on top of the underlying Internet. Mapping EIDs to RLOC-sets is accomplished with a Mapping Database System. EIDs and RLOCs come in many forms than just IP addresses, using a general syntax that includes Address Family Identifier (AFI) [RFC1700]. Not only IP addresses, but other addressing information have privacy requirements. Access to private information is granted only to those who are authorized and authenticated. Using asymmetric keying with public key cryptography enforces authentication for entities that read from and write to the mapping system. The proposal described in this document takes advantage of the latest in Elliptic Curve Cryptography.

In this proposal the EID is derived from a public key, and the corresponding private key is used to authenticate and authorize Map-Register messages. Thus only the owner of the corresponding private key can create and update mapping entries from the EID. Furthermore,

the same approach is used to authenticate Map-Request messages. This in combination with the mapping database containing authorization information for Map-Requests is used to restrict which EIDs can lookup up the RLOCs for another EID.

This specification introduces how to use the Distinguished-Name AFI [AFI] and the [RFC8060] LCAF JSON Type to encode public keys and signatures in the LISP mapping database. The information in the mapping database is used to verify cryptographic signatures in LISP control-plane messages such as the Map-Request and Map-Register.

2. Definition of Terms

Crypto-EID: is an IPv6 EID where part of the EID includes a hash value of a public-key. An IPv6 EID is a Crypto-EID when the Map-Server is configured with an Crypto-EID Prefix that matches the IPv6 EID.

Crypto-EID Hash Length: is the number of low-order bits in a Crypto-EID which make up the hash of a public-key. The hash length is determined by the Map-Server when it is configured with a Crypto-EID Prefix.

Crypto-EID Prefix: is a configuration parameter on the Map-Server that indicates which IPv6 EIDs are Crypto-EIDs and what is the Crypto-EID Hash Length for the IPv6 EID. This can be different for different LISP Instance-IDs.

Hash-EID: is a distinguished name EID-record stored in the mapping database. The EID format is 'hash-<pubkey-hash>'. When a key-pair is generated for an endpoint, the produced private-key does not leave the xTR that will register the Crypto-EID. A hash of the public-key is used to produce a Crypto-EID and a Hash-EID. The Crypto-EID is assigned to the endpoint and the xTR that supports the LISP-site registers the Crypto-EID. Another entity registers the Hash-EID mapping with the public-key as an RLOC-record.

Public-Key RLOC: is a JSON string that encodes a public-key as an RLOC-record for a Hash-EID mapping entry. The format of the JSON string is '{ "public-key" : "<pubkey>" }'.

Control-Plane Signature: a Map-Request or Map-Register sender signs the message with its private key. The format of the signature is a JSON string that includes sender information and the signature value. The JSON string is included in Map-Request and Map-Register messages.

Signature-ID: is a Crypto-EID used for a Control-Plane signature to register or request any type of EID. The Signature-ID is included with the JSON-encoded signature in Map-Request and Map-Register messages.

Multi-Signatures: multiple signatures are used in LISP when an entity allows and authorized another entity to register an EID. There can be more than one authorizing entities that allow a registering entity to register an EID. The authorizing entities sign their own RLOC-records that are registered and merged into the registering entity's Hash-EID public-key mapping. And when

the registering entity registers the EID, all authorizing entity signatures must be verified by the Map-Server before the EID is accepted.

3. Overview

LISP already has several message authentication mechanisms. They can be found in [I-D.ietf-lisp-rfc6833bis], [I-D.ietf-lisp-sec], and [RFC8061]. The mechanisms in this draft are providing a more granular level of authentication as well as a simpler way to manage keys and passwords.

A client of the mapping system can be authenticated using public-key cryptography. The client is required to have a private/public key-pair where it uses the private-key to sign Map-Requests and Map-Registers. The server, or the LISP entity, that processes Map-Requests and Map-Registers uses the public-key to verify signatures.

The following describes how the mapping system is used to implement the public-key crypto system:

1. An entity registers Hash-EID to Public-Key RLOC mappings. A third-party entity that provides a service can register or the client itself can register.
 2. Anyone can lookup the Hash-EID mappings. These mappings are not usually authenticated with the mechanisms in this draft but use the shared configured password mechanisms from [I-D.ietf-lisp-rfc6833bis] that provide group level authentication.
 3. When a Crypto-EID, or any EID type, is registered to the mapping system, a signature is included in the Map-Register message. When a non-Crypto-EID is registered a Signature-ID is also included in the Map-Register message.
 4. The Map-Server processes the registration by constructing the Hash-EID from the registered Crypto-EID, looks up the Hash-EID in the mapping system, obtains the public-key from the RLOC-record and verifies the signature. If Hash-EID lookup fails or the signature verification fails, the Map-Register is not accepted.
 5. When a Crypto-EID, or any EID type, is looked up in the mapping system, a signature is included with a Signature-ID in the Map-Request message.
 6. The Map-Server processes the request for a Crypto-EID by constructing the Hash-EID from the Signature-ID included in the Map-Request. The signer-ID is a Crypto-EID that accompanies a signature in the Map-Request. The Hash-EID is looked up in the mapping system, obtains the public-key from the RLOC-record and verifies the Map-Request signature. If the Hash-EID lookup fails or the signature verification fails, the Map-Request is not accepted and a Negative Map-Reply is sent back with an action of "auth-failure".
4. Public-Key Hash

When a private/public key-pair is created for a node, its IPv6 EID is pre-determined based on the public key generated. Note if the key-pair is compromised or is changed for the node, a new IPv6 EID is assigned for the node.

The sha256 [RFC6234] hex digest function is used to compute the hash. The hash is run over the following hex byte string:

<iid><prefix><pubkey>

Where each field is defined to be:

<iid>: is a 4-byte (leading zeroes filled) binary value of the Instance-ID the EID will be registered with in the mapping database. For example, if the instance-id is 171, then the 4-byte value is 0x000000ab.

<prefix>: is a variable length IPv6 prefix in binary format (with no colons) and IS quad-nibble zero-filled. The length of the prefix is 128 minus the Crypto-EID hash bit length. For example, if the prefix is 2001:5:3::/48, then the 6 byte value is 0x200100050003.

<pubkey>: is a DER [RFC7468] encoded public-key.

The public-key hash is used to construct the Crypto-EID and Hash-EID.

5. Hash-EID Mapping Entry

A Hash-EID is formatted in an EID-record as a Distinguished-Name AFI as specified in [I-D.farinacci-lisp-name-encoding]. The format of the string is:

EID-record: 'hash-<hash-eid>'

Where <hash-eid> is a public-key hash as described in Section 4. The RLOC-record to encode and store the public-key is in LCAF JSON Type format of the form:

RLOC-record: '{ "public-key" : "<pubkey-base64>" }'

Where <pubkey-base64> is a base64 [RFC4648] encoding of the public-key generated for the system that is assigned the Hash-EID.

6. Hash-EID Structure

Since the Hash-EID is formatted as a distinguished-name AFI, the format of the <hash-eid> for EID 'hash-<hash-eid>' needs to be specified. The format will be an IPv6 address [RFC3513] where colons are used between quad-nibble characters when the hash bit length is a multiple of 4. And when the hash bit length is not a multiple of 4 but a multiple of 2, a leading 2 character nibble-pair is present. Here are some examples for different hash bit lengths:

Crypto-EID: 2001:5::1111:2222:3333:4444, hash length 64:
Hash-EID is: 'hash-1111:2222:3333:4444'

Crypto-EID: 2001:5::11:22:33:44, hash length 64:
Hash-EID is: 'hash-0011:0022:0033:0044'

Crypto-EID: 2001:5:aaaa:bbbb:1111:2222:3333:4444, hash length 80:
Hash-EID is: 'hash-bbbb:1111:2222:3333:4444'

Crypto-EID: 2001:5:aaaa:bbbb:1111:2222:3333:4444, hash length 72:
Hash-EID is: 'hash-bb:1111:2222:3333:4444'

Crypto-EID: 2001:5:aaaa:bbbb:1111:22:33:4444, hash length 72:
Hash-EID is: 'hash-bb:1111:0022:0033:4444'

Note when leading zeroes exist in a IPv6 encoded quad between colons, the zeros are included in the quad for the Hash-EID string.

The entity that creates the hash, the entity that registers the Crypto-EID and the Map-Server that uses the hash for Hash-EID lookups MUST agree on the hash bit length.

7. Keys and Signatures

Key generation, message authentication with digital signatures, and signature verification will use the Elliptic Curve Digital Signature Algorithm or ECDSA [X9.62]. For key generation curve 'NIST256p' is used and recommended.

Signatures are computed over signature data that depends on the type of LISP message sent. See Section 8 and Section 9 for each message type. The signature data is passed through a sha256 hash function before it is signed or verified.

8. Signed Map-Register Encoding

When a ETR registers its Crypto-EID or any EID type to the mapping system, it builds a LISP Map-Register message. The mapping includes an EID-record which encodes the Crypto-EID, or any EID type, and an RLOC-set. One of the RLOC-records in the RLOC-set includes the the ETR's signature and signature-ID. The RLOC-record is formatted with a LCAF JSON Type, in the following format:

```
{ "signature" : "<signature-base64>", "signature-id" : "<signer-id>" }
```

Where <signature-base64> is a base64 [RFC4648] encoded string over the following ascii [RFC0020] string signature data:

[<iid>]<crypto-eid>

Where <iid> is the decimal value of the instance-ID the Crypto-EID is registering to and the <crypto-eid> is in the form of [RFC3513] where quad-nibbles between colons ARE NOT zero-filled.

The Map-Server that process an EID-record with a Crypto-EID and a RLOC-record with a signature extracts the public-key hash value from the Crypto-EID to build a Hash-EID. The Map-Server looks up the Hash-EID in the mapping system to obtain the public-key RLOC-record. The Map-Server verifies the signature over the signature data to determine if it should accept the EID-record registration.

9. Signed Map-Request Encoding

When an xTR (an ITR, PITR, or RTR), sends a Map-Request to the mapping system to request the RLOC-set for a Crypto-EID, it signs the Map-Request so it can authenticate itself to the Map-Server the Crypto-EID is registered to. The Map-Request target-EID field will contain the Crypto-EID and the source-EID field will contain an LCAF JSON Type string with the following signature information:

```
{
  "source-eid" : "<seid>",
  "signature-id" : "<signer-id>",
  "signature" : "<signature-base64>"
}
```

Where <signer-id> is an IPv6 encoded string according to [RFC3513] where quad-nibbles between colons ARE NOT zero-filled. The <seid> is the source EID from the data packet that is invoking the Map-Request or the entire key/value pair for "source-eid" can be excluded when a data packet did not invoke the Map-Request (i.e. lig or an API request). The <signer-id> is the IPv6 Crypto-EID of the xTR that is providing the Map-Request signature.

The signature string <signature-base64> is a base64 [RFC4648] encoded string over the following signature data:

<nonce><source-eid><crypto-eid>

Where <nonce> is a hex string [RFC0020] of the nonce used in the Map-Request and the <source-eid> and <crypto-eid> are hex strings [RFC0020] of an IPv6 address in the form of [RFC3513] where quad-nibbles between colons ARE NOT zero-filled. When <seid> is not included in the Map-Request, string "0::0" is used for <source-eid>.

10. Signed Map-Notify Encoding

When a Map-Server originates a Map-Notify message either as an acknowledgment to a Map-Register message, as a solicited [I-D.ietf-lisp-pubsub] notification, or an unsolicited [RFC8378] notification, the receiver of the Map-Notify can verify the message is from an authenticated Map-Server.

An RLOC-record similar to the one used to sign Map-Register messages is used to sign the Map-Notify message:

```
{ "signature" : "<signature-base64>", "signature-id" : "<signer-id>" }
```

Where the "signature-id" is an IPv6 crypto-EID used by the Map-Server to sign the RLOC-record. The signature data and the encoding format of the signature is the same as for a Map-Register message. See details in Section 8.

A receiver of a Map-Notify message will lookup the signature-id in the mapping system to obtain a public-key to verify the signature. The Map-Notify is accepted only if the verification is successful.

11. Other Uses

The mechanisms described within this document can be used to sign other types of LISP messages. And for further study is how to use these mechanisms to sign LISP encapsulated data packets in a compressed manner to reduce data packet header overhead.

In addition to authenticating other types of LISP messages, other types of EID-records can be encoded as well and is not limited to IPv6 EIDs. It is possible for a LISP xTR to register and request non IPv6 EIDs but use IPv6 Crypto-EIDs for the sole purpose of signing and verifying EID-records.

Examples of other EID types that can be authenticated in Map-Request and Map-Register messages are:

EID-Type	Format Definition
IPv4 address prefixes	[RFC1123]
Distinguished-Names	[I-D.farinacci-lisp-name-encoding]
Geo-Coordinates	[I-D.farinacci-lisp-geo]
LCAF defined EIDs	[RFC8060]

12. EID Authorization

When a Crypto-EID is being used for IPv6 communication, it is implicit that the owner has the right to use the EID since it was generated from the key-pair provisioned for the owner. For other EID types that are not directly associated with signature keys, they must be validated for use by the mapping system they are registered to. This policy information for the mapping system must be configured in the Map-Servers the EID owner registers to or a signed authorization provided by a third-party entity.

To achieve signed authorization, an entity that allows another entity to register an EID, must authorize the registering entity. It does so by adding RLOC-records to the registering entity's Hash-EID public-key mapping. The format of the RLOC-record is a JSON encoded record as follows:

```
{
  "allow-eid" : "<eid>",
  "signature-id" : "<signer-id>",
  "signature" : "<signature-base64>"
}
```

The format of the <signer-id> and <signature-base64> values are the same as described in Section 8. The <eid> value is in the same string format as the signature data described in Section 8. For other non-IPv6 EID types, the conventions in [RFC8060] are used. In all cases, the string encoding format of instance-ID '[<iid>]' prepended is to the EID string.

This entry is added to the RLOC-set of the registering entity's Hash-EID 'hash-<hash>' registration. The authorizing entity does signs the Map-Register and sends it with merge-semantic. The Map-Server accepts the registration after the signature is verified and merges the RLOC-record into the existing RLOC-set. The 'signature' is optional and when not included means the authorizing entity has not

yet allowed the registering entity to register the EID <eid>. Note multiple entities can register RLOC-records with the same <eid> meaning that signature verification for all of them is required before the Map-Server accepts the registration.

When the Map-Server receives a Map-Register for <eid>, it looks up 'hash-<hash>' EID in the mapping system. If not found, the Map-Register EID-record is not processed and the next EID-record is retrieved from the Map-Register message, if it exists. If the Hash-EID entry is found, the registering entity's signature is verified first. If the verification fails, the Map-Register EID-record is not accepted. Otherwise, a search for the RLOC-set is done to look for all matches of the EID being registered with <eid>, for those entries found, if any of them do not have a "signature" JSON item, the EID-record is not accepted. Otherwise, the signature-id is looked up in the mapping system to retrieve the public-key of the authorizing entity. If the verification is successful, then a lookup for the next RLOC-record signature-id is done. Only when all signature verifications are verified, the Map-Register EID-record is accepted.

The Map-Server should reject an RLOC-record with a signature-id that contains the Hash-EID of the entry disallowing a registering entity to self authorize itself.

Here is an example of a Hash-EID mapping stored in the mapping system:

EID-record: [1000]'hash-1111:2222:3333:4444', RLOC-Set (count is 4):

```
RLOC-record: { "public-key" : "<pubkey-base64>" }
RLOC-record: { "allow-eid" : "[1000]1.1.1.1/32", "signature" : "<sig>",
               "signature-id" : "[1000]2001:5:3::1111" }
RLOC-record: { "allow-eid" : "[1000]1.1.1.1/32", "signature" : "<sig>",
               "signature-id" : "[1000]2001:5:3::2222" }
RLOC-record: { "allow-eid" : "37-16-46-N-121-52-4-W",
               "signature-id" : "[1000]2001:5:3::5555" }
```

This mapping stores the public-key of the registering entity with Hash-EID 1111:2222:3333:4444. The registering entry registered this RLOC-record. There are two authorizing entities, :1111 and :2222, who allow it to register IPv4 EID 1.1.1.1/32. They each registered their respective RLOC-records. And a third authorizing entity :5555 that registers an RLOC-record that has not yet authorized the registering entity to register Geo-Coordinate 37-16-46-N-121-52-4-W. Note the mapping and the signature-IDs are all within the context of instance-ID 1000.

13. Security Considerations

The mechanisms within this specification are intentionally using accepted practices and state of the art public-key cryptography.

Crypto-EIDs can be made private when control messages are encrypted, for instance, using [RFC8061].

The topological or physical location of a Crypto-EID is only available to the other Crypto-EIDs that register in the same LISP Instance-ID and have their corresponding Hash-EIDs registered.

This draft doesn't address reply attacks directly. If a man-in-the-middle captures Map-Register messages, it could send such captured packets at a later time which contains signatures of the source. In which case, the Map-Server verifies the signature as good and interprets the contents to be valid where in fact the contents can contain old mapping information. This problem can be solved by encrypting the contents of Map-Registers using a third-party protocol like DTLS [RFC6347] or LISP-Crypto [RFC8061] directly by encapsulating Map-Registers in LISP data packets (using port 4341).

Map-Reply message signatures and authentication are not in scope for this document. This document focuses on authentication between xTRs and mapping system components. Map-Reply authentication, which is performed between xTRs is described in [I-D.ietf-lisp-sec].

14. IANA Considerations

Since there are no new packet formats introduced for the functionality in this specification, there are no specific requests for IANA.

15. References

15.1. Normative References

- [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>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<https://www.rfc-editor.org/info/rfc1123>>.

- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, DOI 10.17487/RFC1700, October 1994, <<https://www.rfc-editor.org/info/rfc1700>>.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, DOI 10.17487/RFC3513, April 2003, <<https://www.rfc-editor.org/info/rfc3513>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [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>>.
- [RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, DOI 10.17487/RFC6833, January 2013, <<https://www.rfc-editor.org/info/rfc6833>>.
- [RFC7468] Josefsson, S. and S. Leonard, "Textual Encodings of PKIX, PKCS, and CMS Structures", RFC 7468, DOI 10.17487/RFC7468, April 2015, <<https://www.rfc-editor.org/info/rfc7468>>.
- [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>>.
- [RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.

[RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <<https://www.rfc-editor.org/info/rfc8378>>.

15.2. Informative References

[AFI] IANA, "Address Family Identifier (AFIs)", ADDRESS FAMILY NUMBERS <http://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml?>, February 2007.

[I-D.farinacci-lisp-geo] Farinacci, D., "LISP Geo-Coordinate Use-Cases", draft-farinacci-lisp-geo-05 (work in progress), April 2018.

[I-D.farinacci-lisp-name-encoding] Farinacci, D., "LISP Distinguished Name Encoding", draft-farinacci-lisp-name-encoding-06 (work in progress), September 2018.

[I-D.ietf-lisp-pubsub] Rodriguez-Natal, A., Ermagan, V., Leong, J., Maino, F., Cabellos-Aparicio, A., Barkai, S., Farinacci, D., Boucadair, M., Jacquenet, C., and S. Secci, "Publish/Subscribe Functionality for LISP", draft-ietf-lisp-pubsub-00 (work in progress), April 2018.

[I-D.ietf-lisp-rfc6833bis] Fuller, V., Farinacci, D., and A. Cabellos-Aparicio, "Locator/ID Separation Protocol (LISP) Control-Plane", draft-ietf-lisp-rfc6833bis-15 (work in progress), September 2018.

[I-D.ietf-lisp-sec] Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D. Saucez, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-15 (work in progress), April 2018.

[X9.62] American National Standards Institute, "Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA)", NIST ANSI X9.62-2005, November 2005.

Appendix A. Acknowledgments

A special thanks goes to Sameer Merchant and Colin Cantrell for their ideas and technical contributions to the ideas in this draft.

Appendix B. Document Change Log

[RFC Editor: Please delete this section on publication as RFC.]

- B.1. Changes to draft-ietf-lisp-ecdsa-auth-00
 - o Posted mid-September 2018.
 - o Make draft-farinacci-lisp-ecdsa-auth-03 a LISP working group document.
- B.2. Changes to draft-farinacci-lisp-ecdsa-auth-03
 - o Posted September 2018.
 - o Change all occurrences of signature-EID to signature-ID.
 - o Document how Map-Servers sign Map-Notify messages so they can be verified by xTRs.
 - o Add multi-signatures to mappings so a 3rd-party can allow an entity to register any type of EID.
- B.3. Changes to draft-farinacci-lisp-ecdsa-auth-02
 - o Draft posted April 2018.
 - o Generalize text to allow Map-Requesting and Map-Registering for any EID type with a proper signature-EID and signature encoded together.
- B.4. Changes to draft-farinacci-lisp-ecdsa-auth-01
 - o Draft posted October 2017.
 - o Make it more clear what values and format the EID hash is run over.
 - o Update references to newer RFCs and Internet Drafts.
- B.5. Changes to draft-farinacci-lisp-ecdsa-auth-00
 - o Initial draft posted July 2017.

Authors' Addresses

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Erik Nordmark
Zededa
Santa Clara, CA
USA

Email: erik@zededa.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 22 August 2024

D. Farinacci
lispers.net
E. Nordmark
Zededa
19 February 2024

LISP Control-Plane ECDSA Authentication and Authorization
draft-ietf-lisp-ecdsa-auth-12

Abstract

This draft describes how LISP control-plane messages can be individually authenticated and authorized without a a priori shared-key configuration. Public-key cryptography is used with no new PKI infrastructure required.

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 22 August 2024.

Copyright Notice

Copyright (c) 2024 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. Overview	4
4. Public-Key Hash	5
5. Hash-EID Mapping Entry	6
6. Hash-EID Structure	6
7. Keys and Signatures	7
8. Signed Map-Register Encoding	7
9. Signed Map-Request Encoding	8
10. Signed Map-Notify Encoding	9
11. Other Uses	9
12. EID Authorization	10
13. Security Considerations	12
14. IANA Considerations	13
15. References	13
15.1. Normative References	13
15.2. Informative References	14
Appendix A. Acknowledgments	15
Appendix B. Document Change Log	15
B.1. Changes to draft-ietf-lisp-ecdsa-auth-12	15
B.2. Changes to draft-ietf-lisp-ecdsa-auth-11	15
B.3. Changes to draft-ietf-lisp-ecdsa-auth-10	15
B.4. Changes to draft-ietf-lisp-ecdsa-auth-09	16
B.5. Changes to draft-ietf-lisp-ecdsa-auth-08	16
B.6. Changes to draft-ietf-lisp-ecdsa-auth-07	16
B.7. Changes to draft-ietf-lisp-ecdsa-auth-06	16
B.8. Changes to draft-ietf-lisp-ecdsa-auth-05	16
B.9. Changes to draft-ietf-lisp-ecdsa-auth-04	16
B.10. Changes to draft-ietf-lisp-ecdsa-auth-03	16
B.11. Changes to draft-ietf-lisp-ecdsa-auth-02	16
B.12. Changes to draft-ietf-lisp-ecdsa-auth-01	17
B.13. Changes to draft-ietf-lisp-ecdsa-auth-00	17
B.14. Changes to draft-farinacci-lisp-ecdsa-auth-03	17
B.15. Changes to draft-farinacci-lisp-ecdsa-auth-02	17
B.16. Changes to draft-farinacci-lisp-ecdsa-auth-01	17
B.17. Changes to draft-farinacci-lisp-ecdsa-auth-00	17
Authors' Addresses	17

1. Introduction

The LISP architecture and protocols [RFC9300] introduces two new numbering spaces, Endpoint Identifiers (EIDs) and Routing Locators (RLOCs) which provide an architecture to build overlays on top of the underlying Internet. Mapping EIDs to RLOC-sets is accomplished with a Mapping Database System. EIDs and RLOCs come in many forms than just IP addresses, using a general syntax that includes Address

Family Identifier (AFI) [RFC1700]. Not only IP addresses, but other addressing information have privacy requirements. Access to private information is granted only to those who are authorized and authenticated. Using asymmetric keying with public key cryptography enforces authentication for entities that read from and write to the mapping system. The proposal described in this document takes advantage of the latest in Elliptic Curve Cryptography.

In this proposal the EID is derived from a public key, and the corresponding private key is used to authenticate and authorize Map-Register messages. Thus only the owner of the corresponding private key can create and update mapping entries from the EID. Furthermore, the same approach is used to authenticate Map-Request messages. This in combination with the mapping database containing authorization information for Map-Requests is used to restrict which EIDs can lookup up the RLOCs for another EID.

This specification introduces how to use the Distinguished-Name AFI [AFI] and the [RFC8060] LCAF JSON Type to encode public keys and signatures in the LISP mapping database. The information in the mapping database is used to verify cryptographic signatures in LISP control-plane messages such as the Map-Request and Map-Register.

2. Definition of Terms

Crypto-EID: is an IPv6 EID where part of the EID includes a hash value of a public-key. An IPv6 EID is a Crypto-EID when the Map-Server is configured with an Crypto-EID Prefix that matches the IPv6 EID.

Crypto-EID Hash Length: is the number of low-order bits in a Crypto-EID which make up the hash of a public-key. The hash length is determined by the Map-Server when it is configured with a Crypto-EID Prefix.

Crypto-EID Prefix: is a configuration parameter on the Map-Server that indicates which IPv6 EIDs are Crypto-EIDs and what is the Crypto-EID Hash Length for the IPv6 EID. This can be different for different LISP Instance-IDs.

Hash-EID: is a distinguished name EID-record stored in the mapping

database. The EID format is 'hash-<pubkey-hash>'. When a key-pair is generated for an endpoint, the produced private-key does not leave the xTR that will register the Crypto-EID. A hash of the public-key is used to produce a Crypto-EID and a Hash-EID. The Crypto-EID is assigned to the endpoint and the xTR that supports the LISP-site registers the Crypto-EID. Another entity registers the Hash-EID mapping with the public-key as an RLOC-record.

Public-Key RLOC: is a JSON string that encodes a public-key as an RLOC-record for a Hash-EID mapping entry. The format of the JSON string is '{ "public-key" : "<pubkey>" }'.

Control-Plane Signature: a Map-Request or Map-Register sender signs the message with its private key. The format of the signature is a JSON string that includes sender information and the signature value. The JSON string is included in Map-Request and Map-Register messages.

Signature-ID: is a Crypto-EID used for a Control-Plane signature to register or request any type of EID. The Signature-ID is included with the JSON-encoded signature in Map-Request and Map-Register messages.

Multi-Signatures: multiple signatures are used in LISP when an entity allows and authorized another entity to register an EID. There can be more than one authorizing entities that allow a registering entity to register an EID. The authorizing entities sign their own RLOC-records that are registered and merged into the registering entity's Hash-EID public-key mapping. And when the registering entity registers the EID, all authorizing entity signatures must be verified by the Map-Server before the EID is accepted.

3. Overview

LISP already has several message authentication mechanisms. They can be found in [RFC9301], [RFC9303], and [RFC8061]. The mechanisms in this draft are providing a more granular level of authentication as well as a simpler way to manage keys and passwords.

A client of the mapping system can be authenticated using public-key cryptography. The client is required to have a private/public key-pair where it uses the private-key to sign Map-Requests and Map-Registers. The server, or the LISP entity, that processes Map-Requests and Map-Registers uses the public-key to verify signatures.

The following describes how the mapping system is used to implement the public-key crypto system:

1. An entity registers Hash-EID to Public-Key RLOC mappings. A third-party entity that provides a service can register or the client itself can register.
 2. Anyone can lookup the Hash-EID mappings. These mappings are not usually authenticated with the mechanisms in this draft but use the shared configured password mechanisms from [RFC9301] that provide group level authentication.
 3. When a Crypto-EID, or any EID type, is registered to the mapping system, a signature is included in the Map-Register message. When a non-Crypto-EID is registered a Signature-ID is also included in the Map-Register message.
 4. The Map-Server processes the registration by constructing the Hash-EID from the registered Crypto-EID, looks up the Hash-EID in the mapping system, obtains the public-key from the RLOC-record and verifies the signature. If Hash-EID lookup fails or the signature verification fails, the Map-Register is not accepted.
 5. When a Crypto-EID, or any EID type, is looked up in the mapping system, a signature is included with a Signature-ID in the Map-Request message.
 6. The Map-Server processes the request for a Crypto-EID by constructing the Hash-EID from the Signature-ID included in the Map-Request. The signer-ID is a Crypto-EID that accompanies a signature in the Map-Request. The Hash-EID is looked up in the mapping system, obtains the public-key from the RLOC-record and verifies the Map-Request signature. If the Hash-EID lookup fails or the signature verification fails, the Map-Request is not accepted and a Negative Map-Reply is sent back with an action of "auth-failure".
4. Public-Key Hash

When a private/public key-pair is created for a node, its IPv6 EID is pre-determined based on the public key generated. Note if the key-pair is compromised or is changed for the node, a new IPv6 EID is assigned for the node.

The sha256 [RFC6234] hex digest function is used to compute the hash. The hash is run over the following hex byte string:

<iid><prefix><pubkey>

Where each field is defined to be:

<iid>: is a 4-byte (leading zeroes filled) binary value of the Instance-ID the EID will be registered with in the mapping database. For example, if the instance-id is 171, then the 4-byte value is 0x000000ab.

<prefix>: is a variable length IPv6 prefix in binary format (with no colons) and IS quad-nibble zero-filled. The length of the prefix is 128 minus the Crypto-EID hash bit length. For example, if the prefix is 2001:5:3::/48, then the 6 byte value is 0x200100050003.

<pubkey>: is a DER [RFC7468] encoded public-key.

The public-key hash is used to construct the Crypto-EID and Hash-EID.

5. Hash-EID Mapping Entry

A Hash-EID is formatted in an EID-record as a Distinguished-Name AFI as specified in [I-D.ietf-lisp-name-encoding]. The format of the string is:

```
EID-record: 'hash-<hash-eid>'
```

Where <hash-eid> is a public-key hash as described in Section 4. The RLOC-record to encode and store the public-key is in LCAF JSON Type format of the form:

```
RLOC-record: '{ "public-key" : "<pubkey-base64>" }'
```

Where <pubkey-base64> is a base64 [RFC4648] encoding of the public-key generated for the system that is assigned the Hash-EID.

6. Hash-EID Structure

Since the Hash-EID is formatted as a distinguished-name AFI, the format of the <hash-eid> for EID 'hash-<hash-eid>' needs to be specified. The format will be an IPv6 address [RFC3513] where colons are used between quad-nibble characters when the hash bit length is a multiple of 4. And when the hash bit length is not a multiple of 4 but a multiple of 2, a leading 2 character nibble-pair is present.

Here are some examples for different hash bit lengths:

Crypto-EID: 2001:5::1111:2222:3333:4444, hash length 64:
Hash-EID is: 'hash-1111:2222:3333:4444'

Crypto-EID: 2001:5::11:22:33:44, hash length 64:
Hash-EID is: 'hash-0011:0022:0033:0044'

Crypto-EID: 2001:5:aaaa:bbbb:1111:2222:3333:4444, hash length 80:
Hash-EID is: 'hash-bbbb:1111:2222:3333:4444'

Crypto-EID: 2001:5:aaaa:bbbb:1111:2222:3333:4444, hash length 72:
Hash-EID is: 'hash-bb:1111:2222:3333:4444'

Crypto-EID: 2001:5:aaaa:bbbb:1111:22:33:4444, hash length 72:
Hash-EID is: 'hash-bb:1111:0022:0033:4444'

Note when leading zeroes exist in a IPv6 encoded quad between colons, the zeros are included in the quad for the Hash-EID string.

The entity that creates the hash, the entity that registers the Crypto-EID and the Map-Server that uses the hash for Hash-EID lookups MUST agree on the hash bit length.

7. Keys and Signatures

Key generation, message authentication with digital signatures, and signature verification will use the Elliptic Curve Digital Signature Algorithm or ECDSA [X9.62]. For key generation curve 'NIST256p' is used and recommended.

Signatures are computed over signature data that depends on the type of LISP message sent. See Section 8 and Section 9 for each message type. The signature data is passed through a sha256 hash function before it is signed or verified.

8. Signed Map-Register Encoding

When a ETR registers its Crypto-EID or any EID type to the mapping system, it builds a LISP Map-Register message. The mapping includes an EID-record which encodes the Crypto-EID, or any EID type, and an RLOC-set. One of the RLOC-records in the RLOC-set includes the the ETR's signature and signature-ID. The RLOC-record is formatted with a LCAF JSON Type, in the following format:

```
{ "signature" : "<signature-base64>", "signature-id" : "<signer-id>" }
```

Where <signature-base64> is a base64 [RFC4648] encoded string over the following ascii [RFC0020] string signature data:

```
[<iid>]<crypto-eid>
```

Where <iid> is the decimal value of the instance-ID the Crypto-EID is registering to and the <crypto-eid> is in the form of [RFC3513] where quad-nibbles between colons ARE NOT zero-filled.

The Map-Server that process an EID-record with a Crypto-EID and a RLOC-record with a signature extracts the public-key hash value from the Crypto-EID to build a Hash-EID. The Map-Server looks up the Hash-EID in the mapping system to obtain the public-key RLOC-record. The Map-Server verifies the signature over the signature data to determine if it should accept the EID-record registration.

9. Signed Map-Request Encoding

When an xTR (an ITR, Pitr, or RTR), sends a Map-Request to the mapping system to request the RLOC-set for a Crypto-EID, it signs the Map-Request so it can authenticate itself to the Map-Server the Crypto-EID is registered to. The Map-Request target-EID field will contain the Crypto-EID and the source-EID field will contain an LCAF JSON Type string with the following signature information:

```
{
  "source-eid" : "<seid>",
  "signature-id" : "<signer-id>",
  "signature" : "<signature-base64>"
}
```

Where <signer-id> is an IPv6 encoded string according to [RFC3513] where quad-nibbles between colons ARE NOT zero-filled. The <seid> is the source EID from the data packet that is invoking the Map-Request or the entire key/value pair for "source-eid" can be excluded when a data packet did not invoke the Map-Request (i.e. lig or an API request). The <signer-id> is the IPv6 Crypto-EID of the xTR that is providing the Map-Request signature.

The signature string <signature-base64> is a base64 [RFC4648] encoded string over the following signature data:

<nonce><source-eid><crypto-eid>

Where <nonce> is a hex string [RFC0020] of the nonce used in the Map-Request and the <source-eid> and <crypto-eid> are hex strings [RFC0020] of an IPv6 address in the form of [RFC3513] where quad-nibbles between colons ARE NOT zero-filled. When <seid> is not included in the Map-Request, string "0::0" is used for <source-eid>.

10. Signed Map-Notify Encoding

When a Map-Server originates a Map-Notify message either as an acknowledgment to a Map-Register message, as a solicited [I-D.ietf-lisp-pubsub] notification, or an unsolicited [RFC8378] notification, the receiver of the Map-Notify can verify the message is from an authenticated Map-Server.

An RLOC-record similar to the one used to sign Map-Register messages is used to sign the Map-Notify message:

```
{ "signature" : "<signature-base64>", "signature-id" : "<signer-id>" }
```

Where the "signature-id" is an IPv6 crypto-EID used by the Map-Server to sign the RLOC-record. The signature data and the encoding format of the signature is the same as for a Map-Register message. See details in Section 8.

A receiver of a Map-Notify message will lookup the signature-id in the mapping system to obtain a public-key to verify the signature. The Map-Notify is accepted only if the verification is successful.

11. Other Uses

The mechanisms described within this document can be used to sign other types of LISP messages. And for further study is how to use these mechanisms to sign LISP encapsulated data packets in a compressed manner to reduce data packet header overhead.

In addition to authenticating other types of LISP messages, other types of EID-records can be encoded as well and is not limited to IPv6 EIDs. It is possible for a LISP xTR to register and request non IPv6 EIDs but use IPv6 Crypto-EIDs for the sole purpose of signing and verifying EID-records.

Examples of other EID types that can be authenticated in Map-Request and Map-Register messages are:

EID-Type	Format Definition
IPv4 address prefixes	[RFC1123]
Distinguished-Names	[I-D.ietf-lisp-name-encoding]
Geo-Coordinates	[I-D.farinacci-lisp-geo]
LCAF defined EIDs	[RFC8060]

Table 1

12. EID Authorization

When a Crypto-EID is being used for IPv6 communication, it is implicit that the owner has the right to use the EID since it was generated from the key-pair provisioned for the owner. For other EID types that are not directly associated with signature keys, they must be validated for use by the mapping system they are registered to. This policy information for the mapping system must be configured in the Map-Servers the EID owner registers to or a signed authorization provided by a third-party entity.

To achieve signed authorization, an entity that allows another entity to register an EID, must authorize the registering entity. It does so by adding RLOC-records to the registering entity's Hash-EID public-key mapping. The format of the RLOC-record is a JSON encoded record as follows:

```
{
  "allow-eid" : "<eid>",
  "signature-id" : "<signer-id>",
  "signature" : "<signature-base64>"
}
```

The format of the <signer-id> and <signature-base64> values are the same as described in Section 8. The <eid> value is in the same string format as the signature data described in Section 8. For other non-IPv6 EID types, the conventions in [RFC8060] are used. In all cases, the string encoding format of instance-ID '[<iid>]' prepended is to the EID string.

This entry is added to the RLOC-set of the registering entity's Hash-EID 'hash-`<hash>`' registration. The authorizing entity does signs the Map-Register and sends it with merge-semantics. The Map-Server accepts the registration after the signature is verified and merges the RLOC-record into the existing RLOC-set. The 'signature' is optional and when not included means the authorizing entity has not yet allowed the registering entity to register the EID `<eid>`. Note multiple entities can register RLOC-records with the same `<eid>` meaning that signature verification for all of them is required before the Map-Server accepts the registration.

When the Map-Server receives a Map-Register for `<eid>`, it looks up 'hash-`<hash>`' EID in the mapping system. If not found, the Map-Register EID-record is not processed and the next EID-record is retrieved from the Map-Register message, if it exists. If the Hash-EID entry is found, the registering entity's signature is verified first. If the verification fails, the Map-Register EID-record is not accepted. Otherwise, a search for the RLOC-set is done to look for all matches of the EID being registered with `<eid>`, for those entries found, if any of them do not have a "signature" JSON item, the EID-record is not accepted. Otherwise, the signature-id is looked up in the mapping system to retrieve the public-key of the authorizing entity. If the verification is successful, then a lookup for the next RLOC-record signature-id is done. Only when all signature verifications are verified, the Map-Register EID-record is accepted.

The Map-Server should reject an RLOC-record with a signature-id that contains the Hash-EID of the entry disallowing a registering entity to self authorize itself.

Here is an example of a Hash-EID mapping stored in the mapping system:

EID-record: [1000]'hash-1111:2222:3333:4444', RLOC-Set (count is 4):

```
RLOC-record: { "public-key" : "<pubkey-base64>" }
RLOC-record: { "allow-eid" : "[1000]1.1.1.1/32",
               "signature" : "<sig>",
               "signature-id" : "[1000]2001:5:3::1111" }
RLOC-record: { "allow-eid" : "[1000]1.1.1.1/32",
               "signature" : "<sig>",
               "signature-id" : "[1000]2001:5:3::2222" }
RLOC-record: { "allow-eid" : "37-16-46-N-121-52-4-W",
               "signature-id" : "[1000]2001:5:3::5555" }
```

This mapping stores the public-key of the registering entity with Hash-EID 1111:2222:3333:4444. The registering entry registered this RLOC-record. There are two authorizing entities, :1111 and :2222, who allow it to register IPv4 EID 1.1.1.1/32. They each registered their respective RLOC-records. And a third authorizing entity :5555 that registers an RLOC-record that has not yet authorized the registering entity to register Geo-Coordinate 37-16-46-N-121-52-4-W. Note the mapping and the signature-IDs are all within the context of instance-ID 1000.

13. Security Considerations

The mechanisms within this specification are intentionally using accepted practices and state of the art public-key cryptography.

Crypto-EIDs can be made private when control messages are encrypted, for instance, using [RFC8061].

The topological or physical location of a Crypto-EID is only available to the other Crypto-EIDs that register in the same LISP Instance-ID and have their corresponding Hash-EIDs registered.

This draft doesn't address reply attacks directly. If a man-in-the-middle captures Map-Register messages, it could send such captured packets at a later time which contains signatures of the source. In which case, the Map-Server verifies the signature as good and interprets the contents to be valid where in fact the contents can contain old mapping information. This problem can be solved by encrypting the contents of Map-Registers using a third-party protocol like DTLS [RFC6347] or LISP-Crypto [RFC8061] directly by encapsulating Map-Registers in LISP data packets (using port 4341).

Map-Reply message signatures and authentication are not in scope for this document. This document focuses on authentication between xTRs and mapping system components. Map-Reply authentication, which is performed between xTRs is described in [RFC9303].

14. IANA Considerations

Since there are no new packet formats introduced for the functionality in this specification, there are no specific requests for IANA.

15. References

15.1. Normative References

- [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>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<https://www.rfc-editor.org/info/rfc1123>>.
- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, DOI 10.17487/RFC1700, October 1994, <<https://www.rfc-editor.org/info/rfc1700>>.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, DOI 10.17487/RFC3513, April 2003, <<https://www.rfc-editor.org/info/rfc3513>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.

- [RFC7468] Josefsson, S. and S. Leonard, "Textual Encodings of PKIX, PKCS, and CMS Structures", RFC 7468, DOI 10.17487/RFC7468, April 2015, <<https://www.rfc-editor.org/info/rfc7468>>.
- [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>>.
- [RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.
- [RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <<https://www.rfc-editor.org/info/rfc8378>>.
- [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>>.
- [RFC9303] Maino, F., Ermagan, V., Cabellos, A., and D. Saucez, "Locator/ID Separation Protocol Security (LISP-SEC)", RFC 9303, DOI 10.17487/RFC9303, October 2022, <<https://www.rfc-editor.org/info/rfc9303>>.

15.2. Informative References

- [AFI] "Address Family Identifier (AFIs)", ADDRESS FAMILY NUMBERS <http://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml>?, February 2007.
- [I-D.farinacci-lisp-geo]
Farinacci, D., "LISP Geo-Coordinate Use-Cases", Work in Progress, Internet-Draft, draft-farinacci-lisp-geo-15, 10 November 2022, <<https://datatracker.ietf.org/doc/html/draft-farinacci-lisp-geo-15>>.
- [I-D.ietf-lisp-name-encoding]
Farinacci, D., "LISP Distinguished Name Encoding", Work in Progress, Internet-Draft, draft-ietf-lisp-name-encoding-

05, 28 December 2023,
<<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-name-encoding-05>>.

[I-D.ietf-lisp-pubsub]

Rodriguez-Natal, A., Ermagan, V., Cabellos-Aparicio, A., Barkai, S., and M. Boucadair, "Publish/Subscribe Functionality for the Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-pubsub-15, 28 February 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-pubsub-15>>.

[X9.62] "Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA)", NIST ANSI X9.62-2005, November 2005.

Appendix A. Acknowledgments

A special thanks goes to Sameer Merchant and Colin Cantrell for their ideas and technical contributions to the ideas in this draft.

Appendix B. Document Change Log

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-ietf-lisp-ecdsa-auth-12

- * Posted February 2024.
- * Update references and document timer.

B.2. Changes to draft-ietf-lisp-ecdsa-auth-11

- * Posted August 2023.
- * Update references (to proposed standard documents) and document timer.

B.3. Changes to draft-ietf-lisp-ecdsa-auth-10

- * Posted March 2023.
- * Update references (to proposed standard documents) and document timer.

- B.4. Changes to draft-ietf-lisp-ecdsa-auth-09
 - * Posted September 2022.
 - * Update draft-ietf-lisp-name-encoding reference.
- B.5. Changes to draft-ietf-lisp-ecdsa-auth-08
 - * Posted August 2022.
 - * Update references and document timer.
- B.6. Changes to draft-ietf-lisp-ecdsa-auth-07
 - * Posted February 2022.
 - * Update references and document timer.
- B.7. Changes to draft-ietf-lisp-ecdsa-auth-06
 - * Posted August 2021.
 - * Update references and document timer.
- B.8. Changes to draft-ietf-lisp-ecdsa-auth-05
 - * Posted March 2021.
 - * Update references and document timer.
- B.9. Changes to draft-ietf-lisp-ecdsa-auth-04
 - * Posted September 2020.
 - * Update references and document timer.
- B.10. Changes to draft-ietf-lisp-ecdsa-auth-03
 - * Posted March 2020.
 - * Update references and document timer.
- B.11. Changes to draft-ietf-lisp-ecdsa-auth-02
 - * Posted September 2019.
 - * Update references and document timer.

B.12. Changes to draft-ietf-lisp-ecdsa-auth-01

- * Posted March IETF week 2019.
- * Update references and document timer.

B.13. Changes to draft-ietf-lisp-ecdsa-auth-00

- * Posted mid-September 2018.
- * Make draft-farinacci-lisp-ecdsa-auth-03 a LISP working group docuemnt.

B.14. Changes to draft-farinacci-lisp-ecdsa-auth-03

- * Posted September 2018.
- * Change all occurrences of signature-EID to signature-ID.
- * Document how Map-Servers sign Map-Notify messages so they can be verified by xTRs.
- * Add multi-signatures to mappings so a 3rd-party can allow an entity to register any type of EID.

B.15. Changes to draft-farinacci-lisp-ecdsa-auth-02

- * Draft posted April 2018.
- * Generalize text to allow Map-Requesting and Map-Registering for any EID type with a proper signature-EID and signature encoded together.

B.16. Changes to draft-farinacci-lisp-ecdsa-auth-01

- * Draft posted October 2017.
- * Make it more clear what values and format the EID hash is run over.
- * Update references to newer RFCs and Internet Drafts.

B.17. Changes to draft-farinacci-lisp-ecdsa-auth-00

- * Initial draft posted July 2017.

Authors' Addresses

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

Erik Nordmark
Zededa
Santa Clara, CA
United States of America
Email: erik@zededa.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 25, 2019

D. Farinacci
lispers.net
P. Pillay-Esnault
Huawei Technologies
W. Haddad
Ericsson
October 22, 2018

LISP EID Anonymity
draft-ietf-lisp-eid-anonymity-04

Abstract

This specification will describe how ephemeral LISP EIDs can be used to create source anonymity. The idea makes use of frequently changing EIDs much like how a credit-card system uses a different credit-card numbers for each transaction.

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 April 25, 2019.

Copyright Notice

Copyright (c) 2018 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. Overview	3
4. Design Details	4
5. Other Types of Ephemeral-EIDs	5
6. Interworking Considerations	5
7. Multicast Considerations	5
8. Performance Improvements	6
9. Security Considerations	6
10. IANA Considerations	6
11. References	6
11.1. Normative References	6
11.2. Informative References	8
Appendix A. Acknowledgments	8
Appendix B. Document Change Log	8
B.1. Changes to draft-ietf-lisp-eid-anonymity-04	8
B.2. Changes to draft-ietf-lisp-eid-anonymity-03	8
B.3. Changes to draft-ietf-lisp-eid-anonymity-02	9
B.4. Changes to draft-ietf-lisp-eid-anonymity-01	9
B.5. Changes to draft-ietf-lisp-eid-anonymity-00	9
B.6. Changes to draft-farinacci-lisp-eid-anonymity-02	9
B.7. Changes to draft-farinacci-lisp-eid-anonymity-01	9
B.8. Changes to draft-farinacci-lisp-eid-anonymity-00	9
Authors' Addresses	10

1. Introduction

The LISP architecture [RFC6830] specifies two namespaces, End-Point IDs (EIDs) and Routing Locators (RLOCs). An EID identifies a node in the network and the RLOC indicates the EID's topological location. Typically EIDs are globally unique so a end-node system can connect to any other end-node system on the Internet. Privately used EIDs are allowed when scoped within a VPN but must always be unique within that scope. Therefore, address allocation is required by network administration to avoid address collisions or duplicate address use. In a multiple namespace architecture like LISP, typically the EID

will stay fixed while the RLOC can change. This occurs when the EID is mobile or when the LISP site the EID resides in changes its connection to the Internet.

LISP creates the opportunity where EIDs are fixed and won't change. This can create a privacy problem more so than what we have on the Internet today. This draft will examine a technique to allow a end-node system to use a temporary address. The lifetime of a temporary address can be the same as a lifetime of an address in use today on the Internet or can have traditionally shorter lifetimes, possibly on the order of a day or even change as frequent as new connection attempts.

2. Definition of Terms

Ephemeral-EID - is an IP address that is created randomly for use for a temporary period of time. An Ephemeral-EID has all the properties of an EID as defined in [RFC6830]. Ephemeral-EIDs are not stored in the Domain Name System (DNS) and should not be used in long-term address referrals.

Client End-Node - is a network node that originates and consumes packets. It is a system that originates packets or initiates the establishment of transport-layer connections. It does not offer services as a server system would. It accesses servers and attempts to do it anonymously.

3. Overview

A client end-node can assign its own ephemeral EID and use it to talk to any system on the Internet. The system is acting as a client where it initiates communication and desires to be an inaccessible resource from any other system. The ephemeral EID is used as a destination address solely to return packets to resources the ephemeral EID connects to.

Here is the procedure a client end-node would use:

1. Client end-node desires to talk on the network. It creates and assigns an ephemeral-EID on any interface. The client end-node may also assign multiple ephemeral-EIDs on the same interface or across different interfaces.
2. If the client end-node is a LISP xTR, it will register ephemeral-EIDs mapped to underlay routable RLOCs. If the client end-node is not a LISP xTR, it can send packets on the network where a LISP router xTR will register the ephemeral-EIDs with its RLOC-set.

3. The client end-node originates packets with a source address equal to the ephemeral-EID and will receive packets addressed to the ephemeral-EID.
4. When the client end-node decides to stop using an ephemeral-EID, it will deregister it from the mapping system and create and assign a new ephemeral-EID, or decide to configure a static global address, or participate in DHCP to get assigned a leased address.

Note that the ephemeral-EID can be mobile just like any other EID so if it is initially registered to the mapping system with one or more RLOCs, later the RLOC-set can change as the ephemeral-EID roams.

4. Design Details

This specification proposes the use of the experimental LISP EID-block 2001:5::/32 [RFC7954] when IPv6 is used. See IANA Considerations section for a specific sub-block allocation request. When IPv4 is used, the Class E block 240.0.0.0/4 is being proposed.

The client end-node system will use the rest of the host bits to allocate a random number to be used as the ephemeral-EID. The EID can be created manually or via a programatic interface. When the EID address is going to change frequently, it is suggested to use a programatic interface. The probability of address collision is unlikely for IPv6 EIDs but could occur for IPv4 EIDs. A client end-node can create a ephemeral-EID and then look it up in the mapping system to see if it exists. If the EID exists in the mapping system, the client end-node can attempt creation of a new random number for the ephemeral-EID. See Section 8 where ephemeral-EIDs can be preallocated and registered to the mapping system before use.

When the client end-node system is co-located with the RLOC and acts as an xTR, it should register the binding before sending packets. This eliminates a race condition for returning packets not knowing where to encapsulate packets to the ephemeral-EID's RLOCs. See Section 8 for alternatives for fixing this race condition problem. When the client end-node system is not acting as an xTR, it should send some packets so its ephemeral-EID can be discovered by an xTR which supports EID-mobility [I-D.ietf-lisp-eid-mobility] so mapping system registration can occur before the destination returns packets. When the end-node system is acting as an xTR, the EID and RLOC-set is co-located in the same node. So when the EID is created, the xTR can register the mapping versus waiting for packet transmission.

5. Other Types of Ephemeral-EIDs

When IPv6 Ephemeral-EIDs are used, an alternative to a random number can be used. For example, the low-order bits of the IPv6 address could be a cryptographic hash of a public-key. Mechanisms from [RFC3972] could be used for EIDs. Using this approach allows the sender with a hashed EID to be authenticated. So packet signatures can be verified by the corresponding public-key. When hashed EIDs are used, the EID can change frequently as rekeying may be required for enhanced security. LISP specific control message signature mechanisms can be found in [I-D.farinacci-lisp-ecdsa-auth].

6. Interworking Considerations

If a client end-node is communicating with a system that is not in a LISP site, the procedures from [RFC6832] should be followed. The PISTR will be required to originate route advertisements for the ephemeral-EID sub-block [RFC7954] so it can attract packets sourced by non-LISP sites destined to ephemeral-EIDs. However, in the general case, the coarse block from [RFC7954] will be advertised which would cover the sub-block. For IPv4, the 240.0.0.0/4 must be advertised into the IPv4 routing system.

7. Multicast Considerations

A client end-node system can be a member of a multicast group fairly easily since its address is not used for multicast communication as a receiver. This is due to the design characteristics of IGMP [RFC3376] [RFC2236] [RFC1112] and MLD [RFC2710] [RFC3810].

When a client end-node system is a multicast source, there is ephemeral (S,G) state that is created and maintained in the network via multicast routing protocols such as PIM [RFC4602] and when PIM is used with LISP [RFC6802]. In addition, when [I-D.ietf-lisp-signal-free-multicast] is used, ephemeral-EID state is created in the mapping database. This doesn't present any problems other than the amount of state that may exist in the network if not timed out and removed promptly.

However, there exists a multicast source discovery problem when PIM-SSM [RFC4607] is used. Members that join (S,G) channels via out of band mechanisms. These mechanisms need to support ephemeral-EIDs. Otherwise, PIM-ASM [RFC4602] or PIM-Bidir [RFC5015] will need to be used.

8. Performance Improvements

An optimization to reduce the race condition between registering ephemeral-EIDs and returning packets as well as reducing the probability of ephemeral-EID address collision is to preload the mapping database with a list of ephemeral-EIDs before using them. It comes at a expense of rebinding all of registered ephemeral-EIDs when there is an RLOC change. There is work in progress to consider adding a level of indirection here so a single entry gets the RLOC update and the list of ephemeral-EIDs point to the single entry.

9. Security Considerations

When LISP-crypto [RFC8061] is used the EID payload is more secure through encryption providing EID obfuscation of the ephemeral-EID as well as the global-EID it is communicating with. But the obfuscation only occurs between xTRs. So the randomness of a ephemeral-EID inside of LISP sites provide a new level of privacy.

10. IANA Considerations

This specification is requesting the sub-block 2001:5:ffff::/48 for ephemeral-EID usage.

11. References

11.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [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>>.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, DOI 10.17487/RFC2236, November 1997, <<https://www.rfc-editor.org/info/rfc2236>>.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <<https://www.rfc-editor.org/info/rfc2710>>.

- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4602] Pusateri, T., "Protocol Independent Multicast - Sparse Mode (PIM-SM) IETF Proposed Standard Requirements Analysis", RFC 4602, DOI 10.17487/RFC4602, August 2006, <<https://www.rfc-editor.org/info/rfc4602>>.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, DOI 10.17487/RFC5015, October 2007, <<https://www.rfc-editor.org/info/rfc5015>>.
- [RFC6802] Baillargeon, S., Flinta, C., and A. Johnsson, "Ericsson Two-Way Active Measurement Protocol (TWAMP) Value-Added Octets", RFC 6802, DOI 10.17487/RFC6802, November 2012, <<https://www.rfc-editor.org/info/rfc6802>>.
- [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>>.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC7954] Iannone, L., Lewis, D., Meyer, D., and V. Fuller, "Locator/ID Separation Protocol (LISP) Endpoint Identifier (EID) Block", RFC 7954, DOI 10.17487/RFC7954, September 2016, <<https://www.rfc-editor.org/info/rfc7954>>.

[RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.

11.2. Informative References

[I-D.farinacci-lisp-ecdsa-auth]

Farinacci, D. and E. Nordmark, "LISP Control-Plane ECDSA Authentication and Authorization", draft-farinacci-lisp-ecdsa-auth-03 (work in progress), September 2018.

[I-D.ietf-lisp-eid-mobility]

Portoles-Comeras, M., Ashtaputre, V., Moreno, V., Maino, F., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", draft-ietf-lisp-eid-mobility-02 (work in progress), May 2018.

[I-D.ietf-lisp-signal-free-multicast]

Moreno, V. and D. Farinacci, "Signal-Free LISP Multicast", draft-ietf-lisp-signal-free-multicast-09 (work in progress), March 2018.

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

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-ietf-lisp-eid-anonymity-04

- o Posted October 2018 before Bangkok IETF deadline.
- o Made Padma requested changes to refer to ephemeral-EIDs allowed to have many on one interface and can be registered with more than 1 RLOC but one RLOC-set.

B.2. Changes to draft-ietf-lisp-eid-anonymity-03

- o Posted October 2018.
- o Update document timer and references.

- B.3. Changes to draft-ietf-lisp-eid-anonymity-02
 - o Posted April 2018.
 - o Update document timer and references.
- B.4. Changes to draft-ietf-lisp-eid-anonymity-01
 - o Posted October 2017.
 - o Add to section 5 that PKI can be used to authenticate EIDs.
 - o Update references.
- B.5. Changes to draft-ietf-lisp-eid-anonymity-00
 - o Posted August 2017.
 - o Made draft-farinacci-lisp-eid-anonymity-02 a LISP working group document.
- B.6. Changes to draft-farinacci-lisp-eid-anonymity-02
 - o Posted April 2017.
 - o Added section describing how ephemeral-EIDs can use a public key hash as an alternative to a random number.
 - o Indciate when an EID/RLOC co-located, that the xTR can register the EID when it is configured or changed versus waiting for a packet to be sent as in the EID/RLOC separated case.
- B.7. Changes to draft-farinacci-lisp-eid-anonymity-01
 - o Posted October 2016.
 - o Update document timer.
- B.8. Changes to draft-farinacci-lisp-eid-anonymity-00
 - o Posted April 2016.
 - o Initial posting.

Authors' Addresses

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Padma Pillay-Esnault
Huawei Technologies
San Clara, CA
USA

Email: padma@huawei.com

Wassim Haddad
Ericsson
San Clara, CA
USA

Email: wassim.haddad@ericsson.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 22 August 2024

D. Farinacci
lispers.net
P. Pillay-Esnault
Independent
W. Haddad
Ericsson
19 February 2024

LISP EID Anonymity
draft-ietf-lisp-eid-anonymity-16

Abstract

This specification will describe how ephemeral LISP EIDs can be used to create source anonymity. The idea makes use of frequently changing EIDs much like how a credit-card system uses a different credit-card numbers for each transaction.

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 22 August 2024.

Copyright Notice

Copyright (c) 2024 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	3
2. Definition of Terms	3
3. Overview	3
4. Design Details	4
5. Other Types of Ephemeral-EIDs	5
6. Interworking Considerations	5
7. Multicast Considerations	5
8. Performance Improvements	6
9. Security Considerations	6
10. IANA Considerations	6
11. References	6
11.1. Normative References	6
11.2. Informative References	8
Appendix A. Acknowledgments	8
Appendix B. Document Change Log	9
B.1. Changes to draft-ietf-lisp-eid-anonymity-16	9
B.2. Changes to draft-ietf-lisp-eid-anonymity-15	9
B.3. Changes to draft-ietf-lisp-eid-anonymity-14	9
B.4. Changes to draft-ietf-lisp-eid-anonymity-13	9
B.5. Changes to draft-ietf-lisp-eid-anonymity-12	9
B.6. Changes to draft-ietf-lisp-eid-anonymity-11	9
B.7. Changes to draft-ietf-lisp-eid-anonymity-10	9
B.8. Changes to draft-ietf-lisp-eid-anonymity-09	10
B.9. Changes to draft-ietf-lisp-eid-anonymity-08	10
B.10. Changes to draft-ietf-lisp-eid-anonymity-07	10
B.11. Changes to draft-ietf-lisp-eid-anonymity-06	10
B.12. Changes to draft-ietf-lisp-eid-anonymity-05	10
B.13. Changes to draft-ietf-lisp-eid-anonymity-04	10
B.14. Changes to draft-ietf-lisp-eid-anonymity-03	10
B.15. Changes to draft-ietf-lisp-eid-anonymity-02	10
B.16. Changes to draft-ietf-lisp-eid-anonymity-01	11
B.17. Changes to draft-ietf-lisp-eid-anonymity-00	11
B.18. Changes to draft-farinacci-lisp-eid-anonymity-02	11
B.19. Changes to draft-farinacci-lisp-eid-anonymity-01	11
B.20. Changes to draft-farinacci-lisp-eid-anonymity-00	11
Authors' Addresses	11

1. Introduction

The LISP architecture [RFC9300] specifies two namespaces, End-Point IDs (EIDs) and Routing Locators (RLOCs). An EID identifies a node in the network and the RLOC indicates the EID's topological location. Typically EIDs are globally unique so an end-node system can connect to any other end-node system on the Internet. Privately used EIDs are allowed when scoped within a VPN but must always be unique within that scope. Therefore, address allocation is required by network administration to avoid address collisions or duplicate address use. In a multiple namespace architecture like LISP, typically the EID will stay fixed while the RLOC can change. This occurs when the EID is mobile or when the LISP site the EID resides in changes its connection to the Internet.

LISP creates the opportunity where EIDs are fixed and won't change. This draft will examine a technique to allow a end-node system to use a temporary address. The lifetime of a temporary address can be the same as a lifetime of an address in use today on the Internet or can have traditionally shorter lifetimes, possibly on the order of a day or even change as frequent as new connection attempts.

2. Definition of Terms

Ephemeral-EID - is an IP address that is created randomly for use for a temporary period of time. An Ephemeral-EID has all the properties of an EID as defined in [RFC9300]. Ephemeral-EIDs are not stored in the Domain Name System (DNS) and should not be used in long-term address referrals.

Client End-Node - is a network node that originates and consumes packets. It is a system that originates packets or initiates the establishment of transport-layer connections. It does not offer services as a server system would. It accesses servers and attempts to do it anonymously.

3. Overview

A client end-node can assign its own ephemeral EID and use it to talk to any system on the Internet. The system is acting as a client where it initiates communication and desires to be an inaccessible resource from any other system. The ephemeral EID is used as a destination address solely to return packets to resources the ephemeral EID connects to. A client-node may simultaneously use a traditional EID along with ephemeral EIDs in parallel and are not mutually exclusive. A client may choose to use the ephemeral EIDs with some peers only where it needs to preserve anonymity.

Here is the procedure a client end-node would use:

1. Client end-node desires to talk on the network. It creates and assigns an ephemeral-EID on any interface. The client end-node may also assign multiple ephemeral-EIDs on the same interface or across different interfaces.
2. If the client end-node is a LISP xTR, it will register ephemeral-EIDs mapped to underlay routable RLOCs. If the client end-node is not a LISP xTR, it can send packets on the network where a LISP router xTR will register the ephemeral-EIDs with its RLOC-set.
3. The client end-node originates packets with a source address equal to the ephemeral-EID and will receive packets addressed to the ephemeral-EID.
4. When the client end-node decides to stop using an ephemeral-EID, it will deregister it from the mapping system and create and assign a new ephemeral-EID, or decide to configure a static global address, or participate in DHCP to get assigned a leased address.

Note that the ephemeral-EID can be mobile just like any other EID so if it is initially registered to the mapping system with one or more RLOCs, later the RLOC-set can change as the ephemeral-EID roams.

4. Design Details

This specification proposes the use of the experimental LISP EID-block 2001:5::/32 [RFC7954] when IPv6 is used. See IANA Considerations section for a specific sub-block allocation request. When IPv4 is used, the Class E block 240.0.0.0/4 is being proposed.

The client end-node system will use the rest of the host bits to allocate a random number to be used as the ephemeral-EID. The EID can be created manually or via a programatic interface. When the EID address is going to change frequently, it is suggested to use a programatic interface. The probability of address collision is unlikely for IPv6 EIDs but could occur for IPv4 EIDs. A client end-node can create an ephemeral-EID and then look it up in the mapping system to see if it exists. If the EID exists in the mapping system, the client end-node can attempt creation of a new random number for the ephemeral-EID. See Section 8 where ephemeral-EIDs can be preallocated and registered to the mapping system before use.

When the client end-node system is co-located with the RLOC and acts as an xTR, it should register the binding before sending packets. This eliminates a race condition for returning packets not knowing where to encapsulate packets to the ephemeral-EID's RLOCs. See Section 8 for alternatives for fixing this race condition problem. When the client end-node system is not acting as an xTR, it should send some packets so its ephemeral-EID can be discovered by an xTR which supports EID-mobility [I-D.ietf-lisp-eid-mobility] so mapping system registration can occur before the destination returns packets. When the end-node system is acting as an xTR, the EID and RLOC-set is co-located in the same node. So when the EID is created, the xTR can register the mapping versus waiting for packet transmission.

5. Other Types of Ephemeral-EIDs

When IPv6 Ephemeral-EIDs are used, an alternative to a random number can be used. For example, the low-order bits of the IPv6 address could be a cryptographic hash of a public-key. Mechanisms from [RFC3972] could be used for EIDs. Using this approach allows the sender with a hashed EID to be authenticated. So packet signatures can be verified by the corresponding public-key. When hashed EIDs are used, the EID can change frequently as rekeying may be required for enhanced security. LISP specific control message signature mechanisms can be found in [I-D.farinacci-lisp-ecdsa-auth].

6. Interworking Considerations

If a client end-node is communicating with a system that is not in a LISP site, the procedures from [RFC6832] should be followed. The Pitr will be required to originate route advertisements for the ephemeral-EID sub-block [RFC7954] so it can attract packets sourced by non-LISP sites destined to ephemeral-EIDs. However, in the general case, the coarse block from [RFC7954] will be advertised which would cover the sub-block. For IPv4, the 240.0.0.0/4 must be advertised into the IPv4 routing system.

7. Multicast Considerations

A client end-node system can be a member of a multicast group fairly easily since its address is not used for multicast communication as a receiver. This is due to the design characteristics of IGMP [RFC3376] [RFC2236] [RFC1112] and MLD [RFC2710] [RFC3810].

When a client end-node system is a multicast source, there is ephemeral (S,G) state that is created and maintained in the network via multicast routing protocols such as PIM [RFC4602] and when PIM is used with LISP [RFC6802]. In addition, when [RFC8378] is used, ephemeral-EID state is created in the mapping database. This doesn't present any problems other than the amount of state that may exist in the network if not timed out and removed promptly.

However, there exists a multicast source discovery problem when PIM-SSM [RFC4607] is used. Members that join (S,G) channels via out of band mechanisms. These mechanisms need to support ephemeral-EIDs. Otherwise, PIM-ASM [RFC4602] or PIM-Bidir [RFC5015] will need to be used.

8. Performance Improvements

An optimization to reduce the race condition between registering ephemeral-EIDs and returning packets as well as reducing the probability of ephemeral-EID address collision is to preload the mapping database with a list of ephemeral-EIDs before using them. It comes at the expense of rebinding all of registered ephemeral-EIDs when there is an RLOC change. There is work in progress to consider adding a level of indirection here so a single entry gets the RLOC update and the list of ephemeral-EIDs point to the single entry.

9. Security Considerations

When LISP-crypto [RFC8061] is used the EID payload is more secure through encryption providing EID obfuscation of the ephemeral-EID as well as the global-EID it is communicating with. But the obfuscation only occurs between xTRs. So the randomness of a ephemeral-EID inside of LISP sites provide a new level of privacy.

10. IANA Considerations

This specification is requesting the sub-block 2001:5:ffff::/48 for ephemeral-EID usage.

11. References

11.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.

- [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>>.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, DOI 10.17487/RFC2236, November 1997, <<https://www.rfc-editor.org/info/rfc2236>>.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <<https://www.rfc-editor.org/info/rfc2710>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4602] Pusateri, T., "Protocol Independent Multicast - Sparse Mode (PIM-SM) IETF Proposed Standard Requirements Analysis", RFC 4602, DOI 10.17487/RFC4602, August 2006, <<https://www.rfc-editor.org/info/rfc4602>>.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, DOI 10.17487/RFC5015, October 2007, <<https://www.rfc-editor.org/info/rfc5015>>.
- [RFC6802] Baillargeon, S., Flinta, C., and A. Johnsson, "Ericsson Two-Way Active Measurement Protocol (TWAMP) Value-Added Octets", RFC 6802, DOI 10.17487/RFC6802, November 2012, <<https://www.rfc-editor.org/info/rfc6802>>.

- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC7954] Iannone, L., Lewis, D., Meyer, D., and V. Fuller, "Locator/ID Separation Protocol (LISP) Endpoint Identifier (EID) Block", RFC 7954, DOI 10.17487/RFC7954, September 2016, <<https://www.rfc-editor.org/info/rfc7954>>.
- [RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.
- [RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <<https://www.rfc-editor.org/info/rfc8378>>.
- [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>>.

11.2. Informative References

- [I-D.farinacci-lisp-ecdsa-auth]
Farinacci, D. and E. Nordmark, "LISP Control-Plane ECDSA Authentication and Authorization", Work in Progress, Internet-Draft, draft-farinacci-lisp-ecdsa-auth-03, 4 September 2018, <<https://datatracker.ietf.org/doc/html/draft-farinacci-lisp-ecdsa-auth-03>>.
- [I-D.ietf-lisp-eid-mobility]
Portoles-Comeras, M., Ashtaputre, V., Maino, F., Moreno, V., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-eid-mobility-13, 6 November 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-eid-mobility-13>>.

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

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-ietf-lisp-eid-anonymity-16

- * Posted February 2024.
- * Update references and document timer.

B.2. Changes to draft-ietf-lisp-eid-anonymity-15

- * Posted August 2023.
- * Update references (to proposed standard documents) and document timer.

B.3. Changes to draft-ietf-lisp-eid-anonymity-14

- * Posted March 2023.
- * Update references (to proposed standard documents) and document timer.

B.4. Changes to draft-ietf-lisp-eid-anonymity-13

- * Posted September 2022.
- * Update document timer and references.

B.5. Changes to draft-ietf-lisp-eid-anonymity-12

- * Posted March 2022.
- * Update document timer and references.

B.6. Changes to draft-ietf-lisp-eid-anonymity-11

- * Posted end of September 2021.
- * Update document timer and references.

B.7. Changes to draft-ietf-lisp-eid-anonymity-10

- * Posted end of March 2021.
- * Update document timer and references.

- B.8. Changes to draft-ietf-lisp-eid-anonymity-09
- * Posted end of October 2020.
 - * Update document timer and references.
- B.9. Changes to draft-ietf-lisp-eid-anonymity-08
- * Posted end of April 2020.
 - * Update document timer and references.
- B.10. Changes to draft-ietf-lisp-eid-anonymity-07
- * Posted end of October 2019.
 - * Update document timer and references.
- B.11. Changes to draft-ietf-lisp-eid-anonymity-06
- * Posted end of March 2019.
 - * Padma had more basic edits and some clarification text.
- B.12. Changes to draft-ietf-lisp-eid-anonymity-05
- * Posted March IETF week 2019.
 - * Do not state that ephemeral EIDs make the privacy problem worse.
- B.13. Changes to draft-ietf-lisp-eid-anonymity-04
- * Posted October 2018 before Bangkok IETF deadline.
 - * Made Padma requested changes to refer to ephemeral-EIDs allowed to have many on one interface and can be registered with more than 1 RLOC but one RLOC-set.
- B.14. Changes to draft-ietf-lisp-eid-anonymity-03
- * Posted October 2018.
 - * Update document timer and references.
- B.15. Changes to draft-ietf-lisp-eid-anonymity-02
- * Posted April 2018.

- * Update document timer and references.
- B.16. Changes to draft-ietf-lisp-eid-anonymity-01
- * Posted October 2017.
 - * Add to section 5 that PKI can be used to authenticate EIDs.
 - * Update references.
- B.17. Changes to draft-ietf-lisp-eid-anonymity-00
- * Posted August 2017.
 - * Made draft-farinacci-lisp-eid-anonymity-02 a LISP working group document.
- B.18. Changes to draft-farinacci-lisp-eid-anonymity-02
- * Posted April 2017.
 - * Added section describing how ephemeral-EIDs can use a public key hash as an alternative to a random number.
 - * Indciate when an EID/RLOC co-located, that the xTR can register the EID when it is configured or changed versus waiting for a packet to be sent as in the EID/RLOC separated case.
- B.19. Changes to draft-farinacci-lisp-eid-anonymity-01
- * Posted October 2016.
 - * Update document timer.
- B.20. Changes to draft-farinacci-lisp-eid-anonymity-00
- * Posted April 2016.
 - * Initial posting.

Authors' Addresses

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

Padma Pillay-Esnault
Independent
Santa Clara, CA
United States of America
Email: padma.ietf@gmail.com

Wassim Haddad
Ericsson
Santa Clara, CA
United States of America
Email: wassim.haddad@ericsson.com

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: April 7, 2019

A. Rodriguez-Natal
V. Ermagan
J. Leong
F. Maino
Cisco Systems
A. Cabellos-Aparicio
Technical University of Catalonia
S. Barkai
Fermi Serverless
D. Farinacci
lispers.net
M. Boucadair
C. Jacquenet
Orange
S. Secci
LIP6 UPMC
October 4, 2018

Publish/Subscribe Functionality for LISP
draft-ietf-lisp-pubsub-01

Abstract

This document specifies an extension to the use of Map-Request to enable Publish/Subscribe (PubSub) operation for 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 April 7, 2019.

Copyright Notice

Copyright (c) 2018 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. Requirements Language	3
3. Definition of Terms	3
4. Deployment Assumptions	4
5. Map-Request Additions	4
6. Mapping Request Subscribe Procedures	6
7. Mapping Notification Publish Procedures	8
8. Security Considerations	9
9. Acknowledgments	9
10. IANA Considerations	9
11. Normative References	10
Authors' Addresses	10

1. Introduction

The Locator/ID Separation Protocol (LISP) [I-D.ietf-lisp-rfc6833bis] splits current IP addresses in two different namespaces, Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). LISP uses a map-and-encap approach that relies on (1) a Mapping System (basically a distributed database) that stores and disseminates EID-RLOC mappings and on (2) LISP tunnel routers (xTRs) that encapsulate and decapsulate data packets based on the content of those mappings.

ITRs/RTRs/PITRs pull EID-to-RLOC mapping information from the Mapping System by means of an explicit request message.

[I-D.ietf-lisp-rfc6833bis] indicates how ETRs can tell ITRs/RTRs/PITRs about mapping changes. This document presents a Publish/Subscribe (PubSub) extension in which the Mapping System can notify ITRs/RTRs/PITRs about mapping changes. When this mechanism is used, mapping changes can be notified faster and can be managed in the Mapping System versus the LISP sites.

In general, when an ITR/RTR/PITR wants to be notified for mapping changes for a given EID-prefix, the following steps occur:

- (1) The ITR/RTR/PITR sends a Map-Request for that EID-prefix.
- (2) The ITR/RTR/PITR sets the Notification-Requested bit (N-bit) on the Map-Request and includes its xTR-ID and Site-ID.
- (3) The Map-Request is forwarded to one of the Map-Servers that the EID-prefix is registered to.
- (4) The Map-Server creates subscription state for the ITR/RTR/PITR on the EID-prefix.
- (5) The Map-Server sends a Map-Notify to the ITR/RTR/PITR to acknowledge the successful subscription.
- (6) When there is an RLOC-set change for the EID-prefix, the Map-Server sends a Map-Notify message to each ITR/RTR/PITR in the subscription list.
- (7) Each ITR/RTR/PITR sends a Map-Notify-Ack to acknowledge the received Map-Notify.

This operation is repeated for all EID-prefixes for which ITR/RTR/PITR want to be notified. The ITR/RTR/PITR can set the N-bit for several EID-prefixes within a single Map-Request.

2. 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].

3. Definition of Terms

LISP xTR-ID: A 128-bit field that is used as a unique identifier for an xTR. The xTR-ID is especially useful for identifying multiple xTRs serving the same site/EID-prefix. A value of all zeros indicate the xTR-ID is unspecified.

LISP Site-ID: A 64-bit field that is used as a unique identifier of a group of xTRs belonging to the same site. A value of 0 indicate the Site-ID is unspecified.

4. Deployment Assumptions

The specification described in this document makes the following deployment assumptions:

- (1) A unique 128-bit xTR-ID (plus a 64-bit Site-ID) identifier is assigned to each xTR.
- (2) Map-Servers are configured in proxy-reply mode, i.e., they are solicited to generate and send Map-Reply messages for the mappings they are serving.
- (3) There can be either a soft-state or hard-state security association between the xTRs and the Map-Servers.

The distribution of xTR-IDs (and Site-IDs) and the management of security associations are out of the scope of this document.

5. Map-Request Additions

Figure 1 shows the format of the updated Map-Request [I-D.ietf-lisp-rfc6833bis] to support the PubSub functionality.

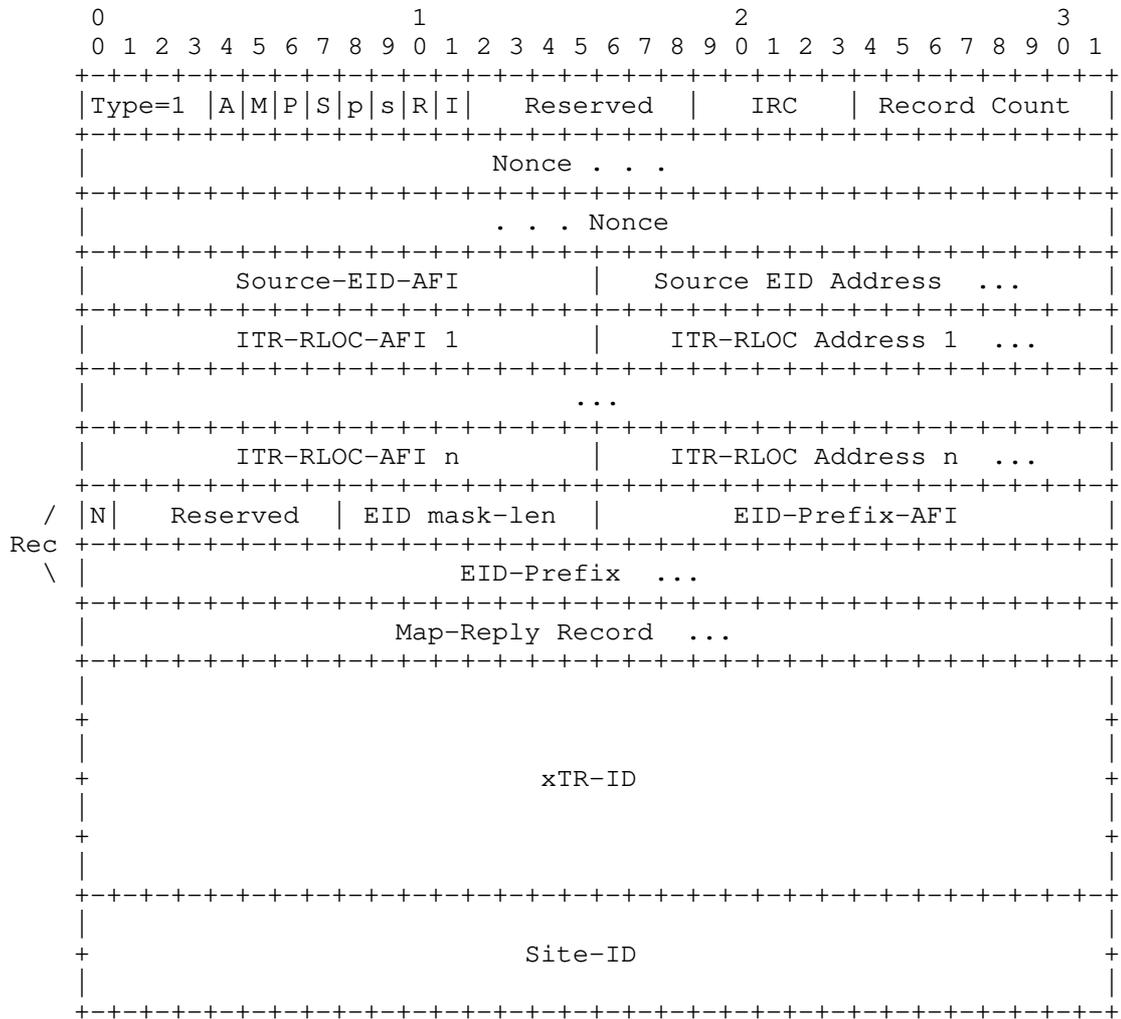


Figure 1: Map-Request with I-bit, N-bit, xTR-ID and Site-ID

The following is added to the Map-Request message defined in [I-D.ietf-lisp-rfc6833bis]:

xTR-ID bit (I-bit): The I-bit of a Map-Request message is set to 1 to indicate that a 128 bit xTR-ID and a 64 bit Site-ID fields are present at the end of the Map-Request message. If an xTR is configured with an xTR-ID or Site-ID, it MUST set the I bit to 1 and include its xTR-ID and Site-ID in the Map-Request messages it generates. If either the xTR-ID or Site-ID is not configured an

unspecified value is encoded for whichever ID that is not configured.

Notification-Requested bit (N-bit): The N-bit of an EID-record is set to 1 to specify that the xTR wants to be notified of updates for that mapping record.

xTR-ID field: xTR-ID is a 128 bit field at the end of the Map-Request message, starting after the final Record in the message (or the Map-Reply Record, if present). The xTR-ID is used to uniquely identify the sender of a Map-Request message, especially in the case where a site has more than one xTR. A value of all zeros indicate that an xTR-ID is not specified, though encoded in the message. This is useful in the case where a Site-ID is specified, but no xTR-ID is configured.

Site-ID field: Site-ID is a 64 bit field at the end of the Map-Request message, following the xTR-ID. Site-ID is used by the Map-Server receiving the Map-Request message to identify which xTRs belong to the same site. A value of 0 indicate that a Site-ID is not specified, though encoded in the message.

6. Mapping Request Subscribe Procedures

The xTR subscribes for RLOC-set changes for a given EID-prefix by sending a Map-Request to the Mapping System with the N-bit set on the EID-Record. The xTR builds a Map-Request according to [I-D.ietf-lisp-rfc6833bis] but also does the following:

- (1) The xTR MUST set the I-bit of the Map-Request message to 1 and append its xTR-ID and Site-ID to the Map-Request. The xTR-ID uniquely identifies the xTR.
- (2) The xTR MUST set the N-bit to 1 for each EID-Record to which the xTR wants to subscribe.

The Map-Request is forwarded to the appropriate Map-Server through the Mapping System. This document does not assume that a Map-Server is pre-assigned to handle the subscription state for a given xTR. The Map-Server that receives the Map-Request will be the Map-Server responsible to notify that specific xTR about future mapping changes for the subscribed mapping records.

Upon reception of the Map-Request, the Map-Server processes it as described in [I-D.ietf-lisp-rfc6833bis]. Upon processing, for each EID-Record that has the N-bit set to 1, the Map-Server proceeds adding the xTR-ID contained in the Map-Request to the list of xTR that have requested to be subscribed to that mapping record.

If the xTR-ID is added to the list, the Map-Server MUST send a Map-Notify message back to the xTR to acknowledge the successful subscription. The Map-Server MUST follow the specification in Section 6.1.7 of [I-D.ietf-lisp-rfc6833bis] to build the Map-Notify with the following considerations.

- (1) The Map-Server MUST use the nonce from the Map-Request as the nonce for the Map-Notify.
- (2) The Map-Server MUST use its security association with the xTR (see Section 4) to compute the authentication data of the Map-Notify.
- (3) The Map-Server MUST send the Map-Notify to one of the ITR-RLOCs received in the Map-Request.

When the xTR receives a Map-Notify with a nonce that matches one in the list of outstanding Map-Request messages sent with an N-bit set, it knows that the Map-Notify is to acknowledge a successful subscription. The xTR processes this Map-Notify as described in [I-D.ietf-lisp-rfc6833bis] with the following considerations. The xTR MUST use its security association with the Map-Server (see Section 4) to validate the authentication data on the Map-Notify. The xTR MUST use the Map-Notify to populate its map-cache with the returned EID-prefix and RLOC-set.

The subscription of an xTR-ID to the list of subscribers for the EID-Record may fail for a number of reasons. For example, because of local configuration policies (such as white/black lists of subscribers), or because the Map-Server has exhausted the resources to dedicate to the subscription of that EID-Record (e.g., the number of subscribers excess the capacity of the Map-Server).

If the subscription fails, the Map-Server MUST send a Map-Reply to the originator of the Map-Request, as described in [I-D.ietf-lisp-rfc6833bis]. This is also the case when the Map-Server does not support PubSub operation. The xTR processes the Map-Reply as specified in [I-D.ietf-lisp-rfc6833bis].

If an xTR-ID is successfully added to the list of subscribers for an EID-Record, the Map-Server MUST extract the ITR-RLOCs present in the Map-Request, and store the association between the xTR-ID and those RLOCs. Any already present state regarding ITR-RLOCs for the same xTR-ID MUST be overwritten.

If the Map-Request only has one ITR-RLOC with AFI = 0 (i.e. Unknown Address), the Map-Server MUST remove the subscription state for that xTR-ID. In this case, the Map-Server MUST send the Map-Notify to the

source RLOC of the Map-Request. When the TTL for the EID-record expires, the EID-prefix is removed from the Map-Server's subscription cache. On EID-Record removal, the Map-Server notifies the subscribers via a Map-Notify with TTL equal 0.

7. Mapping Notification Publish Procedures

The publish procedure is implemented via Map-Notify messages that the Map-Server sends to xTRs. The xTRs acknowledge the reception of Map-Notifies via sending Map-Notify-Ack messages back to the Map-Server. The complete mechanism works as follows.

When a mapping stored in a Map-Server is updated (e.g. via a Map-Register from an ETR), the Map-Server MUST notify the subscribers of that mapping via sending Map-Notify messages with the most updated mapping information. The Map-Notify message sent to each of the subscribers as a result of an update event MUST follow the exact encoding and logic defined in [I-D.ietf-lisp-rfc6833bis] for Map-Notify, except for the following:

- (1) The Map-Notify MUST be sent to one of the ITR-RLOCs associated with the xTR-ID of the subscriber.
- (2) The nonce of the Map-Notify MUST be the one the subscriber sent in the Map-Request. If the subscriber sent no Map-Request (e.g. was subscribed via configuration at the Map-Server) the nonce MUST be randomly generated by the Map-Server.
- (3) The Map-Server MUST use its security association with the xTR to compute the authentication data of the Map-Notify.

When the xTR receives a Map-Notify with a nonce sent previously in a Map-Request, or with a nonce not present in any list of previously sent nonces but with an EID not local to the xTR, the xTR knows that the Map-Notify has been received to update an entry on its map-cache. Processing of unsolicited Map-Notify messages MUST be explicitly enabled via configuration at the xTR.

The xTR processes the received Map-Notify as specified in [I-D.ietf-lisp-rfc6833bis], with the following considerations. The xTR MUST use its security association with the Map-Server (see Section 4) to validate the authentication data on the Map-Notify. The xTR MUST use the mapping information carried in the Map-Notify to update its internal map-cache. The xTR MUST acknowledge the Map-Notify by sending back a Map-Notify-Ack (specified in [I-D.ietf-lisp-rfc6833bis]), with the nonce from the Map-Notify, to the Map-Server. If after a configurable timeout, the Map-Server has

not received back the Map-Notify-Ack, it CAN try to send the Map-Notify to a different ITR-RLOC for that xTR-ID.

8. Security Considerations

The way to provide a security association between the ITRs and the Map-Servers must be evaluated according to the size of the deployment. For small deployments, it is possible to have a shared key (or set of keys) between the ITRs and the Map-Servers. For larger and Internet-scale deployments, scalability is a concern and further study is needed.

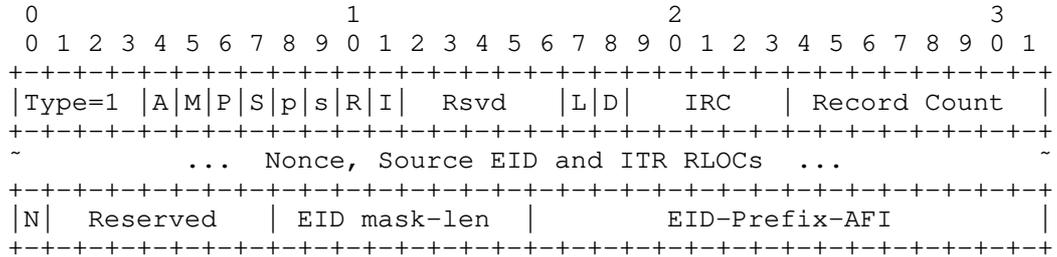
9. Acknowledgments

This work is partly funded by the ANR LISP-Lab project #ANR-13-INFR-009 (<https://lisplab.lip6.fr>).

10. IANA Considerations

This document is requesting bit allocations in the Map-Request message. The registry is introduced in [I-D.ietf-lisp-rfc6833bis] and named "LISP Control Plane Header Bits". This document is adding bits to the sub-registry "Map-Request Header Bits".

Sub-Registry: Map-Request Header Bits:



Spec Name	IANA Name	Bit Position	Description
I	map-request-I	11	xTR-ID Bit
N	map-request-N	... + 0	Notification-Requested Bit

LISP Map-Request Header Bits

11. Normative References

- [I-D.ietf-lisp-rfc6833bis]
Fuller, V., Farinacci, D., and A. Cabellos-Aparicio,
"Locator/ID Separation Protocol (LISP) Control-Plane",
draft-ietf-lisp-rfc6833bis-17 (work in progress), October
2018.
- [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>>.

Authors' Addresses

Alberto Rodriguez-Natal
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: natal@cisco.com

Vina Ermagan
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: vermagan@cisco.com

Johnson Leong
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: joleong@cisco.com

Fabio Maino
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: fmaino@cisco.com

Albert Cabellos-Aparicio
Technical University of Catalonia
Barcelona
Spain

Email: acabello@ac.upc.edu

Sharon Barkai
Fermi Serverless
CA
USA

Email: sharon@fermicloud.io

Dino Farinacci
lisppers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes 35000
France

Email: christian.jacquenet@orange.com

Stefano Secci
LIP6 UPMC
France

Email: stefano.secci@lip6.fr

LISP Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 September 2023

A. Rodriguez-Natal
Cisco
V. Ermagan
Google
A. Cabellos
UPC/BarcelonaTech
S. Barkai
Nexar
M. Boucadair
Orange
28 February 2023

Publish/Subscribe Functionality for the Locator/ID Separation Protocol
(LISP)
draft-ietf-lisp-pubsub-15

Abstract

This document specifies an extension to the request/reply based Locator/ID Separation Protocol (LISP) control plane to enable Publish/Subscribe (PubSub) operation.

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 1 September 2023.

Copyright Notice

Copyright (c) 2023 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
1.1. Scope of Applicability	4
2. Terminology and Requirements Notation	4
3. Deployment Requirements	4
4. Map-Request PubSub Additions	5
5. Mapping Request Subscribe Procedures	6
6. Mapping Notification Publish Procedures	11
7. Security Considerations	12
7.1. Security Association between ITR and Map-Server	12
7.2. DDoS Attack Mitigation	13
8. IANA Considerations	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15
Appendix A. Sample PubSub Deployment Experiences	17
A.1. PubSub as a Monitoring Tool	17
A.2. Mitigating Negative Map-Cache Entries	17
A.3. Improved Mobility Latency	18
A.4. Enhanced Reachability with Dynamic Redistribution of Prefixes	18
A.5. Better Serviceability	19
Acknowledgments	19
Contributors	20
Authors' Addresses	21

1. Introduction

The Locator/ID Separation Protocol (LISP) [RFC9300] [RFC9301] splits IP addresses into two different namespaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). LISP uses a map and encapsulate (a.k.a., map-and-encap) approach that relies on (1) a Mapping System (basically a distributed database) that stores and disseminates EID-RLOC mappings and on (2) LISP tunnel routers (xTRs) that encapsulate and decapsulate data packets based on the content of those mappings.

Ingress Tunnel Routers (ITRs) / Re-encapsulating Tunnel Routers (RTRs) / Proxy Ingress Tunnel Routers (PITRs) pull EID-to-RLOC mapping information from the Mapping System by means of an explicit request message. Section 6.1 of [RFC9301] indicates how Egress Tunnel Routers (ETRs) can tell ITRs/RTRs/PITRs about mapping changes.

This document presents a Publish/Subscribe (PubSub) extension in which the Mapping System can notify ITRs/RTRs/PITRs about mapping changes. When this mechanism is used, mapping changes can be notified faster and can be managed in the Mapping System versus the LISP sites.

In general, when an ITR/RTR/PITR wants to be notified for mapping changes for a given EID-Prefix, the following main steps occur:

- (1) The ITR/RTR/PITR builds a Map-Request for that EID-Prefix with the Notification-Requested bit (N-bit) set and which also includes its xTR-ID and Site-ID.
- (2) The Map-Request is forwarded to one of the Map-Servers that the EID-Prefix is registered to.
- (3) The Map-Server creates subscription state for the ITR/RTR/PITR on the EID-Prefix.
- (4) The Map-Server sends a Map-Notify to the ITR/RTR/PITR to confirm that the subscription has been created and then waits for an acknowledgement of the notification.
- (5) The ITR/RTR/PITR sends back a Map-Notify-Ack to acknowledge the successful receipt of the Map-Notify.
- (6) When there is a change in the mapping of the EID-Prefix, the Map-Server sends a Map-Notify message to each ITR/RTR/PITR in the subscription list.
- (7) Each ITR/RTR/PITR sends a Map-Notify-Ack to acknowledge the received Map-Notify.

This operation is repeated for all EID-Prefixes for which ITRs/RTRs/PITRs want to be notified. An ITR/RTR/PITR can set the N-bit for several EID-Prefixes within a single Map-Request. Please note that the steps above illustrate only the simplest scenario and that details for this and other scenarios are described later in the document.

The reader may refer to [I-D.boucadair-lisp-pubsub-flow-examples] for sample flows to illustrate the use of the PubSub specification.

1.1. Scope of Applicability

The PubSub procedure specified in this document is intended to be used in contexts with controlled access to the Map-Server. How a deployment controls access to a Map-Server is deployment specific, and therefore out of the scope of this document. However, the Map-Resolvers and Map-Servers need to be configured with the required information to at least ensure the following:

- (1) Map-Resolvers MUST verify that an xTR is allowed to (1) set the N-bit to 1 and (2) use the xTR-ID, Site-ID, and ITR-RLOCs included in a Map-Request.
- (2) Map-Servers MUST only accept subscription requests from Map-Resolvers that verify Map-Requests as previously described.

2. Terminology and Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The document uses the terms defined in Section 3 of [RFC9300].

3. Deployment Requirements

In addition to the general assumptions and expectations that [RFC9301] makes for LISP deployments, this document makes the following deployment requirements:

- (1) A unique 128-bit xTR-ID (plus a 64-bit Site-ID) identifier is assigned to each xTR.
- (2) Map-Servers are configured to proxy Map-Replying (i.e., they are solicited to generate and send Map-Reply messages) for the mappings they are serving.
- (3) A security association (e.g., a PubSubKey) is required between the ITRs and the Map-Servers (see Section 7.1).

If a requirement is not met, a subscription cannot be established, and the network will continue operating without this enhancement. The configuration of xTR-IDs and Site-IDs is out of the scope of this document. The reader may refer to [I-D.ietf-lisp-yang] for an example of how these identifiers can be provisioned to LISP nodes.

4. Map-Request PubSub Additions

Figure 1 shows the format of the updated Map-Request to support the PubSub functionality. In particular, this document associates a meaning with one of the reserved bits (see Section 8).

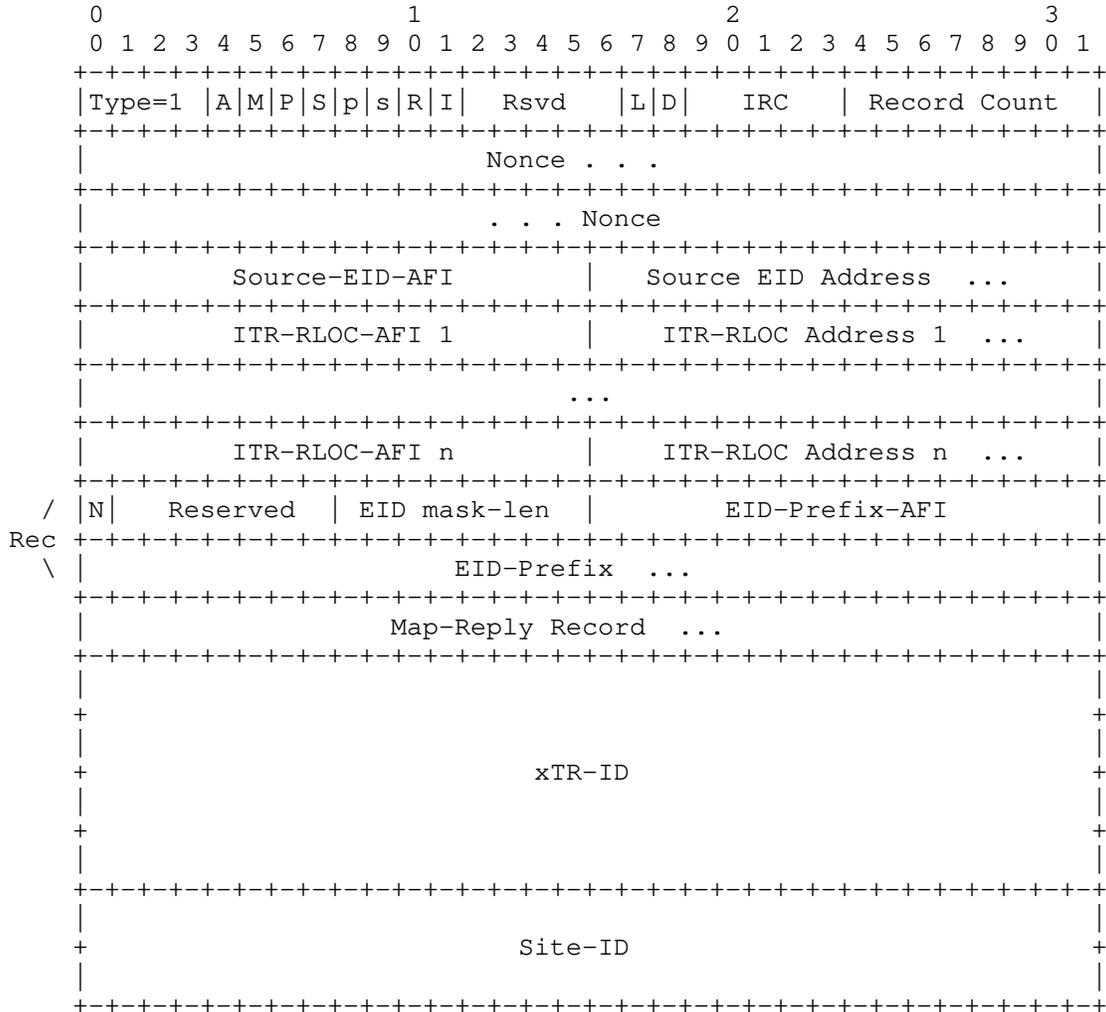


Figure 1: Map-Request with I-bit, N-bit, xTR-ID, and Site-ID

The following is added to the Map-Request message defined in Section 5.2 of [RFC9301]:

xTR-ID bit (I-bit): This bit is set to 1 to indicate that 128-bit

xTR-ID and 64-bit Site-ID fields are present in the Map-Request message. For PubSub operation, an xTR MUST be configured with an xTR-ID and Site-ID, and it MUST set the I-bit to 1 and include its xTR-ID and Site-ID in the Map-Request messages it generates. If the I-bit is set, but the Site-ID and/or xTR-ID are not included, a receiver can detect the error because, after processing that last EID-record, there are no bytes left from processing the message. In this case, the receiver SHOULD log a malformed Map-Request and MUST drop the message.

Notification-Requested bit (N-bit): The N-bit of an EID-Record is set to 1 to specify that the xTR wants to be notified of updates for that EID-Prefix.

xTR-ID field: If the I-bit is set, this field is added to the Map-Request message as shown in Figure 1, starting right after the final Record in the message (or the Map-Reply Record, if present). The xTR-ID is specified in Section 5.6 of [RFC9301].

Site-ID field: If the I-bit is set, this field is added to the Map-Request message as shown in Figure 1, following the xTR-ID field. The Site-ID is defined in Section 5.6 of [RFC9301].

5. Mapping Request Subscribe Procedures

The xTR subscribes for changes, to a given EID-Prefix, by sending a Map-Request to the Mapping System with the N-bit set on the EID-Record. The xTR builds a Map-Request according to Section 5.3 of [RFC9301] but also does the following:

- (1) The xTR MUST set the I-bit to 1 and append its xTR-ID and Site-ID to the Map-Request.
- (2) The xTR MUST set the N-bit to 1 for the EID-Record to which the xTR wants to subscribe.
- (3) If the xTR has a nonce associated with the EID-Prefix, it MUST use this nonce increased by one in the Map-Request. Otherwise, it generates a nonce following Section 5.2 of [RFC9301]. It is RECOMMENDED that the xTR uses persistent storage to keep nonce state. If the xTR does not have persistent storage and does not have a nonce associated with the EID-Prefix, it MUST reset the nonce by using the procedure described in Section 7.1 to successfully create a new security association with the Map-Server.

The Map-Request is forwarded to the appropriate Map-Server through the Mapping System. This document does not assume that a Map-Server is pre-assigned to handle the subscription state for a given xTR. The Map-Server that receives the Map-Request will be the Map-Server responsible for notifying that specific xTR about future mapping changes for the subscribed mapping records.

Upon receipt of the Map-Request, the Map-Server processes it as described in Section 8.3 of [RFC9301]. In addition, unless the xTR is using the procedure described in Section 7.1 to create a new security association, the Map-Server MUST verify that the nonce in the Map-Request is greater than the stored nonce (if any) associated with the xTR-ID (and EID-Prefix, when applicable). Otherwise, the Map-Server MUST silently drop the Map-Request message and SHOULD log the event to record that a replay attack could have occurred. Furthermore, upon processing, for the EID-Record that has the N-bit set to 1, the Map-Server proceeds to add the xTR-ID contained in the Map-Request to the list of xTRs that have requested to be subscribed to that EID-Prefix.

If an xTR-ID is successfully added to the list of subscribers for an EID-Prefix, the Map-Server MUST extract the nonce and ITR-RLOCs present in the Map-Request, and store the association between the EID-Prefix, xTR-ID, ITR-RLOCs, and nonce. Any already present state regarding ITR-RLOCs and/or nonce for the same xTR-ID MUST be overwritten. When the LISP deployment has a single Map-Server, the Map-Server can be configured to keep a single nonce per xTR-ID for all EID-Prefixes (when used, this option MUST be enabled at the Map-Server and all xTRs).

If the xTR-ID is added to the list, the Map-Server MUST send a Map-Notify message back to the xTR to acknowledge the successful subscription. The Map-Server builds the Map-Notify according to Sections 5.5 and 5.7 of [RFC9301] with the following considerations:

- (1) The Map-Server MUST use the nonce from the Map-Request as the nonce for the Map-Notify.
- (2) The Map-Server MUST use its security association with the xTR (Section 7.1) to sign the authentication data of the Map-Notify. The xTR MUST use the security association to verify the received authentication data.
- (3) The Map-Server MUST send the Map-Notify to one of the ITR-RLOCs received in the Map-Request (which one is implementation specific).

As a reminder, the initial transmission and retransmission of Map-Notify messages by a Map-Server follow the procedure specified in Section 5.7 of [RFC9301]. Some state changes may trigger an overload that would impact, e.g., the outbound capacity of a Map-Server. A similar problem may be experienced when a large number of state entries are simultaneously updated. To prevent such phenomena, Map-Servers SHOULD be configured with policies to control the maximum number of subscriptions and also the pace of Map-Notify messages. For example, the Map-Server may be instructed to limit the resources that are dedicated to unsolicited Map-Notify messages to a small fraction (e.g., less than 10%) of its overall processing and forwarding capacity. The exact details to characterize such policies are deployment and implementation specific. Likewise, this document does not specify which notifications take precedence when these policies are enforced.

When the xTR receives a Map-Notify with a nonce that matches one in the list of outstanding Map-Request messages sent with an N-bit set, it knows that the Map-Notify is to acknowledge a successful subscription. The xTR processes this Map-Notify, as described in Section 5.7 of [RFC9301], and MUST use the Map-Notify to populate its Map-Cache with the returned EID-Prefix and RLOC-set. As a reminder, following Section 5.7 of [RFC9301], the xTR has to send a Map-Notify-Ack back to the Map-Server. If the Map-Server does not receive the Map-Notify-Ack after exhausting the Map-Notify retransmissions described in Section 5.7 of [RFC9301], the Map-Server can remove the subscription state. If the Map-Server removes the subscription state, and absent explicit policy, it SHOULD notify the xTR by sending a single Map-Notify with the same nonce but with Loc-Count = 0 (and Loc-AFI = 0), and ACT bits set to 5 "Drop/Auth-Failure". It is OPTIONAL for the xTR to update its map-cache entry for the EID-Prefix (if any) based on this Map-Notify. This message is specifically useful for cases where Map-Notifies are successfully received by an xTR but the corresponding Map-Notify-Acks were lost when forwarded to the Map-Server. xTR implementations can use this signal to try to reinstall their subscription state instead of maintaining stale mappings.

The subscription of an xTR-ID may fail for a number of reasons. For example, it fails because of local configuration policies (such as accept and drop lists of subscribers), because the Map-Server has exhausted the resources to dedicate to the subscription of that EID-Prefix (e.g., the number of subscribers excess the capacity of the Map-Server), or because the xTR tried but was not successful in establishing a new security association (Section 7.1).

If the subscription request fails, the Map-Server sends a Map-Reply to the originator of the Map-Request, as described in Section 8.3 of [RFC9301], with the following considerations:

- * If the subscription request fails due to policy (e.g. for explicitly configured subscriptions, as described later in this section) the Map-Server MUST respond to the Map-Request with a Negative Map-Reply (Loc-Count = 0 and Loc-AFI = 0) with ACT bits set to 4 "Drop/Policy-Denied".
- * If the subscription request fails due to authentication (e.g. when a new security association is being established, as described in Section 7.1), the Map-Server MUST respond to the Map-Request with a Negative Map-Reply (Loc-Count = 0 and Loc-AFI = 0) with ACT bits set to 5 "Drop/Auth-Failure".
- * If the subscription request fails due to any other reason, the Map-Server MUST follow Section 8.3 of [RFC9301] with no changes.

The xTR processes any (Negative) Map-Reply as specified in Section 8.1 of [RFC9301], with the following considerations: if the xTR receives a Negative Map-Reply with ACT bits set to 4 "Drop/Policy-Denied" or 5 "Drop/Auth-Failure" as a response to a subscription request, it is OPTIONAL for the xTR to update its map-cache entry for the EID-Prefix (if any) based on this Negative Map-Reply. If the subscription request fails (for whichever reason), it is up to the implementation of the xTR to try to subscribe again.

If the Map-Server receives a subscription request for an EID-Prefix not present in the mapping database, it SHOULD follow the same logic described in Section 8.4 of [RFC9301] and create a temporary subscription state for the xTR-ID to the least-specific prefix that both matches the original query and does not match any EID-Prefix known to exist in the LISP-capable infrastructure. Alternatively, the Map-Server can instead determine that such a subscription request fails, and send a Negative Map-Reply following Section 8.3 of [RFC9301]. In both cases, the TTL of the temporary subscription state or the Negative Map-Reply SHOULD be configurable, with a value of 15-minutes being RECOMMENDED.

The subscription state can also be created explicitly by configuration at the Map-Server (possible when a pre-shared security association exists, see Section 7) using a variety of means that are out of scope. If at the time the explicit subscription is configured there is no nonce that can be used for the explicit subscription state (e.g., from a different subscription already established with the same xTR when a single nonce is kept per xTR-ID), then both the xTR and Map-Server MUST be configured with the initial nonce to use.

It is RECOMMENDED to have a configuration option to enable (or disable) the xTR to accept publication information for EID-Prefixes the xTR did not explicitly subscribe to. By default, the xTR is allowed to modify explicitly configured subscription state following the procedures described in this section, however this MAY be disabled at the Map-Server via configuration. If the Map-Server is instructed to not allow xTRs to modify explicitly configured subscriptions, and an xTR tries to do so, this triggers a Negative Map-Reply with ACT bits set to 4 "Drop/Policy-Denied" as described earlier in this section.

The following specifies the procedure to remove a subscription: If a valid Map-Request with the N-bit set to 1 only has one ITR-RLOC with AFI = 0 (i.e., Unknown Address), the Map-Server MUST remove the subscription state for that xTR-ID (unless this is disabled via configuration, see previous paragraph). If the subscription state is removed, the Map-Server MUST send a Map-Notify to the source RLOC of the Map-Request. If the subscription removal fails due to configuration, this triggers a Negative Map-Reply with with ACT bits set to 4 "Drop/Policy-Denied" as described earlier in this section; the Map-Server sends the Negative Map-Reply to the source RLOC of the Map-Request in this case. Removing subscription state at the Map-Server can lead to replay attacks. To soften this, the Map-Server SHOULD keep the last nonce seen per xTR-ID (and EID-Prefix, when applicable). If the Map-Server does not keep last nonces seen, then the Map-Server MUST require the xTRs to subscribe using the procedure described in Section 7.1 to create a new security association with the Map-Server.

If the Map-Server receives a Map-Request asking to remove a subscription for an EID-Prefix without subscription state for that xTR-ID, but covered by a less-specific EID-Prefix for which subscription state exists for the xTR-ID, the Map-Server SHOULD stop publishing updates about this more-specific EID-Prefix to that xTR, until the xTR subscribes to the more-specific EID-Prefix. The same considerations regarding authentication, integrity protection, and nonce checks described in this section and Section 7 for Map-Requests used to update subscription state, apply for Map-Requests used to remove subscription state.

When an EID-Prefix is removed from the Map-Server (either when explicitly withdrawn or when its TTL expires), the Map-Server notifies its subscribers (if any) via a Map-Notify with TTL equal 0.

6. Mapping Notification Publish Procedures

The publish procedure is implemented via Map-Notify messages that the Map-Server sends to xTRs. The xTRs acknowledge the reception of Map-Notifies via sending Map-Notify-Ack messages back to the Map-Server. The complete mechanism works as follows:

When a mapping stored in a Map-Server is updated (e.g., via a Map-Register from an ETR), the Map-Server MUST notify the subscribers of that mapping via sending Map-Notify messages with the most updated mapping information. If subscription state in the Map-Server exists for a less-specific EID-Prefix and a more-specific EID-Prefix is updated, then the Map-Notify is sent with the more-specific EID-Prefix mapping to the subscribers of the less-specific EID-Prefix mapping. The Map-Notify message sent to each of the subscribers as a result of an update event follows the encoding and logic defined in Section 5.7 of [RFC9301] for Map-Notify, except for the following:

- (1) The Map-Notify MUST be sent to one of the ITR-RLOCs associated with the xTR-ID of the subscriber (which one is implementation specific).
- (2) The Map-Server increments the nonce by one every time it sends a Map-Notify as publication to an xTR-ID for a particular EID-Prefix.
- (3) The Map-Server MUST use its security association with the xTR to compute the authentication data of the Map-Notify.

When the xTR receives a Map-Notify with an EID not local to the xTR, the xTR knows that the Map-Notify has been received to update an entry on its Map-Cache. The xTR MUST keep track of the last nonce seen in a Map-Notify received as a publication from the Map-Server for the EID-Prefix. When the LISP deployment has a single Map-Server, the xTR can be configured to keep track of a single nonce for all EID-Prefix (when used, this option MUST be enabled at the Map-Server and all xTRs). If a Map-Notify received as a publication has a nonce value that is not greater than the saved nonce, the xTR drops the Map-Notify message and logs the fact a replay attack could have occurred. The same considerations discussed in Section 5.6 of [RFC9301] regarding Map-Register nonces apply here for Map-Notify nonces.

The xTR processes the received Map-Notify as specified in Section 5.7 of [RFC9301], with the following considerations: The xTR MUST use its security association with the Map-Server (Section 7.1) to validate the authentication data on the Map-Notify. The xTR MUST use the mapping information carried in the Map-Notify to update its internal

Map-Cache. If after following Section 5.7 of [RFC9301] regarding retransmission of Map-Notify messages, the Map-Server has not received back the Map-Notify-Ack, it can try to send the Map-Notify to a different ITR-RLOC for that xTR-ID. If the Map-Server tries all the ITR-RLOCs without receiving a response, it may stop trying to send the Map-Notify.

7. Security Considerations

Generic security considerations related to LISP control messages are discussed in Section 9 of [RFC9301].

In the particular case of PubSub, cache poisoning via malicious Map-Notify messages is avoided by the use of nonce and the security association between the ITRs and the Map-Servers.

It is RECOMMENDED to follow guidance from the last paragraph of Section 9 of [RFC9301] to ensure integrity protection of Map-Request messages (e.g., to prevent xTR-ID hijacking).

7.1. Security Association between ITR and Map-Server

Since Map-Notifies from the Map-Server to the ITR need to be authenticated, there is a need for a soft-state or hard-state security association (e.g., a PubSubKey) between the ITRs and the Map-Servers. For some controlled deployments, it might be possible to have a shared PubSubKey (or set of keys) between the ITRs and the Map-Servers. However, if pre-shared keys are not used in the deployment, LISP-SEC [RFC9303] can be used as follows to create a security association between the ITR and the Map-Server.

First, when the ITR is sending a Map-Request with the N-bit set following Section 5, the ITR also performs the steps described in Section 6.4 of [RFC9303]. The ITR can then generate a PubSubKey by deriving a key from the One-Time Key (OTK) and Map-Request's nonce as follows: $\text{PubSubKey} = \text{KDF}(\text{OTK} + \text{nonce})$, where KDF is the Key Derivation Function indicated by the OTK Wrapping ID. If OTK Wrapping ID equals NULL-KEY-WRAP-128, then the PubSubKey is the OTK. Note that as opposed to the pre-shared PubSubKey, this generated PubSubKey is different per EID-Prefix to which an ITR subscribes (since the ITR will use a different OTK per Map-Request).

When the Map-Server receives the Map-Request it follows the procedure specified in Section 5 with the following considerations: The Map-Server MUST verify that the OTK has not been used before. If the Map-Server verifies the OTK and cannot determine that the OTK has not been used before, the subscription request fails due to authentication and this triggers a Negative Map-Reply with ACT bits

set to 5 "Drop/Auth-Failure", as described in Section 5. The xTR might try again with a different OTK upon reception of this Negative Map-Reply. Note that a Map-Server implementation might decide to not keep full past OTKs and instead use some form of hash. In that case, hash collisions are handled as if the OTK has been reused. Such an implementation needs to balance the hash length with the rate of collisions expected for the particular deployment; this is implementation specific. If the Map-Server has to reply with a Map-Reply for any other reason (e.g., if PubSub is not supported or a subscription is not accepted), then it follows normal LISP-SEC procedure described in Section 5.7 of [RFC9303]. No PubSubKey, security association, or subscription state is created when the Map-Server responds with any Map-Reply message.

Otherwise, if the Map-Server has to reply with a Map-Notify (e.g., due to subscription accepted) to a received Map-Request, the following extra steps take place:

- * The Map-Server extracts the OTK and OTK Wrapping ID from the LISP-SEC ECM Authentication Data.
- * The Map-Server generates a PubSubKey by deriving a key from the OTK as described before for the ITR. This is the same PubSubKey derived at the ITR which is used to establish a security association between the ITR and the Map-Server.
- * The PubSubKey can now be used to sign and authenticate any Map-Notify between the Map-Server and the ITR for the subscribed EID-Prefix. This includes the Map-Notify sent as a confirmation to the subscription. When the ITR wants to update the security association for that Map-Server and EID-Prefix, it once again follows the procedure described in this section.

Note that if the Map-Server replies with a Map-Notify, none of the regular LISP-SEC steps regarding Map-Reply described in Section 5.7 of [RFC9303] occur.

7.2. DDoS Attack Mitigation

If PubSub is deployed under the scope of applicability defined in Section 1.1 only known nodes can participate on the PubSub deployment. DDoS attacks based on replayed messages by unknown nodes are avoided by the use of nonce and the security association between the ITRs and the Map-Servers. Misbehaving known nodes may send massive subscription requests which may lead to exhausting the resources of a Map-Server. Furthermore, frequently changing the state of a subscription may also be considered as an attack vector. To mitigate such issues, Section 5.3 of [RFC9301] discusses rate-

limiting Map-Requests and Section 5.7 of [RFC9301] discusses rate-limiting Map-Notifies. Note that when the Map-Notify rate-limit threshold is met for a particular xTR-ID, the Map-Server will discard additional subscription requests from that xTR-ID and will fall back to [RFC9301] behavior when receiving a Map-Request from that xTR-ID (i.e., the Map-Server will send a Map-Reply).

8. IANA Considerations

This document requests IANA to assign a new bit from the "LISP Control Plane Header Bits: Map-Request" sub-registry under the "Locator/ID Separation Protocol (LISP) Parameters" registry available at [IANA-LISP]. The suggested position of this bit in the Map-Request message can be found in Figure 1.

Spec Name	IANA Name	Bit Position	Description	Reference
I	Map-Request-I	11	xTR-ID Bit	This-Document

Table 1: Additions to the Map-Request Header Bits Sub-Registry

This document also requests the creation of a new sub-registry entitled "LISP Control Plane Header Bits: Map-Request-Record" under the "Locator/ID Separation Protocol (LISP) Parameters" registry available at [IANA-LISP].

The initial content of this sub-registry is shown in Table 2:

Spec Name	IANA Name	Bit Position	Description	Reference
N	Map-Request-N	1	Notification-Requested Bit	This-Document

Table 2: Initial Content of LISP Control Plane Header Bits: Map-Request-Record Sub-Registry

The remaining bits (i.e., Bit positions 2-8) are Unassigned.

The policy for allocating new bits from this sub-registry is Specification Required (Section 4.6 of [RFC8126]).

Review requests are evaluated on the advice of one or more designated experts. Criteria that should be applied by the designated experts include determining whether the proposed registration duplicates existing entries and whether the registration description is sufficiently detailed and fits the purpose of this registry. These criteria are considered in addition to those already provided in Section 4.6 of [RFC8126] (e.g., the proposed registration must be documented in a permanent and readily available public specification). The designated experts will either approve or deny the registration request, communicating this decision to IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the request successful.

9. References

9.1. 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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC9303] Maino, F., Ermagan, V., Cabellos, A., and D. Saucez, "Locator/ID Separation Protocol Security (LISP-SEC)", RFC 9303, DOI 10.17487/RFC9303, October 2022, <<https://www.rfc-editor.org/info/rfc9303>>.

9.2. Informative References

- [I-D.boucadair-lisp-pubsub-flow-examples]
Boucadair, M., "LISP PubSub Flow Examples", Work in Progress, Internet-Draft, draft-boucadair-lisp-pubsub-flow-examples-03, 10 February 2023, <<https://datatracker.ietf.org/doc/html/draft-boucadair-lisp-pubsub-flow-examples-03>>.
- [I-D.haindl-lisp-gb-atn]
Haindl, B., Lindner, M., Moreno, V., Portoles-Comeras, M., Maino, F., and B. Venkatachalapathy, "Ground-Based LISP for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-haindl-lisp-gb-atn-08, 23 September 2022, <<https://datatracker.ietf.org/doc/html/draft-haindl-lisp-gb-atn-08>>.
- [I-D.ietf-lisp-eid-mobility]
Portoles-Comeras, M., Ashtaputre, V., Maino, F., Moreno, V., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-eid-mobility-11, 10 January 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-eid-mobility-11>>.
- [I-D.ietf-lisp-yang]
Ermagan, V., Rodriguez-Natal, A., Coras, F., Moberg, C., Rahman, R., Cabellos-Aparicio, A., and F. Maino, "LISP YANG Model", Work in Progress, Internet-Draft, draft-ietf-lisp-yang-18, 29 August 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-yang-18>>.
- [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>>.
- [IANA-LISP]
IANA, "Locator/ID Separation Protocol (LISP) Parameters", <<https://www.iana.org/assignments/lisp-parameters/lisp-parameters.xhtml>>.
- [RFC6835] Farinacci, D. and D. Meyer, "The Locator/ID Separation Protocol Internet Groper (LIG)", RFC 6835, DOI 10.17487/RFC6835, January 2013, <<https://www.rfc-editor.org/info/rfc6835>>.

Appendix A. Sample PubSub Deployment Experiences

Some LISP production networks have been running different forms of PubSub for some time. The following subsections provide an inventory of some experience lessons from these deployments.

A.1. PubSub as a Monitoring Tool

Some LISP deployments are using PubSub as a way to monitor EID-Prefixes (particularly, EID-to-RLOC mappings). To that aim, some LISP implementations have extended the LISP Internet Groper (lig) [RFC6835] tool to use PubSub. Such an extension is meant to support an interactive mode with lig, and request subscription for the EID of interest. If there are RLOC changes, the Map-Server sends a notification and then the lig client displays that change to the user.

A.2. Mitigating Negative Map-Cache Entries

Section 8.1 of [RFC9301] suggests two TTL values for Negative Map-Replies: either 15-minute (if the EID-Prefix does not exist) or 1-minute (if the prefix exists but has not been registered). While these values are based on the original operational experience of the LISP protocol designers, negative cache entries have two unintended effects that were observed in production.

First, if the xTR keeps receiving traffic for a negative EID destination (i.e., an EID-Prefix with no RLOCs associated with it), it will try to resolve the destination again once the cached state expires, even if the state has not changed in the Map-Server. It was observed in production that this is happening often in networks that have a significant amount of traffic addressed for outside of the LISP network. This might result on excessive resolution signaling to keep retrieving the same state due to the cache expiring. PubSub is used to relax TTL values and cache negative mapping entries for longer periods of time, avoiding unnecessary refreshes of these forwarding entries, and drastically reducing signaling in these scenarios. In general, a TTL-based schema is a polling mechanism that leads to more signaling where PubSub provides an "event triggered mechanism" at the cost of state.

Second, if the state does indeed change in the Map-Server, updates based on TTL timeouts might prevent the cached state at the xTR from being updated until the TTL expires. This behavior was observed during configuration (or reconfiguration) periods on the network, where no-longer-negative EID-Prefixes do not receive the traffic yet due to stale Map-Cache entries present in the network. With the activation of PubSub, stale caches can be updated as soon as the state changes.

A.3. Improved Mobility Latency

An improved convergence time was observed on the presence of mobility events on LISP networks running PubSub as compared with running LISP [RFC9301]. As described in Section 4.1.2.1 of [I-D.ietf-lisp-eid-mobility], LISP can rely on data-driven Solicit-Map-Requests (SMRs) to ensure eventual network converge. Generally, PubSub offers faster convergence due to (1) no need to wait for a data triggered event and (2) less signaling as compared with the SMR-based flow. Note that when a Map-Server running PubSub has to update a large number of subscribers at once (i.e., when a popular mapping is updated) SMR based convergence may be faster for a small subset of the subscribers (those receiving PubSub updates last). Deployment experience reveals that data-driven SMRs and PubSub mechanisms complement each other and provide a fast and resilient network infrastructure in the presence of mobility events.

Furthermore, experience showed that not all LISP entities on the network need to implement PubSub for the network to get the benefits. In scenarios with significant traffic coming from outside of the LISP network, the experience showed that enabling PubSub in the border routers significantly improves mobility latency overall. Even if edge xTRs do not implement PubSub, and traffic is exchanged between EID-Prefixes at the edge, xTRs still converge based on data-driven events and SMR-triggered updates.

A.4. Enhanced Reachability with Dynamic Redistribution of Prefixes

There is a need to interconnect LISP networks with other networks that might or might not run LISP. Some of those scenarios are similar to the ones described in [I-D.haindl-lisp-gb-atn] and [I-D.moreno-lisp-uberlay]. When connecting LISP to other networks, the experience revealed that in many deployments the point of interaction with the other domains is not the Mapping System but rather the border router of the LISP site. For those cases the border router needs to be aware of the LISP prefixes to redistribute them to the other networks. Over the years different solutions have been used.

First, Map-Servers were collocated with the border routers, but this was hard to scale since border routers scale at a different pace than Map-Servers. Second, decoupled Map-Servers and border routers were used with static configuration of LISP entries on the border, which was problematic when modifications were made. Third, a routing protocol (e.g., BGP) can be used to redistribute LISP prefixes from the Map-Servers to a border router, but this comes with some implications, particularly the Map-Servers needs to implement an additional protocol which consumes resources and needs to be properly configured. Therefore, once PubSub was available, deployments started to adapt it to enable border routers to dynamically learn the prefixes they need to redistribute without the need of extra protocols or extra configuration on the network.

In other words, PubSub can be used to discover EID-Prefixes so they can be imported into other routing domains that do not use LISP. Similarly, PubSub can also be used to discover when EID-Prefixes need to be withdrawn from other routing domains. That is, in a typical deployment, a border router will withdraw an EID-Prefix it has been announcing to external routing domains, if it receives a notification that the RLOC-set for that EID-Prefix is now empty.

A.5. Better Serviceability

EID-to-RLOC mappings can have very long TTL, sometimes in the order of several hours. Upon the expiry of that TTL, the xTR checks if these entries are being used and removes any entry that is not being used. The problem with very long Map-Cache TTL is that (in the absence of PubSub) if a mapping changes, but it is not being used, the cache remains but it is stale. This is due to no data traffic being sent to the old location to trigger an SMR based Map-Cache update as described in Section 4.1.2.1 of [I-D.ietf-lisp-eid-mobility]. If the network operator runs a show command on a router to track the state of the Map-Cache, the router will display multiple entries waiting to expire but with stale RLOC information. This might be confusing for operators sometimes, particularly when they are debugging problems. With PubSub, the Map-Cache is updated with the correct RLOC information, even when it is not being used or waiting to expire, and this helps with debugging.

Acknowledgments

We would like to thank Marc Portoles, Balaji Venkatachalapathy, Bernhard Haindl, Luigi Iannone, and Padma Pillay-Esnault for their great suggestions and help regarding this document.

Many thanks to Alvaro Retana for the careful AD review.

Thanks to Chris M. Lonvick for the security directorate review, Al Morton for the OPS-DIR review, Roni Even for the Gen-ART review, Mike McBride for the rtg-dir review, Magnus Westerlund for the tsv directorate review, and Sheng Jiang for the int-dir review.

Thanks to John Scudder, Erik Kline, Lars Eggert, Warren Kumari, Martin Duke, Murray Kucherawy, Éric Vyncke, Robert Wilton, Zaheduzzaman Sarker, and Roman Danyliw for the IESG review.

This work was partly funded by the ANR LISP-Lab project #ANR-13-INFR-009 (<https://anr.fr/Projet-ANR-13-INFR-0009>).

Contributors

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Johnson Leong

Email: johnsonleong@gmail.com

Fabio Maino
Cisco
San Jose, CA
USA

Email: fmaino@cisco.com

Christian Jacquenet
Orange
Rennes
France

Email: christian.jacquenet@orange.com

Stefano Secci
Cnam
France

Email: stefano.secci@cnam.fr

Authors' Addresses

Alberto Rodriguez-Natal
Cisco
Barcelona
Spain
Email: natal@cisco.com

Vina Ermagan
Google
United States of America

Email: ermagan@gmail.com

Albert Cabellos
UPC/BarcelonaTech
Barcelona
Spain
Email: acabello@ac.upc.edu

Sharon Barkai
Nexar
Email: sharon.barkai@getnexar.com

Mohamed Boucadair
Orange
Rennes
France
Email: mohamed.boucadair@orange.com

Network Working Group
Internet-Draft
Obsoletes: 6830 (if approved)
Intended status: Standards Track
Expires: 8 November 2022

D. Farinacci
lispers.net
V. Fuller
vaf.net Internet Consulting
D. Meyer
1-4-5.net
D. Lewis
Cisco Systems
A. Cabellos (Ed.)
UPC/BarcelonaTech
7 May 2022

The Locator/ID Separation Protocol (LISP)
draft-ietf-lisp-rfc6830bis-38

Abstract

This document describes the Data-Plane protocol for the Locator/ID Separation Protocol (LISP). LISP defines two namespaces, End-point Identifiers (EIDs) that identify end-hosts and Routing Locators (RLOCs) that identify network attachment points. With this, LISP effectively separates control from data, and allows routers to create overlay networks. LISP-capable routers exchange encapsulated packets according to EID-to-RLOC mappings stored in a local Map-Cache.

LISP requires no change to either host protocol stacks or to underlay routers and offers Traffic Engineering, multihoming and mobility, among other features.

This document obsoletes RFC 6830.

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 8 November 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	3
1.1. Scope of Applicability	4
2. Requirements Notation	5
3. Definition of Terms	5
4. Basic Overview	9
4.1. Deployment on the Public Internet	11
4.2. Packet Flow Sequence	11
5. LISP Encapsulation Details	13
5.1. LISP IPv4-in-IPv4 Header Format	13
5.2. LISP IPv6-in-IPv6 Header Format	14
5.3. Tunnel Header Field Descriptions	15
6. LISP EID-to-RLOC Map-Cache	20
7. Dealing with Large Encapsulated Packets	20
7.1. A Stateless Solution to MTU Handling	21
7.2. A Stateful Solution to MTU Handling	22
8. Using Virtualization and Segmentation with LISP	23
9. Routing Locator Selection	23
10. Routing Locator Reachability	25
10.1. Echo Nonce Algorithm	27
11. EID Reachability within a LISP Site	28
12. Routing Locator Hashing	29
13. Changing the Contents of EID-to-RLOC Mappings	30
13.1. Locator-Status-Bits	30
13.2. Database Map-Versioning	31
14. Multicast Considerations	31
15. Router Performance Considerations	32
16. Security Considerations	33
17. Network Management Considerations	34
18. Changes since RFC 6830	34
19. IANA Considerations	35
19.1. LISP UDP Port Numbers	35
20. References	35

20.1. Normative References	35
20.2. Informative References	37
Appendix A. Acknowledgments	40
Appendix B. Document Change Log	41
B.1. Changes to draft-ietf-lisp-rfc6830bis-38	41
B.2. Changes to draft-ietf-lisp-rfc6830bis-37	41
B.3. Changes to draft-ietf-lisp-rfc6830bis-28	41
B.4. Changes to draft-ietf-lisp-rfc6830bis-27	42
B.5. Changes to draft-ietf-lisp-rfc6830bis-26	42
B.6. Changes to draft-ietf-lisp-rfc6830bis-25	42
B.7. Changes to draft-ietf-lisp-rfc6830bis-24	42
B.8. Changes to draft-ietf-lisp-rfc6830bis-23	42
B.9. Changes to draft-ietf-lisp-rfc6830bis-22	43
B.10. Changes to draft-ietf-lisp-rfc6830bis-21	43
B.11. Changes to draft-ietf-lisp-rfc6830bis-20	43
B.12. Changes to draft-ietf-lisp-rfc6830bis-19	43
B.13. Changes to draft-ietf-lisp-rfc6830bis-18	43
B.14. Changes to draft-ietf-lisp-rfc6830bis-17	43
B.15. Changes to draft-ietf-lisp-rfc6830bis-16	43
B.16. Changes to draft-ietf-lisp-rfc6830bis-15	44
B.17. Changes to draft-ietf-lisp-rfc6830bis-14	44
B.18. Changes to draft-ietf-lisp-rfc6830bis-13	44
B.19. Changes to draft-ietf-lisp-rfc6830bis-12	44
B.20. Changes to draft-ietf-lisp-rfc6830bis-11	44
B.21. Changes to draft-ietf-lisp-rfc6830bis-10	44
B.22. Changes to draft-ietf-lisp-rfc6830bis-09	45
B.23. Changes to draft-ietf-lisp-rfc6830bis-08	45
B.24. Changes to draft-ietf-lisp-rfc6830bis-07	45
B.25. Changes to draft-ietf-lisp-rfc6830bis-06	45
B.26. Changes to draft-ietf-lisp-rfc6830bis-05	46
B.27. Changes to draft-ietf-lisp-rfc6830bis-04	46
B.28. Changes to draft-ietf-lisp-rfc6830bis-03	46
B.29. Changes to draft-ietf-lisp-rfc6830bis-02	46
B.30. Changes to draft-ietf-lisp-rfc6830bis-01	47
B.31. Changes to draft-ietf-lisp-rfc6830bis-00	47
Authors' Addresses	47

1. Introduction

This document describes the Locator/Identifier Separation Protocol (LISP). LISP is an encapsulation protocol built around the fundamental idea of separating the topological location of a network attachment point from the node's identity [CHIAPPA]. As a result LISP creates two namespaces: Endpoint Identifiers (EIDs), that are used to identify end-hosts (e.g., nodes or Virtual Machines) and routable Routing Locators (RLOCs), used to identify network attachment points. LISP then defines functions for mapping between the two namespaces and for encapsulating traffic originated by

devices using non-routable EIDs for transport across a network infrastructure that routes and forwards using RLOCs. LISP encapsulation uses a dynamic form of tunneling where no static provisioning is required or necessary.

LISP is an overlay protocol that separates control from Data-Plane, this document specifies the Data-Plane as well as how LISP-capable routers (Tunnel Routers) exchange packets by encapsulating them to the appropriate location. Tunnel routers are equipped with a cache, called Map-Cache, that contains EID-to-RLOC mappings. The Map-Cache is populated using the LISP Control-Plane protocol [I-D.ietf-lisp-rfc6833bis].

LISP does not require changes to either the host protocol stack or to underlay routers. By separating the EID from the RLOC space, LISP offers native Traffic Engineering, multihoming and mobility, among other features.

Creation of LISP was initially motivated by discussions during the IAB-sponsored Routing and Addressing Workshop held in Amsterdam in October 2006 (see [RFC4984]).

This document specifies the LISP Data-Plane encapsulation and other LISP forwarding node functionality while [I-D.ietf-lisp-rfc6833bis] specifies the LISP control plane. LISP deployment guidelines can be found in [RFC7215] and [RFC6835] describes considerations for network operational management. Finally, [I-D.ietf-lisp-introduction] describes the LISP architecture.

This document obsoletes RFC 6830.

1.1. Scope of Applicability

LISP was originally developed to address the Internet-wide route scaling problem [RFC4984]. While there are a number of approaches of interest for that problem, as LISP as been developed and refined, a large number of other LISP uses have been found and are being used. As such, the design and development of LISP has changed so as to focus on these use cases. The common property of these uses is a large set of cooperating entities seeking to communicate over the public Internet or other large underlay IP infrastructures, while keeping the addressing and topology of the cooperating entities separate from the underlay and Internet topology, routing, and addressing.

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Definition of Terms

Address Family Identifier (AFI): AFI is a term used to describe an address encoding in a packet. An address family that pertains to addresses found in Data-Plane headers. See [AFN] and [RFC3232] for details. An AFI value of 0 used in this specification indicates an unspecified encoded address where the length of the address is 0 octets following the 16-bit AFI value of 0.

Anycast Address: Anycast Address refers to the same IPv4 or IPv6 address configured and used on multiple systems at the same time. An EID or RLOC can be an anycast address in each of their own address spaces.

Client-side: Client-side is a term used in this document to indicate a connection initiation attempt by an end-system represented by an EID.

Egress Tunnel Router (ETR): An ETR is a router that accepts an IP packet where the destination address in the "outer" IP header is one of its own RLOCs. The router strips the "outer" header and forwards the packet based on the next IP header found. In general, an ETR receives LISP-encapsulated IP packets from the Internet on one side and sends decapsulated IP packets to site end-systems on the other side. ETR functionality does not have to be limited to a router device. A server host can be the endpoint of a LISP tunnel as well.

EID-to-RLOC Database: The EID-to-RLOC Database is a distributed database that contains all known EID-Prefix-to-RLOC mappings. Each potential ETR typically contains a small piece of the database: the EID-to-RLOC mappings for the EID-Prefixes "behind" the router. These map to one of the router's own IP addresses that are routable on the underlay. Note that there MAY be transient conditions when the EID-Prefix for the LISP site and Locator-Set for each EID-Prefix may not be the same on all ETRs. This has no negative implications, since a partial set of Locators can be used.

EID-to-RLOC Map-Cache: The EID-to-RLOC Map-Cache is generally short-

lived, on-demand table in an ITR that stores, tracks, and is responsible for timing out and otherwise validating EID-to-RLOC mappings. This cache is distinct from the full "database" of EID-to-RLOC mappings; it is dynamic, local to the ITR(s), and relatively small, while the database is distributed, relatively static, and much more widely scoped to LISP nodes.

EID-Prefix: An EID-Prefix is a power-of-two block of EIDs that are allocated to a site by an address allocation authority. EID-Prefixes are associated with a set of RLOC addresses. EID-Prefix allocations can be broken up into smaller blocks when an RLOC set is to be associated with the larger EID-Prefix block.

End-System: An end-system is an IPv4 or IPv6 device that originates packets with a single IPv4 or IPv6 header. The end-system supplies an EID value for the destination address field of the IP header when communicating outside of its routing domain. An end-system can be a host computer, a switch or router device, or any network appliance.

Endpoint ID (EID): An EID is a 32-bit (for IPv4) or 128-bit (for IPv6) value that identifies a host. EIDs are generally only found in the source and destination address fields of the first (most inner) LISP header of a packet. The host obtains a destination EID the same way it obtains a destination address today, for example, through a Domain Name System (DNS) [RFC1034] lookup or Session Initiation Protocol (SIP) [RFC3261] exchange. The source EID is obtained via existing mechanisms used to set a host's "local" IP address. An EID used on the public Internet MUST have the same properties as any other IP address used in that manner; this means, among other things, that it MUST be unique. An EID is allocated to a host from an EID-Prefix block associated with the site where the host is located. An EID can be used by a host to refer to other hosts. Note that EID blocks MAY be assigned in a hierarchical manner, independent of the network topology, to facilitate scaling of the mapping database. In addition, an EID block assigned to a site MAY have site-local structure (subnetting) for routing within the site; this structure is not visible to the underlay routing system. In theory, the bit string that represents an EID for one device can represent an RLOC for a different device. When used in discussions with other Locator/ID separation proposals, a LISP EID will be called an "LEID". Throughout this document, any references to "EID" refer to an LEID.

Ingress Tunnel Router (ITR): An ITR is a router that resides in a

LISP site. Packets sent by sources inside of the LISP site to destinations outside of the site are candidates for encapsulation by the ITR. The ITR treats the IP destination address as an EID and performs an EID-to-RLOC mapping lookup. The router then prepends an "outer" IP header with one of its routable RLOCs (in the RLOC space) in the source address field and the result of the mapping lookup in the destination address field. Note that this destination RLOC may be an intermediate, proxy device that has better knowledge of the EID-to-RLOC mapping closer to the destination EID. In general, an ITR receives IP packets from site end-systems on one side and sends LISP-encapsulated IP packets toward the Internet on the other side.

LISP Header: LISP header is a term used in this document to refer to the outer IPv4 or IPv6 header, a UDP header, and a LISP-specific 8-octet header that follow the UDP header and that an ITR prepends or an ETR strips.

LISP Router: A LISP router is a router that performs the functions of any or all of the following: ITR, ETR, RTR, Proxy-ITR (PITR), or Proxy-ETR (PETR).

LISP Site: LISP site is a set of routers in an edge network that are under a single technical administration. LISP routers that reside in the edge network are the demarcation points to separate the edge network from the core network.

Locator-Status-Bits (LSBs): Locator-Status-Bits are present in the LISP header. They are used by ITRs to inform ETRs about the up/down status of all ETRs at the local site. These bits are used as a hint to convey up/down router status and not path reachability status. The LSBs can be verified by use of one of the Locator reachability algorithms described in Section 10. An ETR MUST rate-limit the action it takes when it detects changes in the Locator-Status-Bits.

Proxy-ETR (PETR): A PETR is defined and described in [RFC6832]. A PETR acts like an ETR but does so on behalf of LISP sites that send packets to destinations at non-LISP sites.

Proxy-ITR (PITR): A PITR is defined and described in [RFC6832]. A PITR acts like an ITR but does so on behalf of non-LISP sites that send packets to destinations at LISP sites.

Recursive Tunneling: Recursive Tunneling occurs when a packet has

more than one LISP IP header. Additional layers of tunneling MAY be employed to implement Traffic Engineering or other re-routing as needed. When this is done, an additional "outer" LISP header is added, and the original RLOCs are preserved in the "inner" header.

Re-Encapsulating Tunneling Router (RTR): An RTR acts like an ETR to remove a LISP header, then acts as an ITR to prepend a new LISP header. This is known as Re-encapsulating Tunneling. Doing this allows a packet to be re-routed by the RTR without adding the overhead of additional tunnel headers. When using multiple mapping database systems, care must be taken to not create re-encapsulation loops through misconfiguration.

Route-Returnability: Route-returnability is an assumption that the underlying routing system will deliver packets to the destination. When combined with a nonce that is provided by a sender and returned by a receiver, this limits off-path data insertion. A route-returnability check is verified when a message is sent with a nonce, another message is returned with the same nonce, and the destination of the original message appears as the source of the returned message.

Routing Locator (RLOC): An RLOC is an IPv4 [RFC0791] or IPv6 [RFC8200] address of an Egress Tunnel Router (ETR). An RLOC is the output of an EID-to-RLOC mapping lookup. An EID maps to zero or more RLOCs. Typically, RLOCs are numbered from blocks that are assigned to a site at each point to which it attaches to the underlay network; where the topology is defined by the connectivity of provider networks. Multiple RLOCs can be assigned to the same ETR device or to multiple ETR devices at a site.

Server-side: Server-side is a term used in this document to indicate that a connection initiation attempt is being accepted for a destination EID.

xTR: An xTR is a reference to an ITR or ETR when direction of data flow is not part of the context description. "xTR" refers to the router that is the tunnel endpoint and is used synonymously with the term "Tunnel Router". For example, "An xTR can be located at the Customer Edge (CE) router" indicates both ITR and ETR functionality at the CE router.

4. Basic Overview

One key concept of LISP is that end-systems operate the same way they do today. The IP addresses that hosts use for tracking sockets and connections, and for sending and receiving packets, do not change. In LISP terminology, these IP addresses are called Endpoint Identifiers (EIDs).

Routers continue to forward packets based on IP destination addresses. When a packet is LISP encapsulated, these addresses are referred to as Routing Locators (RLOCs). Most routers along a path between two hosts will not change; they continue to perform routing/forwarding lookups on the destination addresses. For routers between the source host and the ITR as well as routers from the ETR to the destination host, the destination address is an EID. For the routers between the ITR and the ETR, the destination address is an RLOC.

Another key LISP concept is the "Tunnel Router". A Tunnel Router prepends LISP headers on host-originated packets and strips them prior to final delivery to their destination. The IP addresses in this "outer header" are RLOCs. During end-to-end packet exchange between two Internet hosts, an ITR prepends a new LISP header to each packet, and an ETR strips the new header. The ITR performs EID-to-RLOC lookups to determine the routing path to the ETR, which has the RLOC as one of its IP addresses.

Some basic rules governing LISP are:

- * End-systems only send to addresses that are EIDs. EIDs are typically IP addresses assigned to hosts (other types of EID are supported by LISP, see [RFC8060] for further information). End-systems don't know that addresses are EIDs versus RLOCs but assume that packets get to their intended destinations. In a system where LISP is deployed, LISP routers intercept EID-addressed packets and assist in delivering them across the network core where EIDs cannot be routed. The procedure a host uses to send IP packets does not change.
- * LISP routers mostly deal with Routing Locator addresses. See details in Section 4.2 to clarify what is meant by "mostly".
- * RLOCs are always IP addresses assigned to routers, preferably topologically oriented addresses from provider CIDR (Classless Inter-Domain Routing) blocks.
- * When a router originates packets, it MAY use as a source address either an EID or RLOC. When acting as a host (e.g., when terminating a transport session such as Secure SHell (SSH),

TELNET, or the Simple Network Management Protocol (SNMP)), it MAY use an EID that is explicitly assigned for that purpose. An EID that identifies the router as a host MUST NOT be used as an RLOC; an EID is only routable within the scope of a site. A typical BGP configuration might demonstrate this "hybrid" EID/RLOC usage where a router could use its "host-like" EID to terminate iBGP sessions to other routers in a site while at the same time using RLOCs to terminate eBGP sessions to routers outside the site.

- * Packets with EIDs in them are not expected to be delivered end-to-end in the absence of an EID-to-RLOC mapping operation. They are expected to be used locally for intra-site communication or to be encapsulated for inter-site communication.
- * EIDs MAY also be structured (subnetted) in a manner suitable for local routing within an Autonomous System (AS).

An additional LISP header MAY be prepended to packets by a TE-ITR when re-routing of the path for a packet is desired. A potential use-case for this would be an ISP router that needs to perform Traffic Engineering for packets flowing through its network. In such a situation, termed "Recursive Tunneling", an ISP transit acts as an additional ITR, and the destination RLOC it uses for the new prepended header would be either a TE-ETR within the ISP (along an intra-ISP traffic engineered path) or a TE-ETR within another ISP (an inter-ISP traffic engineered path, where an agreement to build such a path exists).

In order to avoid excessive packet overhead as well as possible encapsulation loops, this document RECOMMENDS that a maximum of two LISP headers can be prepended to a packet. For initial LISP deployments, it is assumed that two headers is sufficient, where the first prepended header is used at a site for Location/Identity separation and the second prepended header is used inside a service provider for Traffic Engineering purposes.

Tunnel Routers can be placed fairly flexibly in a multi-AS topology. For example, the ITR for a particular end-to-end packet exchange might be the first-hop or default router within a site for the source host. Similarly, the ETR might be the last-hop router directly connected to the destination host. Another example, perhaps for a VPN service outsourced to an ISP by a site, the ITR could be the site's border router at the service provider attachment point. Mixing and matching of site-operated, ISP-operated, and other Tunnel Routers is allowed for maximum flexibility.

4.1. Deployment on the Public Internet

Several of the mechanisms in this document are intended for deployment in controlled, trusted environments, and are insecure for use over the public Internet. In particular, on the public internet xTRs:

- * MUST set the N, L, E, and V bits in the LISP header (Section 5.1) to zero.
- * MUST NOT use Locator-Status-Bits and echo-nonce, as described in Section 10 for Routing Locator Reachability. Instead MUST rely solely on control-plane methods.
- * MUST NOT use Gleaning or Locator-Status-Bits and Map-Versioning, as described in Section 13 to update the EID-to-RLOC Mappings. Instead relying solely on control-plane methods.

4.2. Packet Flow Sequence

This section provides an example of the unicast packet flow, including also Control-Plane information as specified in [I-D.ietf-lisp-rfc6833bis]. The example also assumes the following conditions:

- * Source host "host1.abc.example.com" is sending a packet to "host2.xyz.example.com", exactly as it would if the site was not using LISP.
- * Each site is multihomed, so each Tunnel Router has an address (RLOC) assigned from the service provider address block for each provider to which that particular Tunnel Router is attached.
- * The ITR(s) and ETR(s) are directly connected to the source and destination, respectively, but the source and destination can be located anywhere in the LISP site.
- * A Map-Request is sent for an external destination when the destination is not found in the forwarding table or matches a default route. Map-Requests are sent to the mapping database system by using the LISP Control-Plane protocol documented in [I-D.ietf-lisp-rfc6833bis].
- * Map-Replies are sent on the underlying routing system topology using the [I-D.ietf-lisp-rfc6833bis] Control-Plane protocol.

Client host1.abc.example.com wants to communicate with server host2.xyz.example.com:

1. host1.abc.example.com wants to open a TCP connection to host2.xyz.example.com. It does a DNS lookup on host2.xyz.example.com. An A/AAAA record is returned. This address is the destination EID. The locally assigned address of host1.abc.example.com is used as the source EID. An IPv4 or IPv6 packet is built and forwarded through the LISP site as a normal IP packet until it reaches a LISP ITR.
2. The LISP ITR must be able to map the destination EID to an RLOC of one of the ETRs at the destination site. A method to do this is to send a LISP Map-Request, as specified in [I-D.ietf-lisp-rfc6833bis].
3. The mapping system helps forwarding the Map-Request to the corresponding ETR. When the Map-Request arrives at one of the ETRs at the destination site, it will process the packet as a control message.
4. The ETR looks at the destination EID of the Map-Request and matches it against the prefixes in the ETR's configured EID-to-RLOC mapping database. This is the list of EID-Prefixes the ETR is supporting for the site it resides in. If there is no match, the Map-Request is dropped. Otherwise, a LISP Map-Reply is returned to the ITR.
5. The ITR receives the Map-Reply message, parses the message, and stores the mapping information from the packet. This information is stored in the ITR's EID-to-RLOC Map-Cache. Note that the Map-Cache is an on-demand cache. An ITR will manage its Map-Cache in such a way that optimizes for its resource constraints.
6. Subsequent packets from host1.abc.example.com to host2.xyz.example.com will have a LISP header prepended by the ITR using the appropriate RLOC as the LISP header destination address learned from the ETR. Note that the packet MAY be sent to a different ETR than the one that returned the Map-Reply due to the source site's hashing policy or the destination site's Locator-Set policy.
7. The ETR receives these packets directly (since the destination address is one of its assigned IP addresses), checks the validity of the addresses, strips the LISP header, and forwards packets to the attached destination host.
8. In order to defer the need for a mapping lookup in the reverse direction, an ETR can OPTIONALLY create a cache entry that maps the source EID (inner-header source IP address) to the source RLOC (outer-header source IP address) in a received LISP packet.

Such a cache entry is termed a "glean mapping" and only contains a single RLOC for the EID in question. More complete information about additional RLOCs SHOULD be verified by sending a LISP Map-Request for that EID. Both the ITR and the ETR MAY also influence the decision the other makes in selecting an RLOC.

5. LISP Encapsulation Details

Since additional tunnel headers are prepended, the packet becomes larger and can exceed the MTU of any link traversed from the ITR to the ETR. It is RECOMMENDED in IPv4 that packets do not get fragmented as they are encapsulated by the ITR. Instead, the packet is dropped and an ICMP Unreachable/Fragmentation-Needed message is returned to the source.

In the case when fragmentation is needed, this specification RECOMMENDS that implementations provide support for one of the proposed fragmentation and reassembly schemes. Two existing schemes are detailed in Section 7.

Since IPv4 or IPv6 addresses can be either EIDs or RLOCs, the LISP architecture supports IPv4 EIDs with IPv6 RLOCs (where the inner header is in IPv4 packet format and the outer header is in IPv6 packet format) or IPv6 EIDs with IPv4 RLOCs (where the inner header is in IPv6 packet format and the outer header is in IPv4 packet format). The next sub-sections illustrate packet formats for the homogeneous case (IPv4-in-IPv4 and IPv6-in-IPv6), but all 4 combinations MUST be supported. Additional types of EIDs are defined in [RFC8060].

As LISP uses UDP encapsulation to carry traffic between xTRs across the Internet, implementors should be aware of the provisions of [RFC8085], especially those given in section 3.1.11 on congestion control for UDP tunneling.

Implementors are encouraged to consider UDP checksum usage guidelines in section 3.4 of [RFC8085] when it is desirable to protect UDP and LISP headers against corruption.

5.1. LISP IPv4-in-IPv4 Header Format

ITR. The address fields contain RLOCs obtained from the ingress router's EID-to-RLOC Cache. The IP protocol number is "UDP (17)" from [RFC0768]. The setting of the Don't Fragment (DF) bit 'Flags' field is according to rules listed in Sections 7.1 and 7.2.

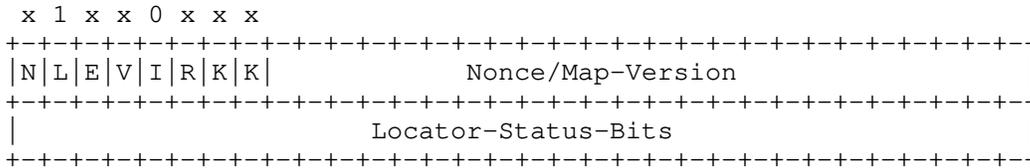
UDP Header: The UDP header contains an ITR selected source port when encapsulating a packet. See Section 12 for details on the hash algorithm used to select a source port based on the 5-tuple of the inner header. The destination port MUST be set to the well-known IANA-assigned port value 4341.

UDP Checksum: The 'UDP Checksum' field SHOULD be transmitted as zero by an ITR for either IPv4 [RFC0768] and IPv6 encapsulation [RFC6935] [RFC6936]. When a packet with a zero UDP checksum is received by an ETR, the ETR MUST accept the packet for decapsulation. When an ITR transmits a non-zero value for the UDP checksum, it MUST send a correctly computed value in this field. When an ETR receives a packet with a non-zero UDP checksum, it MAY choose to verify the checksum value. If it chooses to perform such verification, and the verification fails, the packet MUST be silently dropped. If the ETR chooses not to perform the verification, or performs the verification successfully, the packet MUST be accepted for decapsulation. The handling of UDP zero checksums over IPv6 for all tunneling protocols, including LISP, is subject to the applicability statement in [RFC6936].

UDP Length: The 'UDP Length' field is set for an IPv4-encapsulated packet to be the sum of the inner-header IPv4 Total Length plus the UDP and LISP header lengths. For an IPv6-encapsulated packet, the 'UDP Length' field is the sum of the inner-header IPv6 Payload Length, the size of the IPv6 header (40 octets), and the size of the UDP and LISP headers.

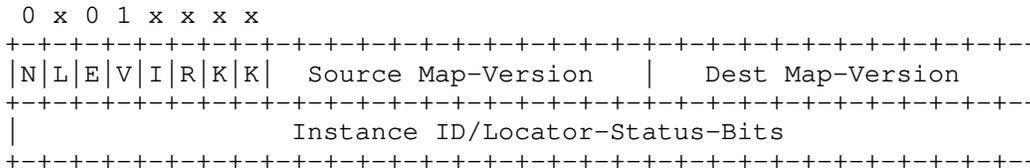
N: The N-bit is the nonce-present bit. When this bit is set to 1, the low-order 24 bits of the first 32 bits of the LISP header contain a Nonce. See Section 10.1 for details. Both N- and V-bits MUST NOT be set in the same packet. If they are, a decapsulating ETR MUST treat the 'Nonce/Map-Version' field as having a Nonce value present.

L: The L-bit is the 'Locator-Status-Bits' field enabled bit. When this bit is set to 1, the Locator-Status-Bits in the second 32 bits of the LISP header are in use.

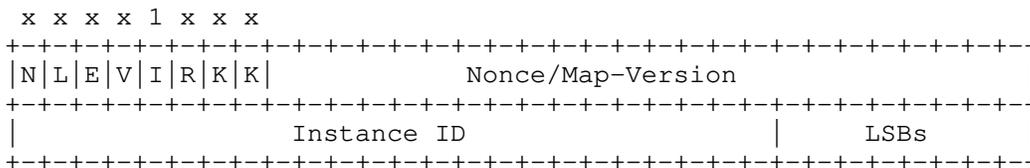


E: The E-bit is the echo-nonce-request bit. This bit MUST be ignored and has no meaning when the N-bit is set to 0. When the N-bit is set to 1 and this bit is set to 1, an ITR is requesting that the nonce value in the 'Nonce' field be echoed back in LISP-encapsulated packets when the ITR is also an ETR. See Section 10.1 for details.

V: The V-bit is the Map-Version present bit. When this bit is set to 1, the N-bit MUST be 0. Refer to the [I-D.ietf-lisp-6834bis] specification for more details on Database Map-Versioning. This bit indicates that the LISP header is encoded in this case as:



I: The I-bit is the Instance ID bit. See Section 8 for more details. When this bit is set to 1, the 'Locator-Status-Bits' field is reduced to 8 bits and the high-order 24 bits are used as an Instance ID. If the L-bit is set to 0, then the low-order 8 bits are transmitted as zero and ignored on receipt. The format of the LISP header would look like this:



R: The R-bit is a Reserved and unassigned bit for future use. It MUST be set to 0 on transmit and MUST be ignored on receipt.

KK: The KK-bits are a 2-bit field used when encapsulated packets are encrypted. The field is set to 00 when the packet is not encrypted. See [RFC8061] for further information.

LISP Nonce: The LISP 'Nonce' field is a 24-bit value that is

randomly generated by an ITR when the N-bit is set to 1. Nonce generation algorithms are an implementation matter but are required to generate different nonces when sending to different RLOCs. The nonce is also used when the E-bit is set to request the nonce value to be echoed by the other side when packets are returned. When the E-bit is clear but the N-bit is set, a remote ITR is either echoing a previously requested echo-nonce or providing a random nonce. See Section 10.1 for more details. Finally, when both the N and V-bit are not set (N=0, V=0), then both the Nonce and Map-Version fields are set to 0 and ignored on receipt.

LISP Locator-Status-Bits (LSBs): When the L-bit is also set, the 'Locator-Status-Bits' field in the LISP header is set by an ITR to indicate to an ETR the up/down status of the Locators in the source site. Each RLOC in a Map-Reply is assigned an ordinal value from 0 to n-1 (when there are n RLOCs in a mapping entry). The Locator-Status-Bits are numbered from 0 to n-1 from the least significant bit of the field. The field is 32 bits when the I-bit is set to 0 and is 8 bits when the I-bit is set to 1. When a Locator-Status-Bit is set to 1, the ITR is indicating to the ETR that the RLOC associated with the bit ordinal has up status. See Section 10 for details on how an ITR can determine the status of the ETRs at the same site. When a site has multiple EID-Prefixes that result in multiple mappings (where each could have a different Locator-Set), the Locator-Status-Bits setting in an encapsulated packet MUST reflect the mapping for the EID-Prefix that the inner-header source EID address matches (longest-match). If the LSB for an anycast Locator is set to 1, then there is at least one RLOC with that address, and the ETR is considered 'up'.

When doing ITR/PITR encapsulation:

- * The outer-header 'Time to Live' field (or 'Hop Limit' field, in the case of IPv6) SHOULD be copied from the inner-header 'Time to Live' field.
- * The outer-header IPv4 'Differentiated Services Code Point' (DSCP) field or the 'Traffic Class' field, in the case of IPv6, SHOULD be copied from the inner-header IPv4 DSCP field or 'Traffic Class' field in the case of IPv6, to the outer-header. Guidelines for this can be found at [RFC2983].
- * The IPv4 'Explicit Congestion Notification' (ECN) field and bits 6 and 7 of the IPv6 'Traffic Class' field requires special treatment in order to avoid discarding indications of congestion as specified in [RFC6040].

When doing ETR/PETR decapsulation:

- * The inner-header IPv4 'Time to Live' field or 'Hop Limit' field in the case of IPv6, MUST be copied from the outer-header 'Time to Live'/'Hop Limit' field, when the 'Time to Live'/'Hop Limit' value of the outer header is less than the 'Time to Live'/'Hop Limit' value of the inner header. Failing to perform this check can cause the 'Time to Live'/'Hop Limit' of the inner header to increment across encapsulation/decapsulation cycles. This check is also performed when doing initial encapsulation, when a packet comes to an ITR or PITR destined for a LISP site.
- * The outer-header IPv4 'Differentiated Services Code Point' (DSCP) field or the 'Traffic Class' field in the case of IPv6, SHOULD be copied from the outer-header IPv4 DSCP field or 'Traffic Class' field in the case of IPv6, to the inner-header. Guidelines for this can be found at [RFC2983].
- * The IPv4 'Explicit Congestion Notification' (ECN) field and bits 6 and 7 of the IPv6 'Traffic Class' field, requires special treatment in order to avoid discarding indications of congestion as specified in [RFC6040]. Note that implementations exist that copy the 'ECN' field from the outer header to the inner header even though [RFC6040] does not recommend this behavior. It is RECOMMENDED that implementations change to support the behavior in [RFC6040].

Note that if an ETR/PETR is also an ITR/PITR and chooses to re-encapsulate after decapsulating, the net effect of this is that the new outer header will carry the same Time to Live as the old outer header minus 1.

Copying the Time to Live (TTL) serves two purposes: first, it preserves the distance the host intended the packet to travel; second, and more importantly, it provides for suppression of looping packets in the event there is a loop of concatenated tunnels due to misconfiguration.

Some xTRs and PxTRs performs re-encapsulation operations and need to treat the 'Explicit Congestion Notification' (ECN) in a special way. Because the re-encapsulation operation is a sequence of two operations, namely a decapsulation followed by an encapsulation, the ECN bits MUST be treated as described above for these two operations.

The LISP dataplane protocol is not backwards compatible with [RFC6830] and does not have explicit support for introducing future protocol changes (e.g. an explicit version field). However, the LISP control plane [I-D.ietf-lisp-rfc6833bis] allows an ETR to register

dataplane capabilities by means of new LCAF types [RFC8060]. In this way an ITR can be made aware of the dataplane capabilities of an ETR, and encapsulate accordingly. The specification of the new LCAF types, new LCAF mechanisms, and their use, is out of the scope of this document.

6. LISP EID-to-RLOC Map-Cache

ITRs and PITRs maintain an on-demand cache, referred as LISP EID-to-RLOC Map-Cache, that contains mappings from EID-prefixes to locator sets. The cache is used to encapsulate packets from the EID space to the corresponding RLOC network attachment point.

When an ITR/PITR receives a packet from inside of the LISP site to destinations outside of the site a longest-prefix match lookup of the EID is done to the Map-Cache.

When the lookup succeeds, the Locator-Set retrieved from the Map-Cache is used to send the packet to the EID's topological location.

If the lookup fails, the ITR/PITR needs to retrieve the mapping using the LISP Control-Plane protocol [I-D.ietf-lisp-rfc6833bis]. While the mapping is being retrieved, the ITR/PITR can either drop or buffer the packets. This document does not have specific recommendations about the action to be taken. It is up to the deployer to consider whether or not it is desirable to buffer packets and deploy a LISP implementation that offers the desired behaviour. Once the mapping is resolved it is then stored in the local Map-Cache to forward subsequent packets addressed to the same EID-prefix.

The Map-Cache is a local cache of mappings, entries are expired based on the associated Time to live. In addition, entries can be updated with more current information, see Section 13 for further information on this. Finally, the Map-Cache also contains reachability information about EIDs and RLOCs, and uses LISP reachability information mechanisms to determine the reachability of RLOCs, see Section 10 for the specific mechanisms.

7. Dealing with Large Encapsulated Packets

This section proposes two mechanisms to deal with packets that exceed the path MTU between the ITR and ETR.

It is left to the implementor to decide if the stateless or stateful mechanism SHOULD be implemented. Both or neither can be used, since it is a local decision in the ITR regarding how to deal with MTU issues, and sites can interoperate with differing mechanisms.

Both stateless and stateful mechanisms also apply to Re-encapsulating and Recursive Tunneling, so any actions below referring to an ITR also apply to a TE-ITR.

7.1. A Stateless Solution to MTU Handling

An ITR stateless solution to handle MTU issues is described as follows:

1. Define H to be the size, in octets, of the outer header an ITR prepends to a packet. This includes the UDP and LISP header lengths.
2. Define L to be the size, in octets, of the maximum-sized packet an ITR can send to an ETR without the need for the ITR or any intermediate routers to fragment the packet. The network administrator of the LISP deployment has to determine what is the suitable value of L so to make sure that no MTU issues arise.
3. Define an architectural constant S for the maximum size of a packet, in octets, an ITR MUST receive from the source so the effective MTU can be met. That is, $L = S + H$.

When an ITR receives a packet from a site-facing interface and adds H octets worth of encapsulation to yield a packet size greater than L octets (meaning the received packet size was greater than S octets from the source), it resolves the MTU issue by first splitting the original packet into 2 equal-sized fragments. A LISP header is then prepended to each fragment. The size of the encapsulated fragments is then $(S/2 + H)$, which is less than the ITR's estimate of the path MTU between the ITR and its correspondent ETR.

When an ETR receives encapsulated fragments, it treats them as two individually encapsulated packets. It strips the LISP headers and then forwards each fragment to the destination host of the destination site. The two fragments are reassembled at the destination host into the single IP datagram that was originated by the source host. Note that reassembly can happen at the ETR if the encapsulated packet was fragmented at or after the ITR.

This behavior MUST be performed by the ITR only when the source host originates a packet with the 'DF' field of the IP header set to 0. When the 'DF' field of the IP header is set to 1, or the packet is an IPv6 packet originated by the source host, the ITR will drop the packet when the size (adding in the size of the encapsulation header) is greater than L and send an ICMPv4 ICMP Unreachable/Fragmentation-Needed or ICMPv6 "Packet Too Big" message to the source with a value of S, where S is $(L - H)$.

When the outer-header encapsulation uses an IPv4 header, an implementation SHOULD set the DF bit to 1 so ETR fragment reassembly can be avoided. An implementation MAY set the DF bit in such headers to 0 if it has good reason to believe there are unresolvable path MTU issues between the sending ITR and the receiving ETR.

This specification RECOMMENDS that L be defined as 1500. Additional information about in-network MTU and fragmentation issues can be found at [RFC4459].

7.2. A Stateful Solution to MTU Handling

An ITR stateful solution to handle MTU issues is described as follows:

1. The ITR will keep state of the effective MTU for each Locator per Map-Cache entry. The effective MTU is what the core network can deliver along the path between the ITR and ETR.
2. When an IPv4-encapsulated packet with the DF bit set to 1, exceeds what the core network can deliver, one of the intermediate routers on the path will send an ICMPv4 Unreachable/Fragmentation-Needed to the ITR, respectively. The ITR will parse the ICMP message to determine which Locator is affected by the effective MTU change and then record the new effective MTU value in the Map-Cache entry.
3. When a packet is received by the ITR from a source inside of the site and the size of the packet is greater than the effective MTU stored with the Map-Cache entry associated with the destination EID the packet is for, the ITR will send an ICMPv4 ICMP Unreachable/Fragmentation-Needed message back to the source. The packet size advertised by the ITR in the ICMP message is the effective MTU minus the LISP encapsulation length.

Even though this mechanism is stateful, it has advantages over the stateless IP fragmentation mechanism, by not involving the destination host with reassembly of ITR fragmented packets.

Please note that [RFC1191] and [RFC1981], which describe the use of ICMP packets for PMTU discovery, can behave suboptimally in the presence of ICMP black holes or off-path attackers that spoof ICMP. Possible mitigations include ITRs and ETRs cooperating on MTU probe packets ([RFC4821], [I-D.ietf-tsvwg-datagram-plpmtud]), or ITRs storing the beginning of large packets to verify that they match the echoed packet in ICMP Frag Needed/PTB.

8. Using Virtualization and Segmentation with LISP

There are several cases where segregation is needed at the EID level. For instance, this is the case for deployments containing overlapping addresses, traffic isolation policies or multi-tenant virtualization. For these and other scenarios where segregation is needed, Instance IDs are used.

An Instance ID can be carried in a LISP-encapsulated packet. An ITR that prepends a LISP header will copy a 24-bit value used by the LISP router to uniquely identify the address space. The value is copied to the 'Instance ID' field of the LISP header, and the I-bit is set to 1.

When an ETR decapsulates a packet, the Instance ID from the LISP header is used as a table identifier to locate the forwarding table to use for the inner destination EID lookup.

For example, an 802.1Q VLAN tag or VPN identifier could be used as a 24-bit Instance ID. See [I-D.ietf-lisp-vpn] for LISP VPN use-case details. Please note that the Instance ID is not protected, an on-path attacker can modify the tags and for instance, allow communications between logically isolated VLANs.

Participants within a LISP deployment must agree on the meaning of Instance ID values. The source and destination EIDs MUST belong to the same Instance ID.

Instance ID SHOULD NOT be used with overlapping IPv6 EID addresses.

9. Routing Locator Selection

The Map-Cache contains the state used by ITRs and PITRs to encapsulate packets. When an ITR/PITR receives a packet from inside the LISP site to a destination outside of the site a longest-prefix match lookup of the EID is done to the Map-Cache (see Section 6). The lookup returns a single Locator-Set containing a list of RLOCs corresponding to the EID's topological location. Each RLOC in the Locator-Set is associated with a 'Priority' and 'Weight', this information is used to select the RLOC to encapsulate.

The RLOC with the lowest 'Priority' is selected. An RLOC with 'Priority' 255 means that MUST NOT be used for forwarding. When multiple RLOCs have the same 'Priority' then the 'Weight' states how to load balance traffic among them. The value of the 'Weight' represents the relative weight of the total packets that match the mapping entry.

The following are different scenarios for choosing RLOCs and the controls that are available:

- * The server-side returns one RLOC. The client-side can only use one RLOC. The server-side has complete control of the selection.
- * The server-side returns a list of RLOCs where a subset of the list has the same best Priority. The client can only use the subset list according to the weighting assigned by the server-side. In this case, the server-side controls both the subset list and load-splitting across its members. The client-side can use RLOCs outside of the subset list if it determines that the subset list is unreachable (unless RLOCs are set to a Priority of 255). Some sharing of control exists: the server-side determines the destination RLOC list and load distribution while the client-side has the option of using alternatives to this list if RLOCs in the list are unreachable.
- * The server-side sets a Weight of zero for the RLOC subset list. In this case, the client-side can choose how the traffic load is spread across the subset list. See Section 12 for details on load-sharing mechanisms. Control is shared by the server-side determining the list and the client-side determining load distribution. Again, the client can use alternative RLOCs if the server-provided list of RLOCs is unreachable.
- * Either side (more likely the server-side ETR) decides to "glean" the RLOCs. For example, if the server-side ETR gleans RLOCs, then the client-side ITR gives the client-side ITR responsibility for bidirectional RLOC reachability and preferability. Server-side ETR gleaning of the client-side ITR RLOC is done by caching the inner-header source EID and the outer-header source RLOC of received packets. The client-side ITR controls how traffic is returned and can alternate using an outer-header source RLOC, which then can be added to the list the server-side ETR uses to return traffic. Since no Priority or Weights are provided using this method, the server-side ETR MUST assume that each client-side ITR RLOC uses the same best Priority with a Weight of zero. In addition, since EID-Prefix encoding cannot be conveyed in data packets, the EID-to-RLOC Cache on Tunnel Routers can grow to be very large. Gleaning has several important considerations. A "gleaned" Map-Cache entry is only stored and used for a RECOMMENDED period of 3 seconds, pending verification. Verification MUST be performed by sending a Map-Request to the source EID (the inner-header IP source address) of the received encapsulated packet. A reply to this "verifying Map-Request" is used to fully populate the Map-Cache entry for the "gleaned" EID and is stored and used for the time indicated from the 'TTL' field

of a received Map-Reply. When a verified Map- Cache entry is stored, data gleaning no longer occurs for subsequent packets that have a source EID that matches the EID-Prefix of the verified entry. This "gleaning" mechanism MUST NOT be used over the public Internet and SHOULD only be used in trusted and closed deployments. Refer to Section 16 for security issues regarding this mechanism.

RLOCs that appear in EID-to-RLOC Map-Reply messages are assumed to be reachable when the R-bit [I-D.ietf-lisp-rfc6833bis] for the Locator record is set to 1. When the R-bit is set to 0, an ITR or PITR MUST NOT encapsulate to the RLOC. Neither the information contained in a Map-Reply nor that stored in the mapping database system provides reachability information for RLOCs. Note that reachability is not part of the mapping system and is determined using one or more of the Routing Locator reachability algorithms described in the next section.

10. Routing Locator Reachability

Several Data-Plane mechanisms for determining RLOC reachability are currently defined. Please note that additional Control-Plane based reachability mechanisms are defined in [I-D.ietf-lisp-rfc6833bis].

1. An ETR MAY examine the Locator-Status-Bits in the LISP header of an encapsulated data packet received from an ITR. If the ETR is also acting as an ITR and has traffic to return to the original ITR site, it can use this status information to help select an RLOC.
2. When an ETR receives an encapsulated packet from an ITR, the source RLOC from the outer header of the packet is likely to be reachable. Please note that in some scenarios the RLOC from the outer header can be a spoofable field.
3. An ITR/ETR pair can use the 'Echo-Noncing' Locator reachability algorithms described in this section.

When determining Locator up/down reachability by examining the Locator-Status-Bits from the LISP-encapsulated data packet, an ETR will receive up-to-date status from an encapsulating ITR about reachability for all ETRs at the site. CE-based ITRs at the source site can determine reachability relative to each other using the site IGP as follows:

- * Under normal circumstances, each ITR will advertise a default route into the site IGP.

- * If an ITR fails or if the upstream link to its PE fails, its default route will either time out or be withdrawn.

Each ITR can thus observe the presence or lack of a default route originated by the others to determine the Locator-Status-Bits it sets for them.

When ITRs at the site are not deployed in CE routers, the IGP can still be used to determine the reachability of Locators, provided they are injected into the IGP. This is typically done when a /32 address is configured on a loopback interface.

RLOCs listed in a Map-Reply are numbered with ordinals 0 to n-1. The Locator-Status-Bits in a LISP-encapsulated packet are numbered from 0 to n-1 starting with the least significant bit. For example, if an RLOC listed in the 3rd position of the Map-Reply goes down (ordinal value 2), then all ITRs at the site will clear the 3rd least significant bit (xxxx x0xx) of the 'Locator-Status-Bits' field for the packets they encapsulate.

When an xTR decides to use 'Locator-Status-Bits' to affect reachability information, it acts as follows: ETRs decapsulating a packet will check for any change in the 'Locator-Status-Bits' field. When a bit goes from 1 to 0, the ETR, if acting also as an ITR, will refrain from encapsulating packets to an RLOC that is indicated as down. It will only resume using that RLOC if the corresponding Locator-Status-Bit returns to a value of 1. Locator-Status-Bits are associated with a Locator-Set per EID-Prefix. Therefore, when a Locator becomes unreachable, the Locator-Status-Bit that corresponds to that Locator's position in the list returned by the last Map-Reply will be set to zero for that particular EID-Prefix.

Locator-Status-Bits MUST NOT be used over the public Internet and SHOULD only be used in trusted and closed deployments. In addition Locator-Status-Bits SHOULD be coupled with Map-Versioning [I-D.ietf-lisp-6834bis] to prevent race conditions where Locator-Status-Bits are interpreted as referring to different RLOCs than intended. Refer to Section 16 for security issues regarding this mechanism.

If an ITR encapsulates a packet to an ETR and the packet is received and decapsulated by the ETR, it is implied but not confirmed by the ITR that the ETR's RLOC is reachable. In most cases, the ETR can also reach the ITR but cannot assume this to be true, due to the possibility of path asymmetry. In the presence of unidirectional traffic flow from an ITR to an ETR, the ITR SHOULD NOT use the lack of return traffic as an indication that the ETR is unreachable. Instead, it MUST use an alternate mechanism to determine reachability.

The security considerations of Section 16 related to data-plane reachability applies to the data-plane RLOC reachability mechanisms described in this section.

10.1. Echo Nonce Algorithm

When data flows bidirectionally between Locators from different sites, a Data-Plane mechanism called "nonce echoing" can be used to determine reachability between an ITR and ETR. When an ITR wants to solicit a nonce echo, it sets the N- and E-bits and places a 24-bit nonce [RFC4086] in the LISP header of the next encapsulated data packet.

When this packet is received by the ETR, the encapsulated packet is forwarded as normal. When the ETR is an xTR (co-located as an ITR), it then sends a data packet to the ITR (when it is an xTR co-located as an ETR), it includes the nonce received earlier with the N-bit set and E-bit cleared. The ITR sees this "echoed nonce" and knows that the path to and from the ETR is up.

The ITR will set the E-bit and N-bit for every packet it sends while in the echo-nonce-request state. The time the ITR waits to process the echoed nonce before it determines the path is unreachable is variable and is a choice left for the implementation.

If the ITR is receiving packets from the ETR but does not see the nonce echoed while being in the echo-nonce-request state, then the path to the ETR is unreachable. This decision MAY be overridden by other Locator reachability algorithms. Once the ITR determines that the path to the ETR is down, it can switch to another Locator for that EID-Prefix.

Note that "ITR" and "ETR" are relative terms here. Both devices MUST be implementing both ITR and ETR functionality for the echo nonce mechanism to operate.

The ITR and ETR MAY both go into the echo-nonce-request state at the same time. The number of packets sent or the time during which echo nonce requests are sent is an implementation-specific setting. In this case, an xTR receiving the echo-nonce-request packets will suspend the echo-nonce-request state and setup a 'echo-nonce-request-state' timer. After the 'echo-nonce-request-state' timer expires it will resume the echo-nonce-request state.

This mechanism does not completely solve the forward path reachability problem, as traffic may be unidirectional. That is, the ETR receiving traffic at a site MAY not be the same device as an ITR that transmits traffic from that site, or the site-to-site traffic is unidirectional so there is no ITR returning traffic.

The echo-nonce algorithm is bilateral. That is, if one side sets the E-bit and the other side is not enabled for echo-nonce, then the echoing of the nonce does not occur and the requesting side may erroneously consider the Locator unreachable. An ITR SHOULD set the E-bit in an encapsulated data packet when it knows the ETR is enabled for echo-nonce. This is conveyed by the E-bit in the Map-Reply message.

Many implementations default to not advertising they are echo-nonce capable in Map-Reply messages and so RLOC-probing tends to be used for RLOC reachability.

The echo-nonce mechanism MUST NOT be used over the public Internet and MUST only be used in trusted and closed deployments. Refer to Section 16 for security issues regarding this mechanism.

11. EID Reachability within a LISP Site

A site MAY be multihomed using two or more ETRs. The hosts and infrastructure within a site will be addressed using one or more EID-Prefixes that are mapped to the RLOCs of the relevant ETRs in the mapping system. One possible failure mode is for an ETR to lose reachability to one or more of the EID-Prefixes within its own site. When this occurs when the ETR sends Map-Replies, it can clear the R-bit associated with its own Locator. And when the ETR is also an ITR, it can clear its Locator-Status-Bit in the encapsulation data header.

It is recognized that there are no simple solutions to the site partitioning problem because it is hard to know which part of the EID-Prefix range is partitioned and which Locators can reach any sub-ranges of the EID-Prefixes. Note that this is not a new problem introduced by the LISP architecture. The problem exists today when a multihomed site uses BGP to advertise its reachability upstream.

12. Routing Locator Hashing

When an ETR provides an EID-to-RLOC mapping in a Map-Reply message that is stored in the Map-Cache of a requesting ITR, the Locator-Set for the EID-Prefix MAY contain different Priority and Weight values for each locator address. When more than one best Priority Locator exists, the ITR can decide how to load-share traffic against the corresponding Locators.

The following hash algorithm MAY be used by an ITR to select a Locator for a packet destined to an EID for the EID-to-RLOC mapping:

1. Either a source and destination address hash or the traditional 5-tuple hash can be used. The traditional 5-tuple hash includes the source and destination addresses; source and destination TCP, UDP, or Stream Control Transmission Protocol (SCTP) port numbers; and the IP protocol number field or IPv6 next-protocol fields of a packet that a host originates from within a LISP site. When a packet is not a TCP, UDP, or SCTP packet, the source and destination addresses only from the header are used to compute the hash.
2. Take the hash value and divide it by the number of Locators stored in the Locator-Set for the EID-to-RLOC mapping.
3. The remainder will yield a value of 0 to "number of Locators minus 1". Use the remainder to select the Locator in the Locator-Set.

The specific hash algorithm the ITR uses for load-sharing is out of scope for this document and does not prevent interoperability.

The Source port SHOULD be the same for all packets belonging to the same flow. Also note that when a packet is LISP encapsulated, the source port number in the outer UDP header needs to be set. Selecting a hashed value allows core routers that are attached to Link Aggregation Groups (LAGs) to load-split the encapsulated packets across member links of such LAGs. Otherwise, core routers would see a single flow, since packets have a source address of the ITR, for packets that are originated by different EIDs at the source site. A suggested setting for the source port number computed by an ITR is a 5-tuple hash function on the inner header, as described above. The source port SHOULD be the same for all packets belonging to the same flow.

Many core router implementations use a 5-tuple hash to decide how to balance packet load across members of a LAG. The 5-tuple hash includes the source and destination addresses of the packet and the

source and destination ports when the protocol number in the packet is TCP or UDP. For this reason, UDP encoding is used for LISP encapsulation. In this scenario, when the outer header is IPv6, the flow label MAY also be set following the procedures specified in [RFC6438]. When the inner header is IPv6 then the flow label is not zero, it MAY be used to compute the hash.

13. Changing the Contents of EID-to-RLOC Mappings

Since the LISP architecture uses a caching scheme to retrieve and store EID-to-RLOC mappings, the only way an ITR can get a more up-to-date mapping is to re-request the mapping. However, the ITRs do not know when the mappings change, and the ETRs do not keep track of which ITRs requested its mappings. For scalability reasons, it is desirable to maintain this approach but need to provide a way for ETRs to change their mappings and inform the sites that are currently communicating with the ETR site using such mappings.

This section defines two Data-Plane mechanism for updating EID-to-RLOC mappings. Additionally, the Solicit-Map Request (SMR) Control-Plane updating mechanism is specified in [I-D.ietf-lisp-rfc6833bis].

13.1. Locator-Status-Bits

Locator-Status-Bits (LSB) can also be used to keep track of the Locator status (up or down) when EID-to-RLOC mappings are changing. When LSB are used in a LISP deployment, all LISP tunnel routers MUST implement both ITR and ETR capabilities (therefore all tunnel routers are effectively xTRs). In this section the term "source xTR" is used to refer to the xTR setting the LSB and "destination xTR" is used to refer to the xTR receiving the LSB. The procedure is as follows:

First, when a Locator record is added or removed from the Locator-Set, the source xTR will signal this by sending a Solicit-Map Request (SMR) Control-Plane message [I-D.ietf-lisp-rfc6833bis] to the destination xTR. At this point the source xTR MUST NOT use LSB (L-bit = 0) since the destination xTR site has outdated information. The source xTR will setup a 'use-LSB' timer.

Second and as defined in [I-D.ietf-lisp-rfc6833bis], upon reception of the SMR message the destination xTR will retrieve the updated EID-to-RLOC mappings by sending a Map-Request.

And third, when the 'use-LSB' timer expires, the source xTR can use again LSB with the destination xTR to signal the Locator status (up or down). The specific value for the 'use-LSB' timer depends on the LISP deployment, the 'use-LSB' timer needs to be large enough for the destination xTR to retrieve the updated EID-to-RLOC mappings. A RECOMMENDED value for the 'use-LSB' timer is 5 minutes.

13.2. Database Map-Versioning

When there is unidirectional packet flow between an ITR and ETR, and the EID-to-RLOC mappings change on the ETR, it needs to inform the ITR so encapsulation to a removed Locator can stop and can instead be started to a new Locator in the Locator-Set.

An ETR, can send Map-Reply messages carrying a Map-Version Number in an EID-record. This is known as the Destination Map-Version Number. ITRs include the Destination Map-Version Number in packets they encapsulate to the site.

An ITR, when it encapsulates packets to ETRs, can convey its own Map-Version Number. This is known as the Source Map-Version Number.

When presented in EID-records of Map-Register messages, a Map-Version Number is a good way for the Map-Server to assure that all ETRs for a site registering to it are synchronized according to Map-Version Number.

See [I-D.ietf-lisp-6834bis] for a more detailed analysis and description of Database Map-Versioning.

14. Multicast Considerations

A multicast group address, as defined in the original Internet architecture, is an identifier of a grouping of topologically independent receiver host locations. The address encoding itself does not determine the location of the receiver(s). The multicast routing protocol, and the network-based state the protocol creates, determine where the receivers are located.

In the context of LISP, a multicast group address is both an EID and a Routing Locator. Therefore, no specific semantic or action needs to be taken for a destination address, as it would appear in an IP header. Therefore, a group address that appears in an inner IP header built by a source host will be used as the destination EID. The outer IP header (the destination Routing Locator address), prepended by a LISP router, can use the same group address as the destination Routing Locator, use a multicast or unicast Routing Locator obtained from a Mapping System lookup, or use other means to determine the group address mapping.

With respect to the source Routing Locator address, the ITR prepends its own IP address as the source address of the outer IP header, just like it would if the destination EID was a unicast address. This source Routing Locator address, like any other Routing Locator address, MUST be routable on the underlay.

There are two approaches for LISP-Multicast, one that uses native multicast routing in the underlay with no support from the Mapping System and the other that uses only unicast routing in the underlay with support from the Mapping System. See [RFC6831] and [RFC8378], respectively, for details. Details for LISP-Multicast and interworking with non-LISP sites are described in [RFC6831] and [RFC6832].

15. Router Performance Considerations

LISP is designed to be very "hardware-based forwarding friendly". A few implementation techniques can be used to incrementally implement LISP:

- * When a tunnel-encapsulated packet is received by an ETR, the outer destination address may not be the address of the router. This makes it challenging for the control plane to get packets from the hardware. This may be mitigated by creating special Forwarding Information Base (FIB) entries for the EID-Prefixes of EIDs served by the ETR (those for which the router provides an RLOC translation). These FIB entries are marked with a flag indicating that Control-Plane processing SHOULD be performed. The forwarding logic of testing for particular IP protocol number values is not necessary. There are a few proven cases where no changes to existing deployed hardware were needed to support the LISP Data-Plane.

- * On an ITR, prepending a new IP header consists of adding more octets to a MAC rewrite string and prepending the string as part of the outgoing encapsulation procedure. Routers that support Generic Routing Encapsulation (GRE) tunneling [RFC2784] or 6to4 tunneling [RFC3056] may already support this action.
- * A packet's source address or interface the packet was received on can be used to select VRF (Virtual Routing/Forwarding). The VRF's routing table can be used to find EID-to-RLOC mappings.

For performance issues related to Map-Cache management, see Section 16.

16. Security Considerations

In what follows we highlight security considerations that apply when LISP is deployed in environments such as those specified in Section 1.1.

The optional mechanisms of gleaning is offered to directly obtain a mapping from the LISP encapsulated packets. Specifically, an xTR can learn the EID-to-RLOC mapping by inspecting the source RLOC and source EID of an encapsulated packet, and insert this new mapping into its Map-Cache. An off-path attacker can spoof the source EID address to divert the traffic sent to the victim's spoofed EID. If the attacker spoofs the source RLOC, it can mount a DoS attack by redirecting traffic to the spoofed victim's RLOC, potentially overloading it.

The LISP Data-Plane defines several mechanisms to monitor RLOC Data-Plane reachability, in this context Locator-Status Bits, Nonce-Present and Echo-Nonce bits of the LISP encapsulation header can be manipulated by an attacker to mount a DoS attack. An off-path attacker able to spoof the RLOC and/or nonce of a victim's xTR can manipulate such mechanisms to declare false information about the RLOC's reachability status.

For example of such attacks, an off-path attacker can exploit the echo-nonce mechanism by sending data packets to an ITR with a random nonce from an ETR's spoofed RLOC. Note the attacker must guess a valid nonce the ITR is requesting to be echoed within a small window of time. The goal is to convince the ITR that the ETR's RLOC is reachable even when it may not be reachable. If the attack is successful, the ITR believes the wrong reachability status of the ETR's RLOC until RLOC-probing detects the correct status. This time frame is on the order of 10s of seconds. This specific attack can be mitigated by preventing RLOC spoofing in the network by deploying uRPF BCP 38 [RFC2827]. In addition and in order to exploit this

vulnerability, the off-path attacker must send echo-nonce packets at high rate. If the nonces have never been requested by the ITR, it can protect itself from erroneous reachability attacks.

A LISP-specific uRPF check is also possible. When decapsulating, an ETR can check that the source EID and RLOC are valid EID-to-RLOC mappings by checking the Mapping System.

Map-Versioning is a Data-Plane mechanism used to signal a peering xTR that a local EID-to-RLOC mapping has been updated, so that the peering xTR uses LISP Control-Plane signaling message to retrieve a fresh mapping. This can be used by an attacker to forge the map-versioning field of a LISP encapsulated header and force an excessive amount of signaling between xTRs that may overload them. Further security considerations on Map-Versioning can be found in [I-D.ietf-lisp-6834bis].

Locator-Status-Bits, echo-nonce and map-versioning MUST NOT be used over the public Internet and SHOULD only be used in trusted and closed deployments. In addition Locator-Status-Bits SHOULD be coupled with map-versioning to prevent race conditions where Locator-Status-Bits are interpreted as referring to different RLOCs than intended.

LISP implementations and deployments which permit outer header fragments of IPv6 LISP encapsulated packets as a means of dealing with MTU issues should also use implementation techniques in ETRs to prevent this from being a DoS attack vector. Limits on the number of fragments awaiting reassembly at an ETR, RTR, or PETR, and the rate of admitting such fragments may be used.

17. Network Management Considerations

Considerations for network management tools exist so the LISP protocol suite can be operationally managed. These mechanisms can be found in [RFC7052] and [RFC6835].

18. Changes since RFC 6830

For implementation considerations, the following changes have been made to this document since RFC 6830 was published:

- * It is no longer mandated that a maximum number of 2 LISP headers be prepended to a packet. If there is an application need for more than 2 LISP headers, an implementation can support more. However, it is RECOMMENDED that a maximum of two LISP headers can be prepended to a packet.

- * The 3 reserved flag bits in the LISP header have been allocated for [RFC8061]. The low-order 2 bits of the 3-bit field (now named the KK bits) are used as a key identifier. The 1 remaining bit is still documented as reserved and unassigned.
- * Data-Plane gleaning for creating map-cache entries has been made optional. Any ITR implementations that depend on or assume the remote ETR is gleaning should not do so. This does not create any interoperability problems since the control-plane map-cache population procedures are unilateral and are the typical method for map-cache population.
- * The bulk of the changes to this document which reduces its length are due to moving the LISP control-plane messaging and procedures to [I-D.ietf-lisp-rfc6833bis].

19. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to this Data-Plane LISP specification, in accordance with BCP 26 [RFC8126].

19.1. LISP UDP Port Numbers

The IANA registry has allocated UDP port number 4341 for the LISP Data-Plane. IANA has updated the description for UDP port 4341 as follows:

lisp-data	4341 udp	LISP Data Packets
-----------	----------	-------------------

20. References

20.1. Normative References

[I-D.ietf-lisp-6834bis]
Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", Work in Progress, Internet-Draft, draft-ietf-lisp-6834bis-10, 3 May 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-6834bis-10.txt>>.

[I-D.ietf-lisp-rfc6833bis]
Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, "Locator/ID Separation Protocol (LISP) Control-Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6833bis-31, 2 May 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6833bis-31.txt>>.

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [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>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, DOI 10.17487/RFC2983, October 2000, <<https://www.rfc-editor.org/info/rfc2983>>.
- [RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <<https://www.rfc-editor.org/info/rfc6040>>.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.
- [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>>.

- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <<https://www.rfc-editor.org/info/rfc8378>>.

20.2. Informative References

- [AFN] IANA, "Address Family Numbers", August 2016, <<http://www.iana.org/assignments/address-family-numbers>>.
- [CHIAPPA] Chiappa, J., "Endpoints and Endpoint names: A Proposed", 1999, <<http://mercury.lcs.mit.edu/~jnc/tech/endpoints.txt>>.
- [I-D.ietf-lisp-introduction] Cabellos, A. and D. S. (Ed.), "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-introduction-15, 20 September 2021, <<https://www.ietf.org/archive/id/draft-ietf-lisp-introduction-15.txt>>.
- [I-D.ietf-lisp-vpn] Moreno, V. and D. Farinacci, "LISP Virtual Private Networks (VPNs)", Work in Progress, Internet-Draft, draft-ietf-lisp-vpn-08, 18 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-vpn-08.txt>>.

- [I-D.ietf-tsvwg-datagram-plpmtud]
Fairhurst, G., Jones, T., Tuexen, M., Ruengeler, I., and T. Voelker, "Packetization Layer Path MTU Discovery for Datagram Transports", Work in Progress, Internet-Draft, draft-ietf-tsvwg-datagram-plpmtud-22, 10 June 2020, <<https://www.ietf.org/archive/id/draft-ietf-tsvwg-datagram-plpmtud-22.txt>>.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, DOI 10.17487/RFC1981, August 1996, <<https://www.rfc-editor.org/info/rfc1981>>.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, DOI 10.17487/RFC3056, February 2001, <<https://www.rfc-editor.org/info/rfc3056>>.
- [RFC3232] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, DOI 10.17487/RFC3232, January 2002, <<https://www.rfc-editor.org/info/rfc3232>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<https://www.rfc-editor.org/info/rfc3261>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.

- [RFC4459] Savola, P., "MTU and Fragmentation Issues with In-the-Network Tunneling", RFC 4459, DOI 10.17487/RFC4459, April 2006, <<https://www.rfc-editor.org/info/rfc4459>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, DOI 10.17487/RFC4984, September 2007, <<https://www.rfc-editor.org/info/rfc4984>>.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC6835] Farinacci, D. and D. Meyer, "The Locator/ID Separation Protocol Internet Groper (LIG)", RFC 6835, DOI 10.17487/RFC6835, January 2013, <<https://www.rfc-editor.org/info/rfc6835>>.
- [RFC6935] Eubanks, M., Chimento, P., and M. Westerlund, "IPv6 and UDP Checksums for Tunneled Packets", RFC 6935, DOI 10.17487/RFC6935, April 2013, <<https://www.rfc-editor.org/info/rfc6935>>.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.
- [RFC7052] Schudel, G., Jain, A., and V. Moreno, "Locator/ID Separation Protocol (LISP) MIB", RFC 7052, DOI 10.17487/RFC7052, October 2013, <<https://www.rfc-editor.org/info/rfc7052>>.
- [RFC7215] Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "Locator/Identifier Separation Protocol (LISP) Network Element Deployment Considerations", RFC 7215, DOI 10.17487/RFC7215, April 2014, <<https://www.rfc-editor.org/info/rfc7215>>.
- [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>>.

- [RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.

Appendix A. Acknowledgments

An initial thank you goes to Dave Oran for planting the seeds for the initial ideas for LISP. His consultation continues to provide value to the LISP authors.

A special and appreciative thank you goes to Noel Chiappa for providing architectural impetus over the past decades on separation of location and identity, as well as detailed reviews of the LISP architecture and documents, coupled with enthusiasm for making LISP a practical and incremental transition for the Internet.

The original authors would like to gratefully acknowledge many people who have contributed discussions and ideas to the making of this proposal. They include Scott Brim, Andrew Partan, John Zwiebel, Jason Schiller, Lixia Zhang, Dorian Kim, Peter Schoenmaker, Vijay Gill, Geoff Huston, David Conrad, Mark Handley, Ron Bonica, Ted Seely, Mark Townsley, Chris Morrow, Brian Weis, Dave McGrew, Peter Lothberg, Dave Thaler, Eliot Lear, Shane Amante, Ved Kafle, Olivier Bonaventure, Luigi Iannone, Robin Whittle, Brian Carpenter, Joel Halpern, Terry Manderson, Roger Jorgensen, Ran Atkinson, Stig Venaas, Iljitsch van Beijnum, Roland Bless, Dana Blair, Bill Lynch, Marc Woolward, Damien Saucez, Damian Lezama, Attilla De Groot, Parantap Lahiri, David Black, Roque Gagliano, Isidor Kouvelas, Jesper Skriver, Fred Templin, Margaret Wasserman, Sam Hartman, Michael Hofling, Pedro Marques, Jari Arkko, Gregg Schudel, Srinivas Subramanian, Amit Jain, Xu Xiaohu, Dharendra Trivedi, Yakov Rekhter, John Scudder, John Drake, Dimitri Papadimitriou, Ross Callon, Selina Heimlich, Job Snijders, Vina Ermagan, Fabio Maino, Victor Moreno, Chris White, Clarence Filsfils, Alia Atlas, Florin Coras and Alberto Rodriguez.

This work originated in the Routing Research Group (RRG) of the IRTF. An individual submission was converted into the IETF LISP working group document that became this RFC.

The LISP working group would like to give a special thanks to Jari Arkko, the Internet Area AD at the time that the set of LISP documents were being prepared for IESG last call, and for his meticulous reviews and detailed commentaries on the 7 working group last call documents progressing toward standards-track RFCs.

The current authors would like to give a sincere thank you to the people who help put LISP on standards track in the IETF. They include Joel Halpern, Luigi Iannone, Deborah Brungard, Fabio Maino, Scott Bradner, Kyle Rose, Takeshi Takahashi, Sarah Banks, Pete Resnick, Colin Perkins, Mirja Kuhlewind, Francis Dupont, Benjamin Kaduk, Eric Rescorla, Alvaro Retana, Alexey Melnikov, Alissa Cooper, Suresh Krishnan, Alberto Rodriguez-Natal, Vina Ermagen, Mohamed Boucadair, Brian Trammell, Sabrina Tanamal, and John Drake. The contributions they offered greatly added to the security, scale, and robustness of the LISP architecture and protocols.

Appendix B. Document Change Log

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-ietf-lisp-rfc6830bis-38

- * Posted May 2022.
- * Removed detailed paragraphs in Section 13.2 that is duplicated text from [I-D.ietf-lisp-6834bis], which is the authoritative source for Map Versioning.

B.2. Changes to draft-ietf-lisp-rfc6830bis-37

- * Posted May 2022.
- * Added references to 6834bis instead of pointing text to Section 13.2. This is so we can advance the Map-Versioning draft rfc6834bis to proposed standard.

B.3. Changes to draft-ietf-lisp-rfc6830bis-28

- * Posted November 2019.
- * Fixed how LSB behave in the presence of new/removed locators.
- * Added ETR synchronization requirements when using Map-Versioning.
- * Fixed a large set of minor comments and edits.

B.4. Changes to draft-ietf-lisp-rfc6830bis-27

- * Posted April 2019 post telechat.
- * Made editorial corrections per Warren's suggestions.
- * Put in suggested text from Luigi that Mirja agreed with.
- * LSB, Echo-Nonce and Map-Versioning SHOULD be only used in closed environments.
- * Removed paragraph stating that Instance-ID can be 32-bit in the control-plane.
- * 6831/8378 are now normative.
- * Rewritten Security Considerations according to the changes.
- * Stated that LSB SHOULD be coupled with Map-Versioning.

B.5. Changes to draft-ietf-lisp-rfc6830bis-26

- * Posted late October 2018.
- * Changed description about "reserved" bits to state "reserved and unassigned".

B.6. Changes to draft-ietf-lisp-rfc6830bis-25

- * Posted mid October 2018.
- * Added more to the Security Considerations section with discussion about echo-nonce attacks.

B.7. Changes to draft-ietf-lisp-rfc6830bis-24

- * Posted mid October 2018.
- * Final editorial changes for Eric and Ben.

B.8. Changes to draft-ietf-lisp-rfc6830bis-23

- * Posted early October 2018.
- * Added an applicability statement in section 1 to address security concerns from Telechat.

- B.9. Changes to draft-ietf-lisp-rfc6830bis-22
- * Posted early October 2018.
 - * Changes to reflect comments post Telechat.
- B.10. Changes to draft-ietf-lisp-rfc6830bis-21
- * Posted late-September 2018.
 - * Changes to reflect comments from Sep 27th Telechat.
- B.11. Changes to draft-ietf-lisp-rfc6830bis-20
- * Posted late-September 2018.
 - * Fix old reference to RFC3168, changed to RFC6040.
- B.12. Changes to draft-ietf-lisp-rfc6830bis-19
- * Posted late-September 2018.
 - * More editorial changes.
- B.13. Changes to draft-ietf-lisp-rfc6830bis-18
- * Posted mid-September 2018.
 - * Changes to reflect comments from Secdir review (Mirja).
- B.14. Changes to draft-ietf-lisp-rfc6830bis-17
- * Posted September 2018.
 - * Indicate in the "Changes since RFC 6830" section why the document has been shortened in length.
 - * Make reference to RFC 8085 about UDP congestion control.
 - * More editorial changes from multiple IESG reviews.
- B.15. Changes to draft-ietf-lisp-rfc6830bis-16
- * Posted late August 2018.
 - * Distinguish the message type names between ICMP for IPv4 and ICMP for IPv6 for handling MTU issues.

B.16. Changes to draft-ietf-lisp-rfc6830bis-15

- * Posted August 2018.
- * Final editorial changes before RFC submission for Proposed Standard.
- * Added section "Changes since RFC 6830" so implementers are informed of any changes since the last RFC publication.

B.17. Changes to draft-ietf-lisp-rfc6830bis-14

- * Posted July 2018 IETF week.
- * Put obsolete of RFC 6830 in Intro section in addition to abstract.

B.18. Changes to draft-ietf-lisp-rfc6830bis-13

- * Posted March IETF Week 2018.
- * Clarified that a new nonce is required per RLOC.
- * Removed 'Clock Sweep' section. This text must be placed in a new OAM document.
- * Some references changed from normative to informative

B.19. Changes to draft-ietf-lisp-rfc6830bis-12

- * Posted July 2018.
- * Fixed Luigi editorial comments to ready draft for RFC status.

B.20. Changes to draft-ietf-lisp-rfc6830bis-11

- * Posted March 2018.
- * Removed sections 16, 17 and 18 (Mobility, Deployment and Traceroute considerations). This text must be placed in a new OAM document.

B.21. Changes to draft-ietf-lisp-rfc6830bis-10

- * Posted March 2018.
- * Updated section 'Router Locator Selection' stating that the Data-Plane MUST follow what's stored in the Map-Cache (priorities and weights).

- * Section 'Routing Locator Reachability': Removed bullet point 2 (ICMP Network/Host Unreachable), 3 (hints from BGP), 4 (ICMP Port Unreachable), 5 (receive a Map-Reply as a response) and RLOC probing
- * Removed 'Solicit-Map Request'.

B.22. Changes to draft-ietf-lisp-rfc6830bis-09

- * Posted January 2018.
- * Add more details in section 5.3 about DSCP processing during encapsulation and decapsulation.
- * Added clarity to definitions in the Definition of Terms section from various commenters.
- * Removed PA and PI definitions from Definition of Terms section.
- * More editorial changes.
- * Removed 4342 from IANA section and move to RFC6833 IANA section.

B.23. Changes to draft-ietf-lisp-rfc6830bis-08

- * Posted January 2018.
- * Remove references to research work for any protocol mechanisms.
- * Document scanned to make sure it is RFC 2119 compliant.
- * Made changes to reflect comments from document WG shepherd Luigi Iannone.
- * Ran IDNITs on the document.

B.24. Changes to draft-ietf-lisp-rfc6830bis-07

- * Posted November 2017.
- * Rephrase how Instance-IDs are used and don't refer to [RFC1918] addresses.

B.25. Changes to draft-ietf-lisp-rfc6830bis-06

- * Posted October 2017.
- * Put RTR definition before it is used.

- * Rename references that are now working group drafts.
 - * Remove "EIDs MUST NOT be used as used by a host to refer to other hosts. Note that EID blocks MAY LISP RLOCs".
 - * Indicate what address-family can appear in data packets.
 - * ETRs may, rather than will, be the ones to send Map-Replies.
 - * Recommend, rather than mandate, max encapsulation headers to 2.
 - * Reference VPN draft when introducing Instance-ID.
 - * Indicate that SMRs can be sent when ITR/ETR are in the same node.
 - * Clarify when private addresses can be used.
- B.26. Changes to draft-ietf-lisp-rfc6830bis-05
- * Posted August 2017.
 - * Make it clear that a Re-encapsulating Tunnel Router is an RTR.
- B.27. Changes to draft-ietf-lisp-rfc6830bis-04
- * Posted July 2017.
 - * Changed reference of IPv6 RFC2460 to RFC8200.
 - * Indicate that the applicability statement for UDP zero checksums over IPv6 adheres to RFC6936.
- B.28. Changes to draft-ietf-lisp-rfc6830bis-03
- * Posted May 2017.
 - * Move the control-plane related codepoints in the IANA Considerations section to RFC6833bis.
- B.29. Changes to draft-ietf-lisp-rfc6830bis-02
- * Posted April 2017.
 - * Reflect some editorial comments from Damien Sausez.

B.30. Changes to draft-ietf-lisp-rfc6830bis-01

- * Posted March 2017.
- * Include references to new RFCs published.
- * Change references from RFC6833 to RFC6833bis.
- * Clarified LCAF text in the IANA section.
- * Remove references to "experimental".

B.31. Changes to draft-ietf-lisp-rfc6830bis-00

- * Posted December 2016.
- * Created working group document from draft-farinacci-lisp-rfc6830-00 individual submission. No other changes made.

Authors' Addresses

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

Vince Fuller
vaf.net Internet Consulting
Email: vince.fuller@gmail.com

Dave Meyer
1-4-5.net
Email: dmm@1-4-5.net

Darrel Lewis
Cisco Systems
170 Tasman Drive
San Jose, CA
United States of America
Email: darlewis@cisco.com

Albert Cabellos
UPC/BarcelonaTech
Campus Nord, C. Jordi Girona 1-3
Barcelona Catalunya
Spain
Email: acabello@ac.upc.edu

Network Working Group
Internet-Draft
Obsoletes: 6830, 6833 (if approved)
Intended status: Standards Track
Expires: November 3, 2022

D. Farinacci
lispers.net
F. Maino
Cisco Systems
V. Fuller
vaf.net Internet Consulting
A. Cabellos (Ed.)
UPC/BarcelonaTech
May 2, 2022

Locator/ID Separation Protocol (LISP) Control-Plane
draft-ietf-lisp-rfc6833bis-31

Abstract

This document describes the Control-Plane and Mapping Service for the Locator/ID Separation Protocol (LISP), implemented by two types of LISP-speaking devices -- the LISP Map-Resolver and LISP Map-Server -- that provides a simplified "front end" for one or more Endpoint ID to Routing Locator mapping databases.

By using this Control-Plane service interface and communicating with Map-Resolvers and Map-Servers, LISP Ingress Tunnel Routers (ITRs) and Egress Tunnel Routers (ETRs) are not dependent on the details of mapping database systems, which facilitates modularity with different database designs. Since these devices implement the "edge" of the LISP Control-Plane infrastructure, connecting EID addressable nodes of a LISP site, it the implementation and operational complexity of the overall cost and effort of deploying LISP.

This document obsoletes RFC 6830 and RFC 6833.

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 3, 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	3
1.1. Scope of Applicability	5
2. Requirements Notation	5
3. Definition of Terms	5
4. Basic Overview	7
5. LISP IPv4 and IPv6 Control-Plane Packet Formats	8
5.1. LISP Control Packet Type Allocations	11
5.2. Map-Request Message Format	12
5.3. EID-to-RLOC UDP Map-Request Message	14
5.4. Map-Reply Message Format	16
5.5. EID-to-RLOC UDP Map-Reply Message	20
5.6. Map-Register Message Format	23
5.7. Map-Notify/Map-Notify-Ack Message Format	27
5.8. Encapsulated Control Message Format	29
6. Changing the Contents of EID-to-RLOC Mappings	31
6.1. Solicit-Map-Request (SMR)	31
7. Routing Locator Reachability	32
7.1. RLOC-Probing Algorithm	33
8. Interactions with Other LISP Components	34
8.1. ITR EID-to-RLOC Mapping Resolution	34
8.2. EID-Prefix Configuration and ETR Registration	35
8.3. Map-Server Processing	37
8.4. Map-Resolver Processing	37
8.4.1. Anycast Operation	38
9. Security Considerations	38
10. Privacy Considerations	40
11. Changes since RFC 6833	41
12. IANA Considerations	41
12.1. LISP UDP Port Numbers	42

12.2.	LISP Packet Type Codes	42
12.3.	LISP Map-Reply EID-Record Action Codes	42
12.4.	LISP Address Type Codes	43
12.5.	LISP Algorithm ID Numbers	43
12.6.	LISP Bit Flags	44
13.	References	47
13.1.	Normative References	47
13.2.	Informative References	48
Appendix A.	Acknowledgments	53
Appendix B.	Document Change Log	53
B.1.	Changes to draft-ietf-lisp-rfc6833bis-31	53
B.2.	Changes to draft-ietf-lisp-rfc6833bis-26	53
B.3.	Changes to draft-ietf-lisp-rfc6833bis-25	53
B.4.	Changes to draft-ietf-lisp-rfc6833bis-24	54
B.5.	Changes to draft-ietf-lisp-rfc6833bis-23	54
B.6.	Changes to draft-ietf-lisp-rfc6833bis-22	54
B.7.	Changes to draft-ietf-lisp-rfc6833bis-21	54
B.8.	Changes to draft-ietf-lisp-rfc6833bis-20	55
B.9.	Changes to draft-ietf-lisp-rfc6833bis-19	55
B.10.	Changes to draft-ietf-lisp-rfc6833bis-18	55
B.11.	Changes to draft-ietf-lisp-rfc6833bis-17	55
B.12.	Changes to draft-ietf-lisp-rfc6833bis-16	55
B.13.	Changes to draft-ietf-lisp-rfc6833bis-15	56
B.14.	Changes to draft-ietf-lisp-rfc6833bis-14	56
B.15.	Changes to draft-ietf-lisp-rfc6833bis-13	56
B.16.	Changes to draft-ietf-lisp-rfc6833bis-12	56
B.17.	Changes to draft-ietf-lisp-rfc6833bis-11	56
B.18.	Changes to draft-ietf-lisp-rfc6833bis-10	56
B.19.	Changes to draft-ietf-lisp-rfc6833bis-09	57
B.20.	Changes to draft-ietf-lisp-rfc6833bis-08	57
B.21.	Changes to draft-ietf-lisp-rfc6833bis-07	57
B.22.	Changes to draft-ietf-lisp-rfc6833bis-06	58
B.23.	Changes to draft-ietf-lisp-rfc6833bis-05	58
B.24.	Changes to draft-ietf-lisp-rfc6833bis-04	58
B.25.	Changes to draft-ietf-lisp-rfc6833bis-03	58
B.26.	Changes to draft-ietf-lisp-rfc6833bis-02	59
B.27.	Changes to draft-ietf-lisp-rfc6833bis-01	59
B.28.	Changes to draft-ietf-lisp-rfc6833bis-00	59
B.29.	Changes to draft-farinacci-lisp-rfc6833bis-00	59
Authors' Addresses	60

1. Introduction

The Locator/ID Separation Protocol [I-D.ietf-lisp-rfc6830bis] (see also [I-D.ietf-lisp-introduction]) specifies an architecture and mechanism for dynamic tunneling by logically separating the addresses currently used by IP in two separate name spaces: Endpoint IDs (EIDs), used within sites; and Routing Locators (RLOCs), used on the

transit networks that make up the Internet infrastructure. To achieve this separation, LISP defines protocol mechanisms for mapping from EIDs to RLOCs. In addition, LISP assumes the existence of a database to store and propagate those mappings across mapping system nodes. Several such databases have been proposed; among them are the Content distribution Overlay Network Service for LISP-NERD (a Not-so-novel EID-to-RLOC Database) [RFC6837], LISP Alternative Logical Topology (LISP-ALT) [RFC6836], and LISP Delegated Database Tree (LISP-DDT) [RFC8111].

The LISP Mapping Service defines two types of LISP-speaking devices: the Map-Resolver, which accepts Map-Requests from an Ingress Tunnel Router (ITR) and "resolves" the EID-to-RLOC mapping using a mapping database; and the Map-Server, which learns authoritative EID-to-RLOC mappings from an Egress Tunnel Router (ETR) and publishes them in a database.

This LISP Control-Plane Mapping Service can be used by many different encapsulation-based or translation-based Data-Planes which include but are not limited to the ones defined in LISP RFC 6830bis [I-D.ietf-lisp-rfc6830bis], LISP-GPE [I-D.ietf-lisp-gpe], VXLAN [RFC7348], VXLAN-GPE [I-D.ietf-nvo3-vxlan-gpe], GRE [RFC2890], GTP [GTP-3GPP], ILA [I-D.herbert-intarea-ila], and Segment Routing (SRv6) [RFC8402].

Conceptually, LISP Map-Servers share some of the same basic configuration and maintenance properties as Domain Name System (DNS) [RFC1035] servers; likewise, Map-Resolvers are conceptually similar to DNS caching resolvers. With this in mind, this specification borrows familiar terminology (resolver and server) from the DNS specifications.

Note this document doesn't assume any particular database mapping infrastructure to illustrate certain aspects of Map-Server and Map-Resolver operation. The Mapping Service interface can (and likely will) be used by ITRs and ETRs to access other mapping database systems as the LISP infrastructure evolves.

LISP is not intended to address problems of connectivity and scaling on behalf of arbitrary communicating parties. Relevant situations are described in the scoping section of the introduction to [I-D.ietf-lisp-rfc6830bis].

This document obsoletes RFC 6830 and 6833.

1.1. Scope of Applicability

LISP was originally developed to address the Internet-wide route scaling problem [RFC4984]. While there are a number of approaches of interest for that problem, as LISP has been developed and refined, a large number of other LISP uses have been found and are being used. As such, the design and development of LISP has changed so as to focus on these use cases. The common property of these uses is a large set of cooperating entities seeking to communicate over the public Internet or other large underlay IP infrastructures, while keeping the addressing and topology of the cooperating entities separate from the underlay and Internet topology, routing, and addressing.

When communicating over the public Internet, deployers MUST consider the following guidelines:

1. LISP-SEC MUST be implemented [I-D.ietf-lisp-sec]. This means that the S-bit MUST be set in the Map-Reply (Section 5.4), Map-Register (Section 5.6) and Encapsulated Control messages (Section 5.8).
2. Implementations SHOULD use the 'HMAC-SHA256-128+HKDF-SHA256' as the Algorithm ID (Section 12.5) in Map-Register message (Section 5.6), and MUST NOT use 'None' or 'HMAC-SHA-1-96-None' as Algorithm ID (Section 12.5) in the Map-Register message (Section 5.6)

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Definition of Terms

Map-Server: A network infrastructure component that learns of EID-Prefix mapping entries from an ETR, via the registration mechanism described below, or some other authoritative source if one exists. A Map-Server publishes these EID-Prefixes in a mapping database.

Map-Request: A LISP Map-Request is a Control-Plane message to query the mapping system to resolve an EID. A LISP Map-Request can also be sent to an RLOC to test for reachability and to exchange security keys between an encapsulator and a decapsulator. This type of Map-Request is also known as an RLOC-Probe Request.

Map-Reply: A LISP Map-Reply is a Control-Plane message returned in response to a Map-Request sent to the mapping system when resolving an EID. A LISP Map-Reply can also be returned by a decapsulator in response to a Map-Request sent by an encapsulator to test for reachability. This type of Map-Reply is known as a RLOC-Probe Reply.

Encapsulated Map-Request: A LISP Map-Request carried within an Encapsulated Control Message (ECM), which has an additional LISP header prepended. Sent to UDP destination port 4342. The "outer" addresses are routable IP addresses, also known as RLOCs. Used by an ITR when sending to a Map-Resolver and by a Map-Server when forwarding a Map-Request to an ETR.

Map-Resolver: A network infrastructure component that accepts LISP Encapsulated (ECM) Map-Requests, typically from an ITR, and determines whether or not the destination IP address is part of the EID namespace; if it is not, a Negative Map-Reply is returned. Otherwise, the Map-Resolver finds the appropriate EID-to-RLOC mapping by consulting a mapping database system.

Negative Map-Reply: A LISP Map-Reply that contains an empty Locator-Set. Returned in response to a Map-Request if the destination EID is not registered in the mapping system, is policy denied or fails authentication.

Map-Register message: A LISP message sent by an ETR to a Map-Server to register its associated EID-Prefixes. In addition to the set of EID-Prefixes to register, the message includes one or more RLOCs to reach ETR(s). The Map-Server uses these RLOCs when forwarding Map-Requests (re-formatted as Encapsulated Map-Requests). An ETR MAY request that the Map-Server answer Map-Requests on its behalf by setting the "proxy Map-Reply" flag (P-bit) in the message.

Map-Notify message: A LISP message sent by a Map-Server to an ETR to confirm that a Map-Register has been received and processed. An ETR requests that a Map-Notify be returned by setting the "want-map-notify" flag (M-bit) in the Map-Register message. Unlike a Map-Reply, a Map-Notify uses UDP port 4342 for both source and destination. Map-Notify messages are also sent to ITRs by Map-Servers when there are RLOC-set changes.

For definitions of other terms, notably Ingress Tunnel Router (ITR), Egress Tunnel Router (ETR), and Re-encapsulating Tunnel Router (RTR), refer to the LISP Data-Plane specification [I-D.ietf-lisp-rfc6830bis].

4. Basic Overview

A Map-Server is a device that publishes EID-Prefixes in a LISP mapping database on behalf of a set of ETRs. When it receives a Map Request (typically originating from an ITR), it consults the mapping database to find an ETR that can answer with the set of RLOCs for an EID-Prefix. To publish its EID-Prefixes, an ETR periodically sends Map-Register messages to the Map-Server. A Map-Register message contains a list of EID-Prefixes plus a set of RLOCs that can be used to reach the ETRs.

When LISP-ALT [RFC6836] is used as the mapping database, a Map-Server connects to the ALT network and acts as a "last-hop" ALT-Router. Intermediate ALT-Routers forward Map-Requests to the Map-Server that advertises a particular EID-Prefix, and the Map-Server forwards them to the owning ETR, which responds with Map-Reply messages.

When LISP-DDT [RFC8111] is used as the mapping database, a Map-Server sends the final Map-Referral messages from the Delegated Database Tree.

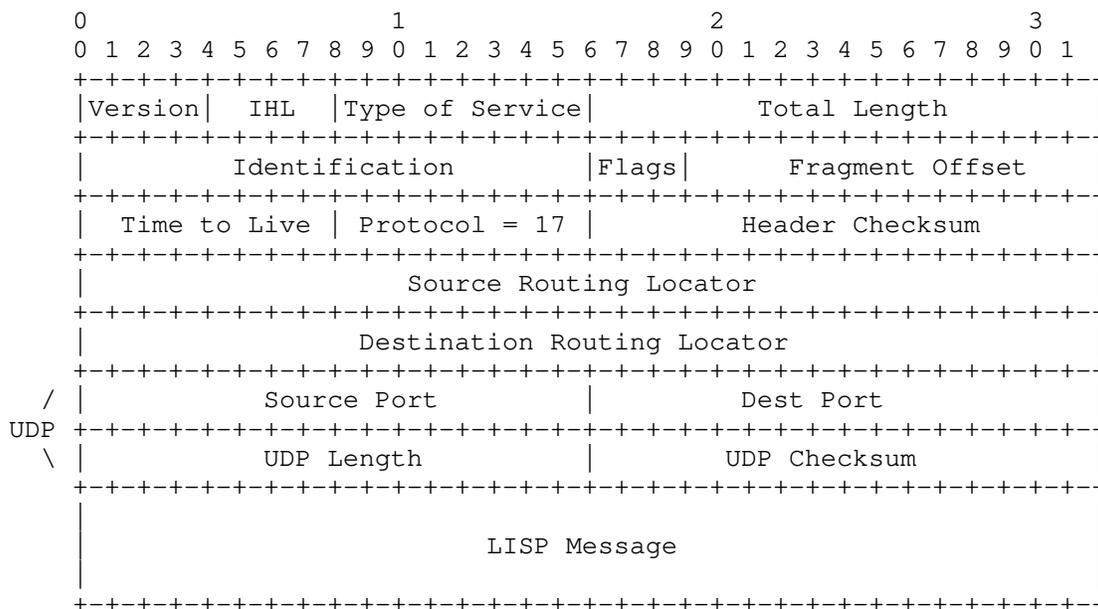
A Map-Resolver receives Encapsulated Map-Requests from its client ITRs and uses a mapping database system to find the appropriate ETR to answer those requests. On a LISP-ALT network, a Map-Resolver acts as a "first-hop" ALT-Router. It has Generic Routing Encapsulation (GRE) tunnels configured to other ALT-Routers and uses BGP to learn paths to ETRs for different prefixes in the LISP-ALT database. The Map-Resolver uses this path information to forward Map-Requests over the ALT to the correct ETRs. On a LISP-DDT network [RFC8111], a Map-Resolver maintains a referral-cache and acts as a "first-hop" DDT-node. The Map-Resolver uses the referral information to forward Map-Requests.

Note that while it is conceivable that a Map-Resolver could cache responses to improve performance, issues surrounding cache management would need to be resolved so that doing so will be reliable and practical. In this specification, Map-Resolvers will operate only in a non-caching mode, decapsulating and forwarding Encapsulated Map Requests received from ITRs. Any specification of caching functionality is out of scope for this document.

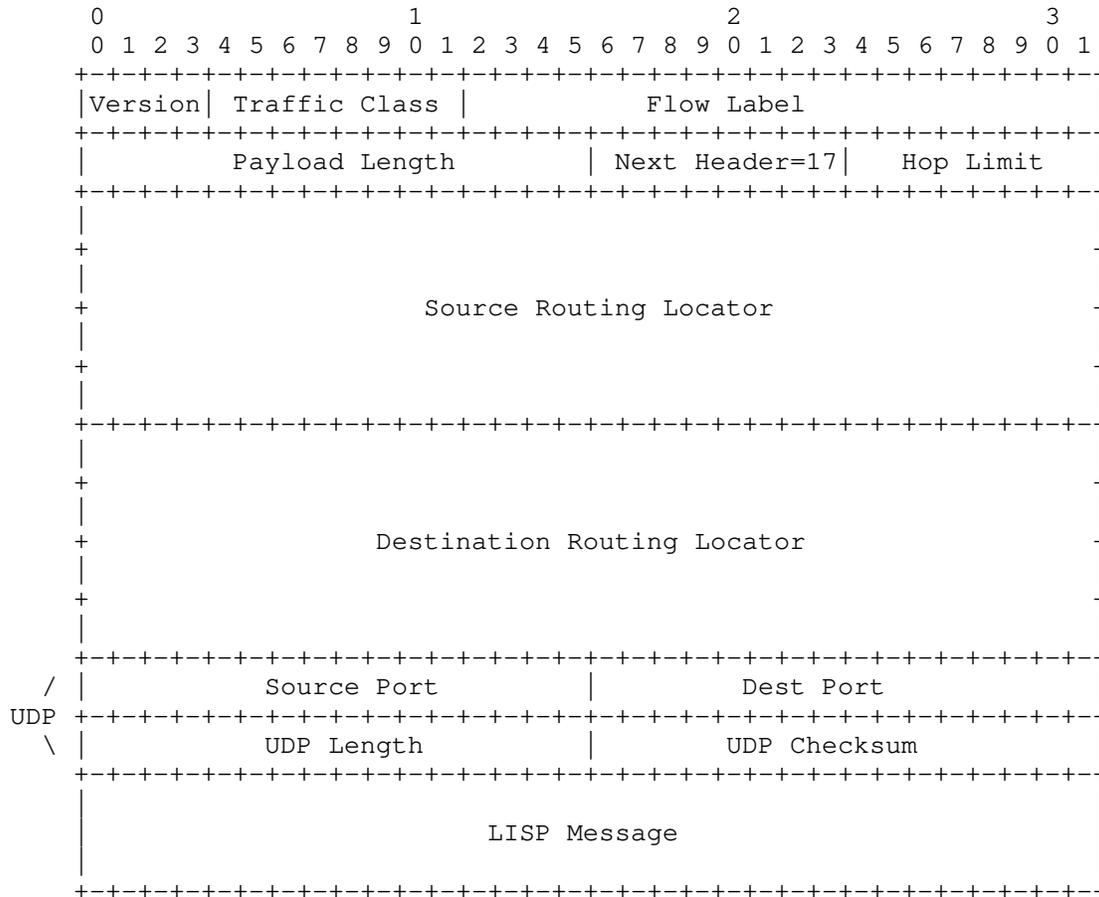
Note that a single device can implement the functions of both a Map-Server and a Map-Resolver, and in many cases the functions will be co-located in that way. Also, there can be ALT-only nodes and DDT-only nodes, when LISP-ALT and LISP-DDT are used, respectively, to connecting Map-Resolvers and Map-Servers together to make up the Mapping System.

5. LISP IPv4 and IPv6 Control-Plane Packet Formats

The following UDP packet formats are used by the LISP control plane.



IPv4 UDP LISP Control Message



IPv6 UDP LISP Control Message

When a UDP Map-Request, Map-Register, or Map-Notify (when used as a notification message) are sent, the UDP source port is chosen by the sender and the destination UDP port number is set to 4342. When a UDP Map-Reply, Map-Notify (when used as an acknowledgement to a Map-Register), or Map-Notify-Ack are sent, the source UDP port number is set to 4342 and the destination UDP port number is copied from the source port of either the Map-Request or the invoking data packet. Implementations MUST be prepared to accept packets when either the source port or destination UDP port is set to 4342 due to NATs changing port number values.

The 'UDP Length' field will reflect the length of the UDP header and the LISP Message payload. LISP is expected to be deployed by cooperating entities communicating over underlays. Deployers are

expected to set the MTU according to the specific deployment guidelines to prevent fragmentation of either the inner packet or the outer encapsulated packet. For deployments not aware of the underlay restrictions on path MTU, the message size MUST be limited to 576 bytes for IPv4 or 1280 bytes for IPv6 -considering the entire IP packet- as outlined in [RFC8085].

The UDP checksum is computed and set to non-zero for all messages sent to or from port 4342. It MUST be checked on receipt, and if the checksum fails, the control message MUST be dropped [RFC1071].

The format of control messages includes the UDP header so the checksum and length fields can be used to protect and delimit message boundaries.

5.1. LISP Control Packet Type Allocations

This section defines the LISP control message formats and summarizes for IANA the LISP Type codes assigned by this document. For completeness, the summary below includes the LISP Shared Extension Message assigned by [I-D.ietf-lisp-rfc8113bis]. Message type definitions are:

Reserved:	0	b'0000'
LISP Map-Request:	1	b'0001'
LISP Map-Reply:	2	b'0010'
LISP Map-Register:	3	b'0011'
LISP Map-Notify:	4	b'0100'
LISP Map-Notify-Ack:	5	b'0101'
LISP Map-Referral:	6	b'0110'
Unassigned	7	b'0111'
LISP Encapsulated Control Message:	8	b'1000'
Unassigned	9-14	b'1001' - b'1110'
LISP Shared Extension Message:	15	b'1111'

Protocol designers experimenting with new message formats are recommended to use the LISP Shared Extension Message Type described in [I-D.ietf-lisp-rfc8113bis].

All LISP Control-Plane messages use Address Family Identifiers (AFI) [AFI] or LISP Canonical Address Format (LCAF) [RFC8060] formats to encode either fixed or variable length addresses. This includes explicit fields in each control message or part of EID-records or RLOC-records in commonly formatted messages. LISP control-plane messages that include an unrecognized AFI MUST be dropped and the event MUST be logged.

The LISP control-plane describes how other data-planes can encode messages to support the Soliciting of Map-Requests as well as RLOC-probing procedures.

- p: This is the Pitr bit. This bit is set to 1 when a Pitr sends a Map-Request. The use of this bit is deployment-specific.
- s: This is the SMR-invoked bit. This bit is set to 1 when an xTR is sending a Map-Request in response to a received SMR-based Map-Request.
- R: This reserved and unassigned bit MUST be set to 0 on transmit and MUST be ignored on receipt.
- Rsvd: This field MUST be set to 0 on transmit and MUST be ignored on receipt.
- L: This is the local-xtr bit. It is used by an xTR in a LISP site to tell other xTRs in the same site that it is part of the RLOC-set for the LISP site. The L-bit is set to 1 when the RLOC is the sender's IP address.
- D: This is the dont-map-reply bit. It is used in the SMR procedure described in Section 6.1. When an xTR sends an SMR message, it doesn't need a Map-Reply returned. When this bit is set, the receiver of the Map-Request does not return a Map-Reply.
- IRC: This 5-bit field is the ITR-RLOC Count, which encodes the additional number of ('ITR-RLOC-AFI', 'ITR-RLOC Address') fields present in this message. At least one (ITR-RLOC-AFI, ITR-RLOC-Address) pair MUST be encoded. Multiple 'ITR-RLOC Address' fields are used, so a Map-Replier can select which destination address to use for a Map-Reply. The IRC value ranges from 0 to 31. For a value of 0, there is 1 ITR-RLOC address encoded; for a value of 1, there are 2 ITR-RLOC addresses encoded, and so on up to 31, which encodes a total of 32 ITR-RLOC addresses.
- Record Count: This is the number of records in this Map-Request message. A record is comprised of the portion of the packet that is labeled 'Rec' above and occurs the number of times equal to Record Count. For this version of the protocol, a receiver MUST accept and process Map-Requests that contain one or more records, but a sender MUST only send Map-Requests containing one record.
- Nonce: This is an 8-octet random value created by the sender of the Map-Request. This nonce will be returned in the Map-Reply. The nonce is used as an index to identify the corresponding Map-Request when a Map-Reply message is received. The nonce MUST be generated by a properly seeded pseudo-random source, see as an example [RFC4086].

Source-EID-AFI: This is the address family of the 'Source EID Address' field.

Source EID Address: This is the EID of the source host that originated the packet that caused the Map-Request. When Map-Requests are used for refreshing a Map-Cache entry or for RLOC-Probing, an AFI value 0 is used and this field is of zero length.

ITR-RLOC-AFI: This is the address family of the 'ITR-RLOC Address' field that follows this field.

ITR-RLOC Address: This is used to give the ETR the option of selecting the destination address from any address family for the Map-Reply message. This address MUST be a routable RLOC address of the sender of the Map-Request message.

EID mask-len: This is the mask length for the EID-Prefix.

EID-Prefix-AFI: This is the address family of the EID-Prefix according to [AFI] and [RFC8060].

EID-Prefix: This prefix address length is 4 octets for an IPv4 address family and 16 octets for an IPv6 address family when the EID-Prefix-AFI is 1 or 2, respectively. For other AFIs [AFI], the address length varies and for the LCAF AFI the format is defined in [RFC8060]. When a Map-Request is sent by an ITR because a data packet is received for a destination where there is no mapping entry, the EID-Prefix is set to the destination IP address of the data packet, and the 'EID mask-len' is set to 32 or 128 for IPv4 or IPv6, respectively. When an xTR wants to query a site about the status of a mapping it already has cached, the EID-Prefix used in the Map-Request has the same mask-length as the EID-Prefix returned from the site when it sent a Map-Reply message.

Map-Reply Record: When the M-bit is set, this field is the size of a single "Record" in the Map-Reply format. This Map-Reply record contains the EID-to-RLOC mapping entry associated with the Source EID. This allows the ETR that will receive this Map-Request to cache the data if it chooses to do so. It is important to note that this mapping has not been validated by the Mapping System.

5.3. EID-to-RLOC UDP Map-Request Message

A Map-Request is sent from an ITR when it needs a mapping for an EID, wants to test an RLOC for reachability, or wants to refresh a mapping before TTL expiration. For the initial case, the destination IP address used for the Map-Request is the data packet's destination address (i.e., the destination EID) that had a mapping cache lookup

failure. For the latter two cases, the destination IP address used for the Map-Request is one of the RLOC addresses from the Locator-Set of the Map-Cache entry. The source address is either an IPv4 or IPv6 RLOC address, depending on whether the Map-Request is using an IPv4 or IPv6 header, respectively. In all cases, the UDP source port number for the Map-Request message is a 16-bit value selected by the ITR/PITR, and the UDP destination port number is set to the well-known destination port number 4342. A successful Map-Reply, which is one that has a nonce that matches an outstanding Map-Request nonce, will update the cached set of RLOCs associated with the EID-Prefix range.

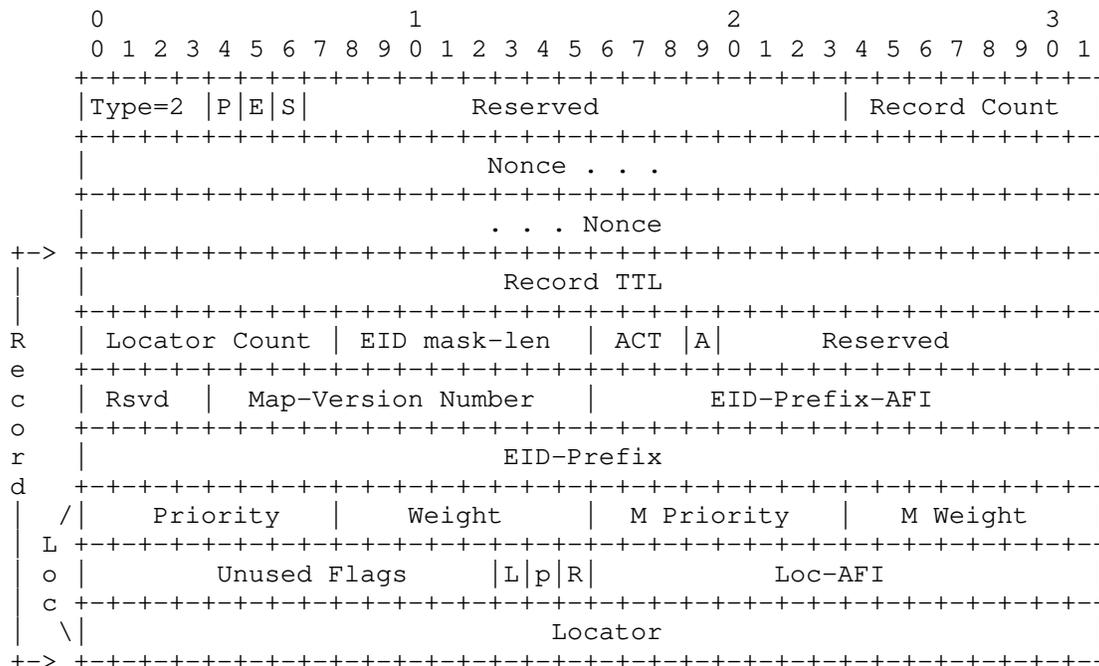
One or more Map-Request ('ITR-RLOC-AFI', 'ITR-RLOC-Address') fields MUST be filled in by the ITR. The number of fields (minus 1) encoded MUST be placed in the 'IRC' field. The ITR MAY include all locally configured Locators in this list or just provide one locator address from each address family it supports. If the ITR erroneously provides no ITR-RLOC addresses, the Map-Replier MUST drop the Map-Request.

Map-Requests can also be LISP encapsulated using UDP destination port 4342 with a LISP Type value set to "Encapsulated Control Message", when sent from an ITR to a Map-Resolver. Likewise, Map-Requests are LISP encapsulated the same way from a Map-Server to an ETR. Details on Encapsulated Map-Requests and Map-Resolvers can be found in Section 5.8.

Map-Requests MUST be rate-limited to 1 per second per EID-prefix. After 10 retransmits without receiving the corresponding Map-Reply the sender MUST wait 30 seconds.

An ITR that is configured with mapping database information (i.e., it is also an ETR) MAY optionally include those mappings in a Map-Request. When an ETR configured to accept and verify such "piggybacked" mapping data receives such a Map-Request and it does not have this mapping in the Map-Cache, it MUST originate a "verifying Map-Request" through the mapping database to validate the "piggybacked" mapping data.

5.4. Map-Reply Message Format



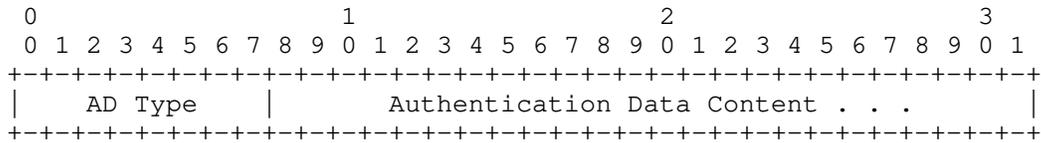
Packet field descriptions:

Type: 2 (Map-Reply)

P: This is the probe-bit, which indicates that the Map-Reply is in response to a Locator reachability probe Map-Request. The 'Nonce' field must contain a copy of the nonce value from the original Map-Request. See RLOC-probing Section 7.1 for more details. When the probe-bit is set to 1 in a Map-Reply message, the A-bit in each EID-record included in the message MUST be set to 1, otherwise MUST be silently discarded.

E: This bit indicates that the ETR that sends this Map-Reply message is advertising that the site is enabled for the Echo-Nonce Locator reachability algorithm. See Echo-Nonce [I-D.ietf-lisp-rtc6830bis] for more details.

S: This is the Security bit. When set to 1, the following authentication information will be appended to the end of the Map-Reply. The details can be found in [I-D.ietf-lisp-sec].



Reserved: This unassigned field MUST be set to 0 on transmit and MUST be ignored on receipt.

Record Count: This is the number of records in this reply message. A record is comprised of that portion of the packet labeled 'Record' above and occurs the number of times equal to Record Count. Note that the reply count can be larger than the requested count, for instance when more-specifics are present.

Nonce: This 64-bit value from the Map-Request is echoed in this 'Nonce' field of the Map-Reply.

Record TTL: This is the time in minutes the recipient of the Map-Reply can store the mapping. If the TTL is 0, the entry MUST be removed from the cache immediately. If the value is 0xffffffff, the recipient can decide locally how long to store the mapping.

Locator Count: This is the number of Locator entries in the given Record. A Locator entry comprises what is labeled above as 'Loc'. The Locator count can be 0, indicating that there are no Locators for the EID-Prefix.

EID mask-len: This is the mask length for the EID-Prefix.

ACT: This 3-bit field describes Negative Map-Reply actions. In any other message type, these bits are set to 0 and ignored on receipt. These bits are used only when the 'Locator Count' field is set to 0. The action bits are encoded only in Map-Reply messages. They are used to tell an ITR or PITR why a empty locator-set was returned from the mapping system and how it stores the map-cache entry. See Section 12.3 for additional information.

(0) No-Action: The Map-Cache is kept alive, and no packet encapsulation occurs.

(1) Natively-Forward: The packet is not encapsulated or dropped but natively forwarded.

- (2) Send-Map-Request: The Map-Cache entry is created and flagged that any packet matching this entry invokes sending a Map-Request.
 - (3) Drop/No-Reason: A packet that matches this Map-Cache entry is dropped. An ICMP Destination Unreachable message SHOULD be sent.
 - (4) Drop/Policy-Denied: A packet that matches this Map-Cache entry is dropped. The reason for the Drop action is that a Map-Request for the target-EID is being policy denied by either an xTR or the mapping system.
 - (5) Drop/Authentication-Failure: A packet that matches this Map-Cache entry is dropped. The reason for the Drop action is that a Map-Request for the target-EID fails an authentication verification-check by either an xTR or the mapping system.
- A: The Authoritative bit MAY only be set to 1 by an ETR. A Map-Server generating Map-Reply messages as a proxy MUST NOT set the A-bit to 1. This bit indicates to the requesting ITRs if the Map-Reply was originated by a LISP node managed at the site that owns the EID-Prefix.

Map-Version Number: When this 12-bit value in an EID-record of a Map-Reply message is non-zero, follow the procedures in [I-D.ietf-lisp-6834bis] for details.

EID-Prefix-AFI: Address family of the EID-Prefix according to [AFI] and [RFC8060].

EID-Prefix: This prefix is 4 octets for an IPv4 address family and 16 octets for an IPv6 address family.

Priority: Each RLOC is assigned a unicast Priority. Lower values are more preferable. When multiple RLOCs have the same Priority, they may be used in a load-split fashion. A value of 255 means the RLOC MUST NOT be used for unicast forwarding.

Weight: When priorities are the same for multiple RLOCs, the Weight indicates how to balance unicast traffic between them. Weight is encoded as a relative weight of total unicast packets that match the mapping entry. For example, if there are 4 Locators in a Locator-Set, where the Weights assigned are 30, 20, 20, and 10, the first Locator will get 37.5% of the traffic, the 2nd and 3rd Locators will each get 25% of the traffic, and the 4th Locator will get 12.5% of the traffic. If all Weights for a Locator-Set are equal, the receiver of the Map-Reply will decide how to load-

split the traffic. See RLOC-hashing [I-D.ietf-lisp-rfc6830bis] for a suggested hash algorithm to distribute the load across Locators with the same Priority and equal Weight values.

M Priority: Each RLOC is assigned a multicast Priority used by an ETR in a receiver multicast site to select an ITR in a source multicast site for building multicast distribution trees. A value of 255 means the RLOC MUST NOT be used for joining a multicast distribution tree. For more details, see [RFC6831].

M Weight: When priorities are the same for multiple RLOCs, the Weight indicates how to balance building multicast distribution trees across multiple ITRs. The Weight is encoded as a relative weight (similar to the unicast Weights) of the total number of trees built to the source site identified by the EID-Prefix. If all Weights for a Locator-Set are equal, the receiver of the Map-Reply will decide how to distribute multicast state across ITRs. For more details, see [RFC6831].

Unused Flags: These are set to 0 when sending and ignored on receipt.

L: When this bit is set, the Locator is flagged as a local Locator to the ETR that is sending the Map-Reply. When a Map-Server is doing proxy Map-Replying for a LISP site, the L-bit is set to 0 for all Locators in this Locator-Set.

p: When this bit is set, an ETR informs the RLOC-Probing ITR that the locator address for which this bit is set is the one being RLOC-probed and may be different from the source address of the Map-Reply. An ITR that RLOC-probes a particular Locator MUST use this Locator for retrieving the data structure used to store the fact that the Locator is reachable. The p-bit is set for a single Locator in the same Locator-Set. If an implementation sets more than one p-bit erroneously, the receiver of the Map-Reply MUST select the first set p-bit Locator. The p-bit MUST NOT be set for Locator-Set records sent in Map-Request and Map-Register messages.

R: This is set when the sender of a Map-Reply has a route to the Locator in the Locator data record. This receiver may find this useful to know if the Locator is up but not necessarily reachable from the receiver's point of view.

Locator: This is an IPv4 or IPv6 address (as encoded by the 'Loc-AFI' field) assigned to an ETR and used by an ITR as a destination RLOC address in the outer header of a LISP encapsulated packet. Note that the destination RLOC address of a LISP encapsulated packet MAY be an anycast address. A source RLOC of a LISP

encapsulated packet can be an anycast address as well. The source or destination RLOC MUST NOT be the broadcast address (255.255.255.255 or any subnet broadcast address known to the router) and MUST NOT be a link-local multicast address. The source RLOC MUST NOT be a multicast address. The destination RLOC SHOULD be a multicast address if it is being mapped from a multicast destination EID.

Map-Reply MUST be rate-limited, it is RECOMMENDED that a Map-Reply for the same destination RLOC be sent no more than one packets per 3 seconds.

The Record format, as defined here, is used both in the Map-Reply and Map-Register messages, this includes all the field definitions.

5.5. EID-to-RLOC UDP Map-Reply Message

A Map-Reply returns an EID-Prefix with a mask-length that is less than or equal to the EID being requested. The EID being requested is either from the destination field of an IP header of a Data-Probe or the EID of a record of a Map-Request. The RLOCs in the Map-Reply are routable IP addresses of all ETRs for the LISP site. Each RLOC conveys status reachability but does not convey path reachability from a requester's perspective. Separate testing of path reachability is required. See RLOC-reachability Section 7.1 for details.

Note that a Map-Reply MAY contain different EID-Prefix granularity (prefix + mask-length) than the Map-Request that triggers it. This might occur if a Map-Request were for a prefix that had been returned by an earlier Map-Reply. In such a case, the requester updates its cache with the new prefix information and granularity. For example, a requester with two cached EID-Prefixes that are covered by a Map-Reply containing one less-specific prefix replaces the entry with the less-specific EID-Prefix. Note that the reverse, replacement of one less-specific prefix with multiple more-specific prefixes, can also occur, not by removing the less-specific prefix but rather by adding the more-specific prefixes that, during a lookup, will override the less-specific prefix.

When an EID moves out of a LISP site [I-D.ietf-lisp-eid-mobility], the database mapping system may have overlapping EID-prefixes. Or when a LISP site is configured with multiple sets of ETRs that support different EID-prefix mask-lengths, the database mapping system may have overlapping EID-prefixes. When overlapping EID-prefixes exist, a Map-Request with an EID that best matches any EID-Prefix MUST be returned in a single Map-Reply message. For instance, if an ETR had database mapping entries for EID-Prefixes:

```
2001:db8::/32
2001:db8:1::/48
2001:db8:1:1::/64
2001:db8:1:2::/64
```

A Map-Request for EID 2001:db8:1:1::1 would cause a Map-Reply with a record count of 1 to be returned with a mapping record EID-Prefix of 2001:db8:1:1::/64.

A Map-Request for EID 2001:db8:1:5::5 would cause a Map-Reply with a record count of 3 to be returned with mapping records for EID-Prefixes 2001:db8:1::/48, 2001:db8:1:1::/64, 2001:db8:1:2::/64, filling out the /48 with more-specifics that exist in the mapping system.

Note that not all overlapping EID-Prefixes need to be returned but only the more-specific entries (note that in the second example above 2001:db8::/32 was not returned for requesting EID 2001:db8:1:5::5) for the matching EID-Prefix of the requesting EID. When more than one EID-Prefix is returned, all SHOULD use the same Time to Live value so they can all time out at the same time. When a more-specific EID-Prefix is received later, its Time to Live value in the Map-Reply record can be stored even when other less-specific entries exist. When a less-specific EID-Prefix is received later, its Map-Cache expiration time SHOULD be set to the minimum expiration time of any more-specific EID-Prefix in the Map-Cache. This is done so the integrity of the EID-Prefix set is wholly maintained and so no more-specific entries are removed from the Map-Cache while keeping less-specific entries.

For scalability, it is expected that aggregation of EID addresses into EID-Prefixes will allow one Map-Reply to satisfy a mapping for the EID addresses in the prefix range, thereby reducing the number of Map-Request messages.

Map-Reply records can have an empty Locator-Set. A Negative Map-Reply is a Map-Reply with an empty Locator-Set. Negative Map-Replies convey special actions by the sender to the ITR or PITR that have solicited the Map-Reply. There are two primary applications for Negative Map-Replies. The first is for a Map-Resolver to instruct an ITR or PITR when a destination is for a LISP site versus a non-LISP site, and the other is to source quench Map-Requests that are sent for non-allocated EIDs.

For each Map-Reply record, the list of Locators in a Locator-Set MUST be sorted in order of ascending IP address where an IPv4 locator address is considered numerically 'less than' an IPv6 locator address.

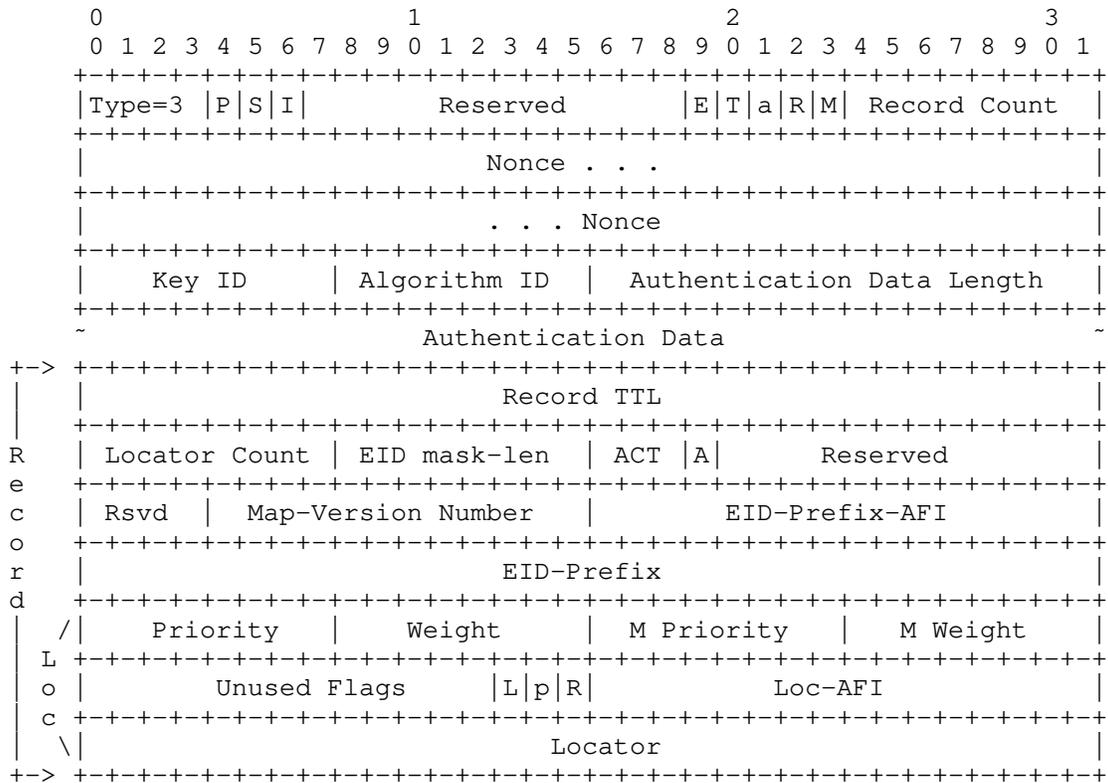
When sending a Map-Reply message, the destination address is copied from one of the 'ITR-RLLOC' fields from the Map-Request. The ETR can choose a locator address from one of the address families it supports. For Data-Probes, the destination address of the Map-Reply is copied from the source address of the Data-Probe message that is invoking the reply. The source address of the Map-Reply is one of the local IP addresses chosen, to allow Unicast Reverse Path Forwarding (uRPF) checks to succeed in the upstream service provider. The destination port of a Map-Reply message is copied from the source port of the Map-Request or Data-Probe, and the source port of the Map-Reply message is set to the well-known UDP port 4342.

5.6. Map-Register Message Format

This section specifies the encoding format for the Map-Register message. The message is sent in UDP with a destination UDP port of 4342 and a randomly selected UDP source port number.

The fields below are used in multiple control messages. They are defined for Map-Register, Map-Notify and Map-Notify-Ack message types.

The Map-Register message format is:



Packet field descriptions:

Type: 3 (Map-Register)

P: This is the proxy Map-Reply bit. When set to 1, the ETR sending the Map-Register message is requesting the Map-Server to proxy a Map-Reply. The Map-Server will send non-authoritative Map-Replies on behalf of the ETR.

- S: This is the security-capable bit. When set, the procedures from [I-D.ietf-lisp-sec] are supported.
- I: This is the ID-present bit. This bit is set to 1 to indicate that a 128 bit xTR-ID and a 64 bit Site-ID fields are present at the end of the Map-Register message. If an xTR is configured with an xTR-ID and Site-ID, it MUST set the I bit to 1 and include its xTR-ID and Site-ID in the Map-Register messages it generates. The combination of Site-ID plus xTR-ID uniquely identifies an xTR in a LISP domain and serves to track its last seen nonce.
- Reserved: This unassigned field MUST be set to 0 on transmit and MUST be ignored on receipt.
- E: This is the Map-Register EID-notify bit. This is used by a First-Hop-Router (FHR) which discovers a dynamic-EID. This EID-notify based Map-Register is sent by the FHR to a same site xTR that propagates the Map-Register to the mapping system. The site xTR keeps state to later Map-Notify the FHR after the EID has moves away. See [I-D.ietf-lisp-eid-mobility] for a detailed use-case.
- T: This is the use-TTL for timeout bit. When set to 1, the xTR wants the Map-Server to time out registrations based on the value in the "Record TTL" field of this message. Otherwise, the default timeout described in Section 8.2 is used.
- a: This is the merge-request bit. When set to 1, the xTR requests to merge RLOC-records from different xTRs registering the same EID-record. See signal-free multicast [RFC8378] for one use case example.
- R: This reserved and unassigned bit MUST be set to 0 on transmit and MUST be ignored on receipt.
- M: This is the want-map-notify bit. When set to 1, an ETR is requesting a Map-Notify message to be returned in response to sending a Map-Register message. The Map-Notify message sent by a Map-Server is used to acknowledge receipt of a Map-Register message.
- Record Count: This is the number of records in this Map-Register message. A record is comprised of that portion of the packet labeled 'Record' above and occurs the number of times equal to Record Count.
- Nonce: This 8-octet 'Nonce' field is incremented each time a Map-Register message is sent. When a Map-Register acknowledgement is requested, the nonce is returned by Map-Servers in Map-Notify

messages. Since the entire Map-Register message is authenticated, the 'Nonce' field serves to protect against Map-Register replay attacks. An ETR that registers to the mapping system SHOULD store the last nonce sent in persistent storage so when it restarts it can continue using an incrementing nonce. If the ETR cannot support saving the nonce, then when it restarts it MUST use a new authentication key to register to the mapping system. A Map-Server MUST track and save in persistent storage the last nonce received for each ETR xTR-ID and key pair. If a Map-Register is received with a nonce value that is not greater than the saved nonce, it MUST drop the Map-Register message and SHOULD log the fact a replay attack could have occurred.

Key ID: A key-id value that identifies a pre-shared secret between an ETR and a Map-Server. Per-message keys are derived from the pre-shared secret to authenticate the origin and protect the integrity of the Map-Register. The Key ID allows to rotate between multiple pre-shared secrets in a non disruptive way. The pre-shared secret MUST be unique per each LISP "Site-ID"

Algorithm ID: This field identifies the Key Derivation Function (KDF) and Message Authentication Code (MAC) algorithms used to derive the key and to compute the Authentication Data of a Map-Register. This 8-bit field identifies the KDF and MAC algorithm pair. See Section 12.5 for codepoint assignments.

Authentication Data Length: This is the length in octets of the 'Authentication Data' field that follows this field. The length of the 'Authentication Data' field is dependent on the MAC algorithm used. The length field allows a device that doesn't know the MAC algorithm to correctly parse the packet.

Authentication Data: This is the output of the MAC algorithm placed in this field after the MAC computation. The MAC output is computed as follows:

- 1: The KDF algorithm is identified by the field 'Algorithm ID' according to the table in Section 12.5. Implementations of this specification MUST implement HMAC-SHA-256-128 [RFC4868] and SHOULD implement HMAC-SHA-256-128+HKDF-SHA256 [RFC5869] .
- 2: The MAC algorithm is identified by the field 'Algorithm ID' according to the table in Section 12.5.

- 3: The pre-shared secret used to derive the per-message key is represented by PSK[Key ID], that is the pre-shared secret identified by the 'Key ID'.
- 4: The derived per-message key is computed as: $\text{per-msg-key} = \text{KDF}(\text{nonce} + \text{PSK}[\text{Key ID}], s)$. Where the nonce is the value in the Nonce field of the Map-Register, '+' denotes concatenation and 's' (the salt) is a string that corresponds to the message type being authenticated. For Map-Register messages, it is equal to "Map-Register Authentication". Similarly, for Map-Notify and Map-Notify-Ack messages, it is "Map-Notify Authentication" and "Map-Notify-Ack Authentication", respectively. For those Algorithm IDs defined in Section 12.5 that specify a 'none' KDF, the per-message key is computed as: $\text{per-msg-key} = \text{PSK}[\text{Key ID}]$. This means that the same key is used across multiple protocol messages.
- 5: The MAC output is computed using the MAC algorithm and the per-msg-key over the entire Map-Register payload (from and including the LISP message type field through the end of the last RLOC record) with the authenticated data field preset to 0.

The definition of the rest of the Map-Register can be found in EID-record description in Section 5.4. When the I-bit is set, the following fields are added to the end of the Map-Register message:

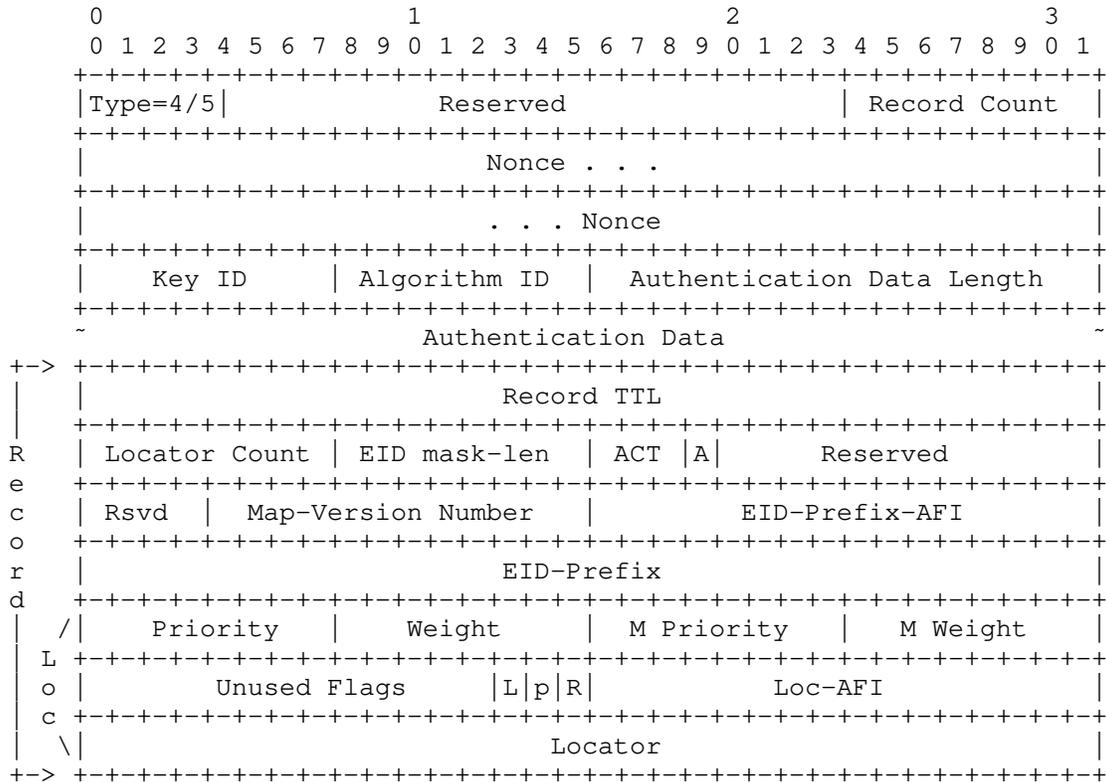
xTR-ID: xTR-ID is a 128 bit field at the end of the Map-Register message, starting after the final Record in the message. The xTR-ID is used to uniquely identify a xTR. The same xTR-ID value MUST NOT be used in two different xTRs in the scope of the Site-ID.

Site-ID: Site-ID is a 64 bit field at the end of the Map-Register message, following the xTR-ID. Site-ID is used to uniquely identify to which site the xTR that sent the message belongs. This document does not specify a strict meaning for the Site-ID field. Informally it provides an indication that a group of xTRs have some relation, either administratively, topologically or otherwise.

5.7. Map-Notify/Map-Notify-Ack Message Format

This section specifies the encoding format for the Map-Notify and Map-Notify-Ack messages. The messages are sent inside a UDP packet with source and destination UDP ports equal to 4342.

The Map-Notify and Map-Notify-Ack message formats are:



Packet field descriptions:

Type: 4/5 (Map-Notify/Map-Notify-Ack)

The Map-Notify message has the same contents as a Map-Register message. See the Map-Register section for field descriptions and the Map-Reply section for EID-record and RLOC-record descriptions.

The fields of the Map-Notify are copied from the corresponding Map-Register to acknowledge its correct processing. In the Map-Notify, the 'Authentication Data' field is recomputed using the corresponding per-message key and according to the procedure defined in the

previous section. The Map-Notify message can also be used, outside the scope of this specification, in an unsolicited manner, such as is specified in [I-D.ietf-lisp-pubsub].

After sending a Map-Register, if a Map-Notify is not received after 1 second the transmitter MUST re-transmit the original Map-Register with an exponential backoff (base of 2, that is, the next backoff timeout interval is doubled), the maximum backoff is 1 minute. Map-Notify messages are only transmitted upon the reception of a Map-Register with the M-bit set, Map-Notify messages are not retransmitted. The only exception to this is for unsolicited Map-Notify messages, see below.

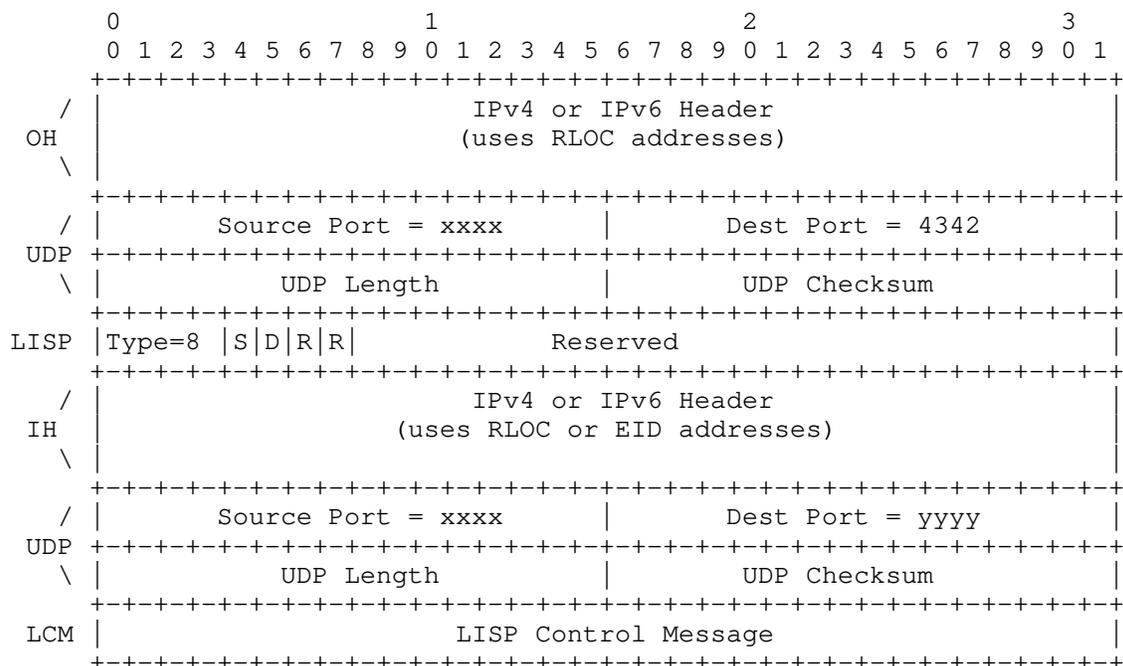
A Map-Server sends an unsolicited Map-Notify message (one that is not used as an acknowledgment to a Map-Register message) only in conformance with the Congestion Control And Reliability Guideline sections of [RFC8085]. A Map-Notify is retransmitted until a Map-Notify-Ack is received by the Map-Server with the same nonce used in the Map-Notify message. An implementation SHOULD retransmit up to 3 times at 3 second retransmission intervals, after which time the retransmission interval is exponentially backed-off (base of 2, that is, the next backoff timeout interval is doubled) for another 3 retransmission attempts. Map-Notify-Ack messages are only transmitted upon the reception of an unsolicited Map-Notify, Map-Notify-Ack messages are not retransmitted.

The Map-Notify-Ack message has the same contents as a Map-Notify message. It is used to acknowledge the receipt of an unsolicited Map-Notify and, once the authentication data is validated, allows for the sender to stop retransmitting a Map-Notify with the same nonce and the authentication data validates. The fields of the Map-Notify-Ack are copied from the corresponding Map-Notify message to acknowledge its correct processing. The 'Authentication Data' field is recomputed using the corresponding per-message key and according to the procedure defined in the previous section.

Upon reception of Map-Register, Map-Notify or Map-Notify-Ack, the receiver verifies the authentication data. If the authentication data fails to validate, the message is dropped without further processing.

5.8. Encapsulated Control Message Format

An Encapsulated Control Message (ECM) is used to encapsulate control packets sent between xTRs and the mapping database system or internal to the mapping database system.



Packet header descriptions:

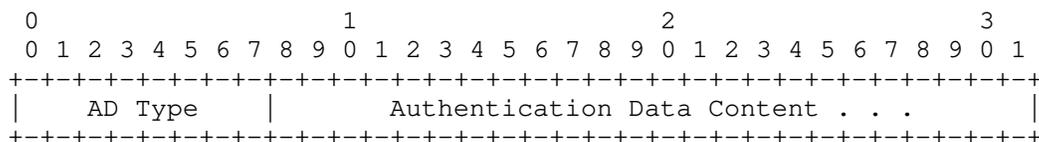
OH: The outer IPv4 or IPv6 header, which uses RLOC addresses in the source and destination header address fields.

UDP: The outer UDP header with destination port 4342. The source port is randomly allocated. The checksum field MUST be non-zero.

LISP: Type 8 is defined to be a "LISP Encapsulated Control Message", and what follows is either an IPv4 or IPv6 header as encoded by the first 4 bits after the 'Reserved' field, or the Authentication Data field [I-D.ietf-lisp-sec] if the S-bit (see below) is set.

Type: 8 (Encapsulated Control Message (ECM))

S: This is the Security bit. When set to 1, the field following the 'Reserved' field will have the following Authentication Data format and follow the procedures from [I-D.ietf-lisp-sec].



D: This is the DDT-bit. When set to 1, the sender is requesting a Map-Referral message to be returned. The details of this procedure are described in [RFC8111].

R: This reserved and unassigned bit MUST be set to 0 on transmit and MUST be ignored on receipt.

IH: The inner IPv4 or IPv6 header, which can use either RLOC or EID addresses in the header address fields. When a Map-Request is encapsulated in this packet format, the destination address in this header is an EID.

UDP: The inner UDP header, where the port assignments depend on the control packet being encapsulated. When the control packet is a Map-Request or Map-Register, the source port is selected by the ITR/PITR and the destination port is 4342. When the control packet is a Map-Reply, the source port is 4342 and the destination port is assigned from the source port of the invoking Map-Request. Port number 4341 MUST NOT be assigned to either port. The checksum field MUST be non-zero.

LCM: The format is one of the control message formats described in Section 5. Map-Request messages are allowed to be Control-Plane (ECM) encapsulated. When Map-Requests are sent for RLOC-Probing purposes (i.e. the probe-bit is set), they MUST NOT be sent inside Encapsulated Control Messages. PIM Join/Prune messages [RFC6831] are also allowed to be Control-Plane (ECM) encapsulated.

6. Changing the Contents of EID-to-RLOC Mappings

In the LISP architecture ITRs/PITRs use a local Map-Cache to store EID-to-RLOC mappings for forwarding. When an ETR updates a mapping a mechanism is required to inform ITRs/PITRs that are using such mappings.

The LISP Data-Plane defines several mechanism to update mappings [I-D.ietf-lisp-rfc6830bis]. This document specifies the Solicit-Map-Request (SMR), a Control-Plane push-based mechanism. An additional Control-Plane mechanism based on the Publish/subscribe paradigm is specified in [I-D.ietf-lisp-pubsub].

6.1. Solicit-Map-Request (SMR)

Soliciting a Map-Request is a selective way for ETRs, at the site where mappings change, to control the rate they receive requests for Map-Reply messages. SMRs are also used to tell remote ITRs to update the mappings they have cached.

Since ETRs are not required to keep track of remote ITRs that have cached their mappings, they do not know which ITRs need to have their mappings updated. As a result, an ETR will solicit Map-Requests to those sites to which it has been sending LISP encapsulated data packets for the last minute. As a result, when an ETR is also acting as ITR, it will send an SMR to an ITR to which it has recently sent encapsulated data.

An SMR message is simply a bit set in a Map-Request message. An ITR or PITR will send a Map-Request (SMR-invoked Map-Request) when they receive an SMR message. While the SMR message is sent through the data-plane, the SMR-invoked Map-Request MUST be sent through the Mapping System (not directly).

Both the SMR sender and the SMR responder MUST rate-limit these messages. It is RECOMMENDED that the SMR sender rate-limits Map-Request for the same destination RLOC to no more than one packet per 3 seconds. It is RECOMMENDED that the SMR responder rate-limits Map-Request for the same EID-Prefix to no more than once per 3 seconds.

When an ITR receives an SMR message for which it does not have a cached mapping for the EID in the SMR message, it SHOULD NOT send an SMR-invoked Map-Request. This scenario can occur when an ETR sends SMR messages to all Locators in the Locator-Set it has stored in its Map-Cache but the remote ITRs that receive the SMR may not be sending packets to the site. There is no point in updating the ITRs until they need to send, in which case they will send Map-Requests to obtain a Map-Cache entry.

7. Routing Locator Reachability

This document defines several Control-Plane mechanisms for determining RLOC reachability. Please note that additional Data-Plane reachability mechanisms are defined in [I-D.ietf-lisp-rfc6830bis].

1. An ITR may receive an ICMP Network Unreachable or Host Unreachable message for an RLOC it is using. This indicates that the RLOC is likely down. Note that trusting ICMP messages may not be desirable, but neither is ignoring them completely. Implementations are encouraged to follow current best practices in treating these conditions [I-D.ietf-opsec-icmp-filtering].
2. When an ITR participates in the routing protocol that operates in the underlay routing system, it can determine that an RLOC is down when no Routing Information Base (RIB) entry exists that matches the RLOC IP address.
3. An ITR may receive an ICMP Port Unreachable message from a destination host. This occurs if an ITR attempts to use interworking [RFC6832] and LISP-encapsulated data is sent to a non-LISP-capable site.
4. An ITR may receive a Map-Reply from an ETR in response to a previously sent Map-Request. The RLOC source of the Map-Reply is likely up, since the ETR was able to send the Map-Reply to the ITR. Please note that in some scenarios the RLOC -from the outer header- can be a spoofable field.
5. An ITR/ETR pair can use the 'RLOC-Probing' mechanism described below.

When ITRs receive ICMP Network Unreachable or Host Unreachable messages as a method to determine unreachability, they will refrain from using Locators that are described in Locator lists of Map-Replies. However, using this approach is unreliable because many network operators turn off generation of ICMP Destination Unreachable messages.

If an ITR does receive an ICMP Network Unreachable or Host Unreachable message, it MAY originate its own ICMP Destination Unreachable message destined for the host that originated the data packet the ITR encapsulated.

This assumption does create a dependency: Locator unreachability is detected by the receipt of ICMP Host Unreachable messages. When a Locator has been determined to be unreachable, it is not used for

active traffic; this is the same as if it were listed in a Map-Reply with Priority 255.

The ITR can test the reachability of the unreachable Locator by sending periodic Map-Requests. Both Map-Requests and Map-Replies MUST be rate-limited, see Section 5.3 and Section 5.4 for information about rate-limiting. Locator reachability testing is never done with data packets, since that increases the risk of packet loss for end-to-end sessions.

7.1. RLOC-Probing Algorithm

RLOC-Probing is a method that an ITR or PITR can use to determine the reachability status of one or more Locators that it has cached in a Map-Cache entry. The probe-bit of the Map-Request and Map-Reply messages is used for RLOC-Probing.

RLOC-Probing is done in the control plane on a timer basis, where an ITR or PITR will originate a Map-Request destined to a locator address from one of its own locator addresses. A Map-Request used as an RLOC-probe is NOT encapsulated and NOT sent to a Map-Server or to the mapping database system as one would when requesting mapping data. The EID record encoded in the Map-Request is the EID-Prefix of the Map-Cache entry cached by the ITR or PITR. The ITR MAY include a mapping data record for its own database mapping information that contains the local EID-Prefixes and RLOCs for its site. RLOC-probes are sent periodically using a jittered timer interval.

When an ETR receives a Map-Request message with the probe-bit set, it returns a Map-Reply with the probe-bit set. The source address of the Map-Reply is set to the IP address of the outgoing interface the Map-Reply destination address routes to. The Map-Reply SHOULD contain mapping data for the EID-Prefix contained in the Map-Request. This provides the opportunity for the ITR or PITR that sent the RLOC-probe to get mapping updates if there were changes to the ETR's database mapping entries.

There are advantages and disadvantages of RLOC-Probing. The main benefit of RLOC-Probing is that it can handle many failure scenarios allowing the ITR to determine when the path to a specific Locator is reachable or has become unreachable, thus providing a robust mechanism for switching to using another Locator from the cached Locator. RLOC-Probing can also provide rough Round-Trip Time (RTT) estimates between a pair of Locators, which can be useful for network management purposes as well as for selecting low delay paths. The major disadvantage of RLOC-Probing is in the number of control messages required and the amount of bandwidth used to obtain those

benefits, especially if the requirement for failure detection times is very small.

8. Interactions with Other LISP Components

8.1. ITR EID-to-RLOC Mapping Resolution

An ITR is configured with one or more Map-Resolver addresses. These addresses are "Locators" (or RLOCs) and MUST be routable on the underlying core network; they MUST NOT need to be resolved through LISP EID-to-RLOC mapping, as that would introduce a circular dependency. When using a Map-Resolver, an ITR does not need to connect to any other database mapping system.

An ITR sends an Encapsulated Map-Request to a configured Map-Resolver when it needs an EID-to-RLOC mapping that is not found in its local Map-Cache. Using the Map-Resolver greatly reduces both the complexity of the ITR implementation and the costs associated with its operation.

In response to an Encapsulated Map-Request, the ITR can expect one of the following:

- o An immediate Negative Map-Reply (with action code of "Natively-Forward", 15-minute Time to Live (TTL)) from the Map-Resolver if the Map-Resolver can determine that the requested EID does not exist. The ITR saves the EID-Prefix returned in the Map-Reply in its cache, marks it as non-LISP-capable, and knows not to attempt LISP encapsulation for destinations matching it.
- o A Negative Map-Reply, with action code of "Natively-Forward", from a Map-Server that is authoritative (within the LISP deployment Section 1.1) for an EID-Prefix that matches the requested EID but that does not have an actively registered, more-specific EID-prefix. In this case, the requested EID is said to match a "hole" in the authoritative EID-Prefix. If the requested EID matches a more-specific EID-Prefix that has been delegated by the Map-Server but for which no ETRs are currently registered, a 1-minute TTL is returned. If the requested EID matches a non-delegated part of the authoritative EID-Prefix, then it is not a LISP EID and a 15-minute TTL is returned. See Section 8.2 for discussion of aggregate EID-Prefixes and details of Map-Server EID-Prefix matching.
- o A LISP Map-Reply from the ETR that owns the EID-to-RLOC mapping or possibly from a Map-Server answering on behalf of the ETR. See Section 8.4 for more details on Map-Resolver message processing.

Note that an ITR may be configured to both use a Map-Resolver and to participate in a LISP-ALT logical network. In such a situation, the ITR SHOULD send Map-Requests through the ALT network for any EID-Prefix learned via ALT BGP. Such a configuration is expected to be very rare, since there is little benefit to using a Map-Resolver if an ITR is already using LISP-ALT. There would be, for example, no need for such an ITR to send a Map-Request to a possibly non-existent EID (and rely on Negative Map-Replies) if it can consult the ALT database to verify that an EID-Prefix is present before sending that Map-Request.

8.2. EID-Prefix Configuration and ETR Registration

An ETR publishes its EID-Prefixes on a Map-Server by sending LISP Map-Register messages. A Map-Register message includes authentication data, so prior to sending a Map-Register message, the ETR and Map-Server MUST be configured with a pre-shared secret used to derive Map-Register authentication keys. A Map-Server's configuration SHOULD also include a list of the EID-Prefixes for which each ETR is authoritative. Upon receipt of a Map-Register from an ETR, a Map-Server accepts only EID-Prefixes that are configured for that ETR. Failure to implement such a check would leave the mapping system vulnerable to trivial EID-Prefix hijacking attacks.

In addition to the set of EID-Prefixes defined for each ETR that may register, a Map-Server is typically also configured with one or more aggregate prefixes that define the part of the EID numbering space assigned to it. When LISP-ALT is the database in use, aggregate EID-Prefixes are implemented as discard routes and advertised into ALT BGP. The existence of aggregate EID-Prefixes in a Map-Server's database means that it may receive Map Requests for EID-Prefixes that match an aggregate but do not match a registered prefix; Section 8.3 describes how this is handled.

Map-Register messages are sent periodically from an ETR to a Map-Server with a suggested interval between messages of one minute. A Map-Server SHOULD time out and remove an ETR's registration if it has not received a valid Map-Register message within the past three minutes. When first contacting a Map-Server after restart or changes to its EID-to-RLOC database mappings, an ETR MAY initially send Map-Register messages at an increased frequency, up to one every 20 seconds. This "quick registration" period is limited to five minutes in duration.

An ETR MAY request that a Map-Server explicitly acknowledge receipt and processing of a Map-Register message by setting the "want-map-notify" (M-bit) flag. A Map-Server that receives a Map-Register with this flag set will respond with a Map-Notify message. Typical use of

this flag by an ETR would be to set it for Map-Register messages sent during the initial "quick registration" with a Map-Server but then set it only occasionally during steady-state maintenance of its association with that Map-Server. Note that the Map-Notify message is sent to UDP destination port 4342, not to the source port specified in the original Map-Register message.

Note that a one-minute minimum registration interval during maintenance of an ETR-Map-Server association places a lower bound on how quickly and how frequently a mapping database entry can be updated. This may have implications for what sorts of mobility can be supported directly by the mapping system; shorter registration intervals or other mechanisms might be needed to support faster mobility in some cases. For a discussion on one way that faster mobility may be implemented for individual devices, please see [I-D.ietf-lisp-mn].

An ETR MAY also request, by setting the "proxy Map-Reply" flag (P-bit) in the Map-Register message, that a Map-Server answer Map-Requests instead of forwarding them to the ETR. See Section 7.1 for details on how the Map-Server sets certain flags (such as those indicating whether the message is authoritative and how returned Locators SHOULD be treated) when sending a Map-Reply on behalf of an ETR. When an ETR requests proxy reply service, it SHOULD include all RLOCs for all ETRs for the EID-Prefix being registered, along with the routable flag ("R-bit") setting for each RLOC. The Map-Server includes all of this information in Map-Reply messages that it sends on behalf of the ETR. This differs from a non-proxy registration, since the latter need only provide one or more RLOCs for a Map-Server to use for forwarding Map-Requests; the registration information is not used in Map-Replies, so it being incomplete is not incorrect.

An ETR that uses a Map-Server to publish its EID-to-RLOC mappings does not need to participate further in the mapping database protocol(s). When using a LISP-ALT mapping database, for example, this means that the ETR does not need to implement GRE or BGP, which greatly simplifies its configuration and reduces its cost of operation.

Note that use of a Map-Server does not preclude an ETR from also connecting to the mapping database (i.e., it could also connect to the LISP-ALT network), but doing so doesn't seem particularly useful, as the whole purpose of using a Map-Server is to avoid the complexity of the mapping database protocols.

8.3. Map-Server Processing

Once a Map-Server has EID-Prefixes registered by its client ETRs, it can accept and process Map-Requests for them.

In response to a Map-Request, the Map-Server first checks to see if the destination EID matches a configured EID-Prefix. If there is no match, the Map-Server returns a Negative Map-Reply with action code "Natively-Forward" and a 15-minute TTL. This can occur if a Map Request is received for a configured aggregate EID-Prefix for which no more-specific EID-Prefix exists; it indicates the presence of a non-LISP "hole" in the aggregate EID-Prefix.

Next, the Map-Server checks to see if any ETRs have registered the matching EID-Prefix. If none are found, then the Map-Server returns a Negative Map-Reply with action code "Natively-Forward" and a 1-minute TTL.

If the EID-prefix is either registered or not registered to the mapping system and there is a policy in the Map-Server to have the requestor drop packets for the matching EID-prefix, then a Drop/Policy-Denied action is returned. If the EID-prefix is registered or not registered and there is an authentication failure, then a Drop/Authentication-failure action is returned. If either of these actions result as a temporary state in policy or authentication then a Send-Map-Request action with 1-minute TTL MAY be returned to allow the requestor to retry the Map-Request.

If any of the registered ETRs for the EID-Prefix have requested proxy reply service, then the Map-Server answers the request instead of forwarding it. It returns a Map-Reply with the EID-Prefix, RLOCs, and other information learned through the registration process.

If none of the ETRs have requested proxy reply service, then the Map-Server re-encapsulates and forwards the resulting Encapsulated Map-Request to one of the registered ETRs. It does not otherwise alter the Map-Request, so any Map-Reply sent by the ETR is returned to the RLOC in the Map-Request, not to the Map-Server. Unless also acting as a Map-Resolver, a Map-Server should never receive Map-Replies; any such messages SHOULD be discarded without response, perhaps accompanied by the logging of a diagnostic message if the rate of Map-Replies is suggestive of malicious traffic.

8.4. Map-Resolver Processing

Upon receipt of an Encapsulated Map-Request, a Map-Resolver decapsulates the enclosed message and then searches for the requested EID in its local database of mapping entries (statically configured

or learned from associated ETRs if the Map-Resolver is also a Map-Server offering proxy reply service). If it finds a matching entry, it returns a LISP Map-Reply with the known mapping.

If the Map-Resolver does not have the mapping entry and if it can determine that the EID is not in the mapping database (for example, if LISP-ALT is used, the Map-Resolver will have an ALT forwarding table that covers the full EID space), it immediately returns a negative LISP Map-Reply, with action code "Natively-Forward" and a 15-minute TTL. To minimize the number of negative cache entries needed by an ITR, the Map-Resolver SHOULD return the least-specific prefix that both matches the original query and does not match any EID-Prefix known to exist in the LISP-capable infrastructure.

If the Map-Resolver does not have sufficient information to know whether the EID exists, it needs to forward the Map-Request to another device that has more information about the EID being requested. To do this, it forwards the unencapsulated Map-Request, with the original ITR RLOC as the source, to the mapping database system. Using LISP-ALT, the Map-Resolver is connected to the ALT network and sends the Map-Request to the next ALT hop learned from its ALT BGP neighbors. The Map-Resolver does not send any response to the ITR; since the source RLOC is that of the ITR, the ETR or Map-Server that receives the Map-Request over the ALT and responds will do so directly to the ITR.

8.4.1. Anycast Operation

A Map-Resolver can be set up to use "anycast", where the same address is assigned to multiple Map-Resolvers and is propagated through IGP routing, to facilitate the use of a topologically close Map-Resolver by each ITR.

ETRs MAY have anycast RLOC addresses which are registered as part of their RLOC-set to the mapping system. However, registrations MUST use their unique RLOC addresses, distinct authentication keys or different XTR-IDs to identify security associations with the Map-Servers.

9. Security Considerations

A LISP threat analysis can be found in [RFC7835]. In what follows we highlight security considerations that apply when LISP is deployed in environments such as those specified in Section 1.1, where the following assumptions hold:

1. The Mapping System is secure and trusted, and for the purpose of this security considerations the Mapping System is considered as one trusted element.
2. The ETRs have a pre-configured trust relationship with the Mapping System, which includes some form of shared secret, and the Mapping System is aware of which EIDs an ETR can advertise. How those keys and mappings gets established is out of the scope of this document.
3. LISP-SEC [I-D.ietf-lisp-sec] MUST be implemented. Network operators should carefully weight how the LISP-SEC threat model applies to their particular use case or deployment. If they decide to ignore a particular recommendation, they should make sure the risk associated with the corresponding threats is well understood.

The Map-Request/Map-Reply message exchange can be exploited by an attacker to mount DoS and/or amplification attacks. Attackers can send Map-Requests at high rates to overload LISP nodes and increase the state maintained by such nodes or consume CPU cycles. Such threats can be mitigated by systematically applying filters and rate limiters.

The Map-Request/Map-Reply message exchange can also be exploited to inject forged mappings directly in the ITR EID-to-RLOC map-cache. This can lead to traffic being redirected to the attacker, see further details in [RFC7835]. In addition, valid ETRs in the system can perform overclaiming attacks. In this case, attackers can claim to own an EID-prefix that is larger than the prefix owned by the ETR. Such attacks can be addressed by using LISP-SEC [I-D.ietf-lisp-sec]. The LISP-SEC protocol defines a mechanism for providing origin authentication, integrity protection, and prevention of 'man-in-the-middle' and 'prefix overclaiming' attacks on the Map-Request/Map-Reply exchange. In addition and while beyond the scope of securing an individual Map-Server or Map-Resolver, it should be noted that LISP-SEC can be complemented by additional security mechanisms defined by the Mapping System Infrastructure. For instance, BGP-based LISP-ALT [RFC6836] can take advantage of standards work on adding security to BGP while LISP-DDT [RFC8111] defines its own additional security mechanisms.

To publish an authoritative EID-to-RLOC mapping with a Map-Server using the Map-Register message, an ETR includes authentication data that is a MAC of the entire message using a key derived from the pre-shared secret. An implementation SHOULD support HMAC-SHA256-128+HKDF-SHA256 [RFC5869]. The Map-Register message includes protection for replay attacks by a man-in-the-middle. However, there

is a potential attack where a compromised ETR could overclaim the prefix it owns and successfully register it on its corresponding Map-Server. To mitigate this and as noted in Section 8.2, a Map-Server MUST verify that all EID-Prefixes registered by an ETR match the configuration stored on the Map-Server.

Deployments concerned about manipulations of Map-Request and Map-Reply messages, and malicious ETR EID prefix overclaiming MUST drop LISP Control Plane messages that do not contain LISP-SEC material (S-bit, EID-AD, OTK-AD, PKT-AD).

Mechanisms to encrypt, support privacy, prevent eavesdropping and packet tampering for messages exchanged between xTRs, xTRs and the mapping system, and nodes that make up the mapping system, SHOULD be deployed. Examples of this are DTLS [RFC6347] or LISP-crypto [RFC8061].

10. Privacy Considerations

As noted by [RFC6973] privacy is a complex issue that greatly depends on the specific protocol use-case and deployment. As noted in section 1.1 of [I-D.ietf-lisp-rfc6830bis] LISP focuses on use-cases where entities communicate over the public Internet while keeping separate addressing and topology. In what follows we detail the privacy threats introduced by the LISP Control Plane, the analysis is based on the guidelines detailed in [RFC6973].

LISP can use long-lived identifiers (EIDs) that survive mobility events. Such identifiers bind to the RLOCs of the nodes, which represents the topological location with respect to the specific LISP deployments. In addition, EID-to-RLOC mappings are typically considered public information within the LISP deployment when control-plane messages are not encrypted, and can be eavesdropped while Map-Request messages are sent to the corresponding Map-Resolvers or Map-Register messages to Map-Servers.

In this context, attackers can correlate the EID with the RLOC and track the corresponding user topological location and/or mobility. This can be achieved by off-path attackers, if they are authenticated, by querying the mapping system. Deployments concerned about this threat can use access control-lists or stronger authentication mechanisms [I-D.ietf-lisp-ecdsa-auth] in the mapping system to make sure that only authorized users can access this information (data minimization). Use of ephemeral EIDs [I-D.ietf-lisp-eid-anonymity] to achieve anonymity is another mechanism to lessen persistency and identity tracking.

11. Changes since RFC 6833

For implementation considerations, the following major changes have been made to this document since RFC 6833 was published:

- o A Map-Notify-Ack message is added in this document to provide reliability for Map-Notify messages. Any receiver of a Map-Notify message must respond with a Map-Notify-Ack message. Map-Servers who are senders of Map-Notify messages, must queue the Map-Notify contents until they receive a Map-Notify-Ack with the nonce used in the Map-Notify message. Note that implementations for Map-Notify-Ack support already exist and predate this document.
- o This document is incorporating the codepoint for the Map-Referral message from the LISP-DDT specification [RFC8111] to indicate that a Map-Server must send the final Map-Referral message when it participates in the LISP-DDT mapping system procedures.
- o The "L" and "D" bits are added to the Map-Request message. See Section 5.3 for details.
- o The "S", "I", "E", "T", "a", "R", and "M" bits are added to the Map-Register message. See Section 5.6 for details.
- o The 16-bit Key-ID field of the Map-Register message has been split into a 8-bit Key-ID field and a 8-bit Algorithm-ID field.
- o The nonce and the authentication data in the Map-Register message have a different behaviour, see Section 5.6 for details.
- o This document adds two new Action values that are in an EID-record that appear in Map-Reply, Map-Register, Map-Notify, and Map-Notify-Ack messages. The Drop/Policy-Denied and Drop/Auth-Failure are the descriptions for the two new action values. See Section 5.4 for details.

12. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to this LISP Control-Plane specification, in accordance with BCP 26 [RFC8126].

There are three namespaces (listed in the sub-sections below) in LISP that have been registered.

- o LISP IANA registry allocations should not be made for purposes unrelated to LISP routing or transport protocols.

- o The following policies are used here with the meanings defined in BCP 26: "Specification Required", "IETF Review", "Experimental Use", and "First Come First Served".

12.1. LISP UDP Port Numbers

The IANA registry has allocated UDP port number 4342 for the LISP Control-Plane. IANA has updated the description for UDP port 4342 as follows:

Keyword	Port	Transport Layer	Description
-----	----	-----	-----
lisp-control	4342	udp	LISP Control Packets

12.2. LISP Packet Type Codes

It is being requested that the IANA be authoritative for LISP Packet Type definitions and it is requested to replace the [RFC6830] registry message references with the RFC number assigned to this document.

Based on deployment experience of [RFC6830], the Map-Notify-Ack message, message type 5, was added by this document. This document requests IANA to add it to the LISP Packet Type Registry.

Name	Number	Defined in
----	-----	-----
LISP Map-Notify-Ack	5	RFC6833bis

12.3. LISP Map-Reply EID-Record Action Codes

New ACT values can be allocated through IETF review or IESG approval. Four values have already been allocated by [RFC6830]. IANA is requested to replace the [RFC6830] reference for this registry with the RFC number assigned to this document and [RFC6830]. This specification changes the name of ACT type 3 value from "Drop" to "Drop/No-Reason" as well as adding two new ACT values, the "Drop/Policy-Denied" (type 4) and "Drop/Authentication-Failure" (type 5).

Value	Action	Description	Reference
4	Drop/Policy-Denied	A packet matching this Map-Cache entry is dropped because the target EWID is policy-denied by the xTR or the mapping system.	RFC6833bis
5	Drop/Auth-Failure	Packet matching the Map-Cache entry is dropped because the Map-Request for the target EID fails an authentication check by the xTR or the mapping system.	RFC6833bis

LISP Map-Reply Action Values

In addition, LISP has a number of flag fields and reserved fields, such as the LISP header flags field [I-D.ietf-lisp-rfc6830bis]. New bits for flags in these fields can be implemented after IETF review or IESG approval, but these need not be managed by IANA.

12.4. LISP Address Type Codes

LISP Canonical Address Format (LCAF) [RFC8060] is an 8-bit field that defines LISP-specific encodings for AFI value 16387. LCAF encodings are used for specific use-cases where different address types for EID-records and RLOC-records are required.

The IANA registry "LISP Canonical Address Format (LCAF) Types" is used for LCAF types. The registry for LCAF types use the Specification Required policy [RFC8126]. Initial values for the registry as well as further information can be found in [RFC8060].

Therefore, there is no longer a need for the "LISP Address Type Codes" registry requested by [RFC6830]. This document requests to remove it.

12.5. LISP Algorithm ID Numbers

In [RFC6830], a request for a "LISP Key ID Numbers" registry was submitted. This document renames the registry to "LISP Algorithm ID Numbers" and requests the IANA to make the name change.

The following Algorithm ID values are defined by this specification as used in any packet type that references a 'Algorithm ID' field:

Name	Number	MAC	KDF
None	0	None	None
HMAC-SHA-1-96-None	1	[RFC2404]	None
HMAC-SHA-256-128-None	2		[RFC4868]
HMAC-SHA256-128+HKDF-SHA256	3	[RFC4868]	[RFC4868]

Number values are in the range of 0 to 255. The allocation of values is on a first come first served basis.

12.6. LISP Bit Flags

This document asks IANA to create a registry for allocation of bits in several headers of the LISP control plane, namely in the Map-Request, Map-Reply, Map-Register, Encapsulated Control Message (ECM) messages. Bit allocations are also requested for EID-records and RLOC-records. The registry created should be named "LISP Control Plane Header Bits". A sub-registry needs to be created per each message and EID-record. The name of each sub-registry is indicated below, along with its format and allocation of bits defined in this document. Any additional bits allocation, requires a specification, according with [RFC8126] policies.

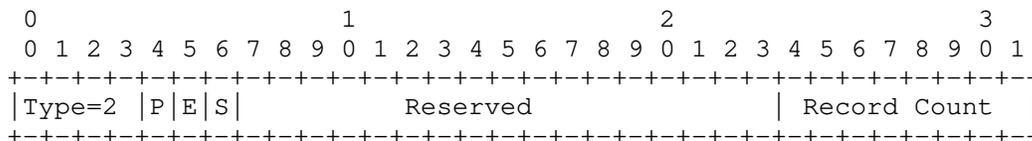
Sub-Registry: Map-Request Header Bits [Section 5.2]:

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	2	3	4	5	6	7	8	9																				
Type=1										A M P S p s R R										Rsvd										L D										IRC										Record Count									

Spec Name	IANA Name	Bit Position	Description
A	map-request-A	4	Authoritative Bit
M	map-request-M	5	Map Data Present Bit
P	map-request-P	6	RLOC-Probe Request Bit
S	map-request-S	7	Solicit Map-Request (SMR) Bit
p	map-request-p	8	Proxy-ITR Bit
s	map-request-s	9	Solicit Map-Request Invoked Bit
L	map-request-L	17	Local xTR Bit
D	map-request-D	18	Don't Map-Reply Bit

LISP Map-Request Header Bits

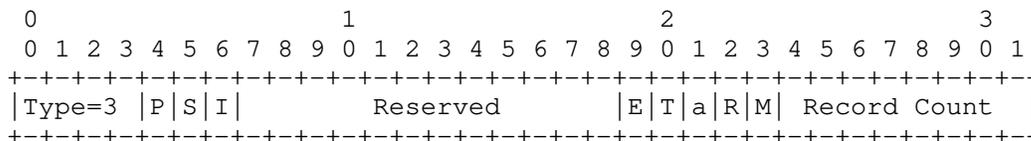
Sub-Registry: Map-Reply Header Bits [Section 5.4]:



Spec Name	IANA Name	Bit Position	Description
P	map-reply-P	4	RLOC-Probe Bit
E	map-reply-E	5	Echo Nonce Capable Bit
S	map-reply-S	6	Security Bit

LISP Map-Reply Header Bits

Sub-Registry: Map-Register Header Bits [Section 5.6]:



Spec Name	IANA Name	Bit Position	Description
P	map-register-P	4	Proxy Map-Reply Bit
S	map-register-S	5	LISP-SEC Capable Bit
I	map-register-I	6	xTR-ID present flag

LISP Map-Register Header Bits

Sub-Registry: Encapsulated Control Message (ECM) Header Bits [Section 5.8]:

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
Type=8 S D E M		Reserved	

Spec Name	IANA Name	Bit Position	Description
S	ecm-S	4	Security Bit
D	ecm-D	5	LISP-DDT Bit
E	ecm-E	6	Forward to ETR Bit
M	ecm-M	7	Destined to Map-Server Bit

LISP Encapsulated Control Message (ECM) Header Bits

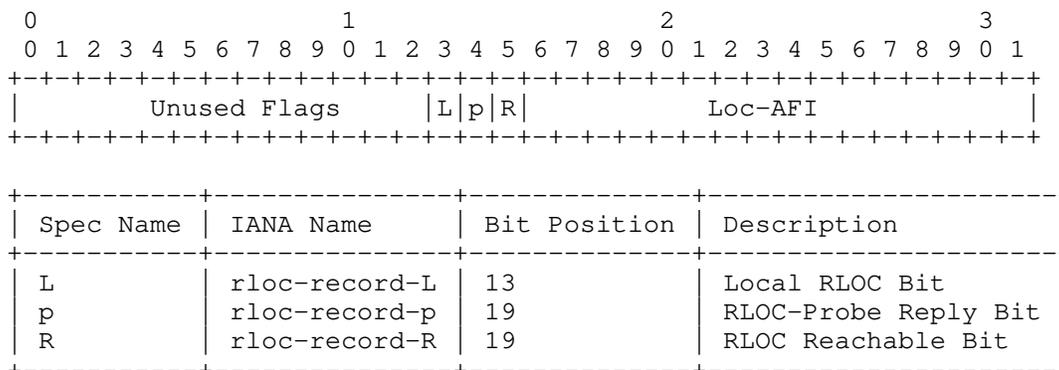
Sub-Registry: EID-Record Header Bits [Section 5.4]:

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
Locator Count		EID mask-len ACT A	
		Reserved	

Spec Name	IANA Name	Bit Position	Description
A	eid-record-A	19	Authoritative Bit

LISP EID-Record Header Bits

Sub-Registry: RLOC-Record Header Bits [Section 5.4]:



LISP RLOC-Record Header Bits

13. References

13.1. Normative References

[I-D.ietf-lisp-6834bis]
 Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", draft-ietf-lisp-6834bis-09 (work in progress), August 2021.

[I-D.ietf-lisp-rfc6830bis]
 Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, "The Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-rfc6830bis-36 (work in progress), November 2020.

[I-D.ietf-lisp-rfc8113bis]
 Boucadair, M. and C. Jacquenet, "Locator/ID Separation Protocol (LISP): Shared Extension Message & IANA Registry for Packet Type Allocations", draft-ietf-lisp-rfc8113bis-03 (work in progress), January 2019.

[I-D.ietf-lisp-sec]
 Maino, F., Ermagan, V., Cabellos, A., and D. Saucez, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-25 (work in progress), December 2021.

[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>>.

- [RFC2404] Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", RFC 2404, DOI 10.17487/RFC2404, November 1998, <<https://www.rfc-editor.org/info/rfc2404>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", RFC 4868, DOI 10.17487/RFC4868, May 2007, <<https://www.rfc-editor.org/info/rfc4868>>.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, DOI 10.17487/RFC5869, May 2010, <<https://www.rfc-editor.org/info/rfc5869>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.

13.2. Informative References

- [AFI] "Address Family Identifier (AFIs)", ADDRESS FAMILY NUMBERS <http://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml?>, February 2007.
- [GTP-3GPP] "General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)", TS.29.281 <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1699>, January 2015.

- [I-D.herbert-intarea-ila]
Herbert, T. and P. Lapukhov, "Identifier-locator addressing for IPv6", draft-herbert-intarea-ila-01 (work in progress), March 2018.
- [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.
- [I-D.ietf-lisp-eid-anonymity]
Farinacci, D., Pillay-Esnault, P., and W. Haddad, "LISP EID Anonymity", draft-ietf-lisp-eid-anonymity-12 (work in progress), March 2022.
- [I-D.ietf-lisp-eid-mobility]
Comeras, M. P., Ashtaputre, V., Maino, F., Moreno, V., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", draft-ietf-lisp-eid-mobility-09 (work in progress), January 2022.
- [I-D.ietf-lisp-gpe]
Maino, F., Lemon, J., Agarwal, P., Lewis, D., and M. Smith, "LISP Generic Protocol Extension", draft-ietf-lisp-gpe-19 (work in progress), July 2020.
- [I-D.ietf-lisp-introduction]
Cabellos, A. and D. S. (Ed.), "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-introduction-15 (work in progress), September 2021.
- [I-D.ietf-lisp-mn]
Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", draft-ietf-lisp-mn-11 (work in progress), January 2022.
- [I-D.ietf-lisp-pubsub]
Rodriguez-Natal, A., Ermagan, V., Cabellos, A., Barkai, S., and M. Boucadair, "Publish/Subscribe Functionality for LISP", draft-ietf-lisp-pubsub-09 (work in progress), June 2021.
- [I-D.ietf-nvo3-vxlan-gpe]
(Editor), F. M., (editor), L. K., and U. E. (editor), "Generic Protocol Extension for VXLAN (VXLAN-GPE)", draft-ietf-nvo3-vxlan-gpe-12 (work in progress), September 2021.

- [I-D.ietf-opsec-icmp-filtering]
Gont, F., Gont, G., and C. Pignataro, "Recommendations for filtering ICMP messages", draft-ietf-opsec-icmp-filtering-04 (work in progress), July 2013.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC1071] Braden, R., Borman, D., and C. Partridge, "Computing the Internet checksum", RFC 1071, DOI 10.17487/RFC1071, September 1988, <<https://www.rfc-editor.org/info/rfc1071>>.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <<https://www.rfc-editor.org/info/rfc2104>>.
- [RFC2890] Dommetry, G., "Key and Sequence Number Extensions to GRE", RFC 2890, DOI 10.17487/RFC2890, September 2000, <<https://www.rfc-editor.org/info/rfc2890>>.
- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, DOI 10.17487/RFC4984, September 2007, <<https://www.rfc-editor.org/info/rfc4984>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [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>>.

- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, DOI 10.17487/RFC6836, January 2013, <<https://www.rfc-editor.org/info/rfc6836>>.
- [RFC6837] Lear, E., "NERD: A Not-so-novel Endpoint ID (EID) to Routing Locator (RLOC) Database", RFC 6837, DOI 10.17487/RFC6837, January 2013, <<https://www.rfc-editor.org/info/rfc6837>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.
- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/info/rfc7835>>.
- [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>>.
- [RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <<https://www.rfc-editor.org/info/rfc8061>>.
- [RFC8111] Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "Locator/ID Separation Protocol Delegated Database Tree (LISP-DDT)", RFC 8111, DOI 10.17487/RFC8111, May 2017, <<https://www.rfc-editor.org/info/rfc8111>>.

- [RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <<https://www.rfc-editor.org/info/rfc8378>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

Appendix A. Acknowledgments

The original authors would like to thank Greg Schudel, Darrel Lewis, John Zwiebel, Andrew Partan, Dave Meyer, Isidor Kouvelas, Jesper Skriver, Fabio Maino, and members of the `lisp@ietf.org` mailing list for their feedback and helpful suggestions.

Special thanks are due to Noel Chiappa for his extensive work and thought about caching in Map-Resolvers.

The current authors would like to give a sincere thank you to the people who help put LISP on standards track in the IETF. They include Joel Halpern, Luigi Iannone, Deborah Brungard, Fabio Maino, Scott Bradner, Kyle Rose, Takeshi Takahashi, Sarah Banks, Pete Resnick, Colin Perkins, Mirja Kuhlewind, Francis Dupont, Benjamin Kaduk, Eric Rescorla, Alvaro Retana, Alexey Melnikov, Alissa Cooper, Suresh Krishnan, Alberto Rodriguez-Natal, Vina Ermagen, Mohamed Boucadair, Brian Trammell, Sabrina Tanamal, and John Drake. The contributions they offered greatly added to the security, scale, and robustness of the LISP architecture and protocols.

Appendix B. Document Change Log

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-ietf-lisp-rfc6833bis-31

- o Posted May 2020.
- o Added reference to 6834bis when describing the syntax and action of the Map-Version field in the Map-Reply section. This is so we can advance the Map-Versioning draft rfc6834bis to proposed standard.

B.2. Changes to draft-ietf-lisp-rfc6833bis-26

- o Posted November 2019.
- o Fixed the required (MUST implement) authentication algorithms.
- o Fixed a large set of minor comments and edits.

B.3. Changes to draft-ietf-lisp-rfc6833bis-25

- o Posted June 2019.
- o Added change requested by Mirja describing Record Count in an EID-record.

- o Fixed Requirements Notation section per Pete.
 - o Added KDF for shared-secret
 - o Specified several rate-limiters for control messages
- B.4. Changes to draft-ietf-lisp-rfc6833bis-24
- o Posted February 2019.
 - o Added suggested text from Albert that Benjamin Kaduk agreed with.
 - o Added suggested editorial comments from Alvaro's review.
 - o Ran document through IDnits. Fixed bugs found.
- B.5. Changes to draft-ietf-lisp-rfc6833bis-23
- o Posted December 2018.
 - o Added to Security Considerations section that deployments that care about prefix over claiming should use LISP-SEC.
 - o Added to Security Considerations section that DTLS or LISP-crypto be used for control-plane privacy.
 - o Make LISP-SEC a normative reference.
 - o Make it more clear where field descriptions are spec'ed when referencing to the same fields in other packet types.
- B.6. Changes to draft-ietf-lisp-rfc6833bis-22
- o Posted week after IETF November 2018.
 - o No longer need to use IPSEC for replay attacks.
- B.7. Changes to draft-ietf-lisp-rfc6833bis-21
- o Posted early November 2018.
 - o Added I-bit back in because its necessary to use for Map-Register replay attack scenarios. The Map-Server tracks the nonce per xTR-ID to detect duplicate or replayed Map-Register messages.

- B.8. Changes to draft-ietf-lisp-rfc6833bis-20
- o Posted late October 2018.
 - o Changed description about "reserved" bits to state "reserved and unassigned".
 - o Make it more clear how Map-Register nonce processing is performed in an ETR and Map-Server.
- B.9. Changes to draft-ietf-lisp-rfc6833bis-19
- o Posted mid October 2018.
 - o Added Fabio text to the Security Considerations section.
- B.10. Changes to draft-ietf-lisp-rfc6833bis-18
- o Posted mid October 2018.
 - o Fixed comments from Eric after more email clarity.
- B.11. Changes to draft-ietf-lisp-rfc6833bis-17
- o Posted early October 2018.
 - o Changes to reflect comments from Sep 27th Telechat.
 - o Added all flag bit definitions as request for allocation in IANA Considerations section.
 - o Added an applicability statement in section 1 to address security concerns from Telechat.
 - o Moved m-bit description and IANA request to draft-ietf-lisp-mn.
 - o Moved I-bit description and IANA request to draft-ietf-lisp-pubsub.
- B.12. Changes to draft-ietf-lisp-rfc6833bis-16
- o Posted Late-September 2018.
 - o Re-wrote Security Considerations section. Thanks Albert.
 - o Added Alvaro text to be more clear about IANA actions.

- B.13. Changes to draft-ietf-lisp-rfc6833bis-15
- o Posted mid-September 2018.
 - o Changes to reflect comments from Colin and Mirja.
- B.14. Changes to draft-ietf-lisp-rfc6833bis-14
- o Posted September 2018.
 - o Changes to reflect comments from Genart, RTGarea, and Secdir reviews.
- B.15. Changes to draft-ietf-lisp-rfc6833bis-13
- o Posted August 2018.
 - o Final editorial changes before RFC submission for Proposed Standard.
 - o Added section "Changes since RFC 6833" so implementators are informed of any changes since the last RFC publication.
- B.16. Changes to draft-ietf-lisp-rfc6833bis-12
- o Posted late July 2018.
 - o Moved RFC6830bis and RFC6834bis to Normative References.
- B.17. Changes to draft-ietf-lisp-rfc6833bis-11
- o Posted July 2018.
 - o Fixed Luigi editorial comments to ready draft for RFC status and ran through IDNITs again.
- B.18. Changes to draft-ietf-lisp-rfc6833bis-10
- o Posted after LISP WG at IETF week March.
 - o Move AD field encoding after S-bit in the ECM packet format description section.
 - o Say more about when the new Drop actions should be sent.

B.19. Changes to draft-ietf-lisp-rfc6833bis-09

- o Posted March IETF week 2018.
- o Fixed editorial comments submitted by document shepherd Luigi Iannone.

B.20. Changes to draft-ietf-lisp-rfc6833bis-08

- o Posted March 2018.
- o Added RLOC-probing algorithm.
- o Added Solicit-Map Request algorithm.
- o Added several mechanisms (from 6830bis) regarding Routing Locator Reachability.
- o Added port 4342 to IANA Considerations section.

B.21. Changes to draft-ietf-lisp-rfc6833bis-07

- o Posted December 2017.
- o Make it more clear in a couple of places that RLOCs are used to locate ETRs more so than for Map-Server Map-Request forwarding.
- o Make it clear that "encapsulated" for a control message is an ECM based message.
- o Make it more clear what messages use source-port 4342 and which ones use destination-port 4342.
- o Don't make DDT references when the mapping transport system can be of any type and the referenced text is general to it.
- o Generalize text when referring to the format of an EID-prefix. Can use other AFIs than IPv4 and IPv6.
- o Many editorial changes to clarify text.
- o Changed some "must", "should", and "may" to capitalized.
- o Added definitions for Map-Request and Map-Reply messages.
- o Ran document through IDNITs.

B.22. Changes to draft-ietf-lisp-rfc6833bis-06

- o Posted October 2017.
- o Spec the I-bit to include the xTR-ID in a Map-Request message to be consistent with the Map-Register message and to anticipate the introduction of pubsub functionality to allow Map-Requests to subscribe to RLOC-set changes.
- o Updated references for individual submissions that became working group documents.
- o Updated references for working group documents that became RFCs.

B.23. Changes to draft-ietf-lisp-rfc6833bis-05

- o Posted May 2017.
- o Update IANA Considerations section based on new requests from this document and changes from what was requested in [RFC6830].

B.24. Changes to draft-ietf-lisp-rfc6833bis-04

- o Posted May 2017.
- o Clarify how the Key-ID field is used in Map-Register and Map-Notify messages. Break the 16-bit field into a 8-bit Key-ID field and a 8-bit Algorithm-ID field.
- o Move the Control-Plane codepoints from the IANA Considerations section of RFC6830bis to the IANA Considerations section of this document.
- o In the "LISP Control Packet Type Allocations" section, indicate how message Types are IANA allocated and how experimental RFC8113 sub-types should be requested.

B.25. Changes to draft-ietf-lisp-rfc6833bis-03

- o Posted April 2017.
- o Add types 9-14 and specify they are not assigned.
- o Add the "LISP Shared Extension Message" type and point to RFC8113.

B.26. Changes to draft-ietf-lisp-rfc6833bis-02

- o Posted April 2017.
- o Clarify that the LISP Control-Plane document defines how the LISP Data-Plane uses Map-Requests with either the SMR-bit set or the P-bit set supporting mapping updates and RLOC-probing. Indicating that other Data-Planes can use the same mechanisms or their own defined mechanisms to achieve the same functionality.

B.27. Changes to draft-ietf-lisp-rfc6833bis-01

- o Posted March 2017.
- o Include references to new RFCs published.
- o Remove references to self.
- o Change references from RFC6830 to RFC6830bis.
- o Add two new action/reasons to a Map-Reply has posted to the LISP WG mailing list.
- o In intro section, add refernece to I-D.ietf-lisp-introduction.
- o Removed Open Issues section and references to "experimental".

B.28. Changes to draft-ietf-lisp-rfc6833bis-00

- o Posted December 2016.
- o Created working group document from draft-farinacci-lisp-rfc6833-00 individual submission. No other changes made.

B.29. Changes to draft-farinacci-lisp-rfc6833bis-00

- o Posted November 2016.
- o This is the initial draft to turn RFC 6833 into RFC 6833bis.
- o The document name has changed from the "Locator/ID Separation Protocol (LISP) Map-Server Interface" to the "Locator/ID Separation Protocol (LISP) Control-Plane".
- o The fundamental change was to move the Control-Plane messages from RFC 6830 to this document in an effort so any IETF developed or industry created Data-Plane could use the LISP mapping system and Control-Plane.

- o Update Control-Plane messages to incorporate what has been implemented in products during the early phase of LISP development but wasn't able to make it into RFC6830 and RFC6833 to make the Experimental RFC deadline.
- o Indicate there may be nodes in the mapping system that are not MRs or MSs, that is a ALT-node or a DDT-node.
- o Include LISP-DDT in Map-Resolver section and explain how they maintain a referral-cache.
- o Removed open issue about additional state in Map-Servers. With [RFC8111], Map-Servers have the same registration state and can give Map-Resolvers complete information in ms-ack Map-Referral messages.
- o Make reference to the LISP Threats Analysis RFC [RFC7835].

Authors' Addresses

Dino Farinacci
lispers.net

EEmail: farinacci@gmail.com

Fabio Maino
Cisco Systems

EEmail: fmaino@cisco.com

Vince Fuller
vaf.net Internet Consulting

EEmail: vaf@vaf.net

Albert Cabellos
UPC/BarcelonaTech
Campus Nord, C. Jordi Girona 1-3
Barcelona, Catalunya
Spain

EEmail: acabello@ac.upc.edu

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: April 21, 2019

V. Ermagan
A. Rodriguez-Natal
F. Coras
C. Moberg
R. Rahman
Cisco Systems
A. Cabellos-Aparicio
Technical University of Catalonia
F. Maino
Cisco Systems
October 18, 2018

LISP YANG Model
draft-ietf-lisp-yang-09

Abstract

This document describes a YANG data model to use with the Locator/ID Separation Protocol (LISP).

The YANG modules in this document conform to the Network Management Datastore Architecture (NMDA).

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 April 21, 2019.

Copyright Notice

Copyright (c) 2018 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
1.2. Tree Diagrams	3
2. LISP Module	3
2.1. Module Structure	3
2.2. Module Definition	6
3. LISP-ITR Module	16
3.1. Module Structure	16
3.2. Module Definition	21
4. LISP-ETR Module	25
4.1. Module Structure	25
4.2. Module Definition	27
5. LISP-Map-Server Module	32
5.1. Module Structure	32
5.2. Module Definition	40
6. LISP-Map-Resolver Module	46
6.1. Module Structure	47
6.2. Module Definition	47
7. LISP-Address-Types Module	49
7.1. Module Definition	49
7.2. Data Model examples	64
7.2.1. LISP protocol instance	64
7.2.2. LISP ITR	65
7.2.3. LISP ETR	66
7.2.4. LISP Map-Server	68
8. Acknowledgments	69
9. IANA Considerations	69
10. Security Considerations	71
11. Normative References	74
Authors' Addresses	76

1. Introduction

The Locator/ID Separation Protocol (LISP) defines several network elements subject to be configured. This document presents the YANG data models required for basic configuration of all major LISP [RFC6830] elements. The models also capture some essential operational data elements as well.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Tree Diagrams

This document uses the graphical representation of data models defined in [RFC8340].

2. LISP Module

This module is the base LISP module that is augmented in multiple models to represent various LISP device roles.

2.1. Module Structure

```

module: ietf-lisp
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol:
      +--rw lisp
        +--rw locator-sets
          +--rw locator-set* [locator-set-name]
            +--rw locator-set-name      string
            +--rw (locator-type)?
              +---:(local-interface)
                +--rw interface* [interface-ref]
                  +--rw interface-ref    if:interface-ref
                  +--rw priority?        uint8
                  +--rw weight?          uint8
                  +--rw multicast-priority? uint8
                  +--rw multicast-weight? uint8
              +---:(general-locator)
                +--rw locator* [id]
                  +--rw id                string
                  +--rw locator-address
                    +--rw address-type
                    |   lisp-address-family-ref
                    +--rw virtual-network-id?
                    |   instance-id-type
                    +--rw (address)?
                      +---:(no-address)
                        | +--rw no-address?          empty
                      +---:(ipv4)
                        | +--rw ipv4?
  
```

```

|             inet:ipv4-address
+---:(ipv4-prefix)
|   +---rw ipv4-prefix?
|             inet:ipv4-prefix
+---:(ipv6)
|   +---rw ipv6?
|             inet:ipv6-address
+---:(ipv6-prefix)
|   +---rw ipv6-prefix?
|             inet:ipv6-prefix
+---:(mac)
|   +---rw mac?
|             yang:mac-address
+---:(distinguished-name)
|   +---rw distinguished-name?
|             distinguished-name-type
+---:(as-number)
|   +---rw as-number?
|             inet:as-number
+---:(null-address)
|   +---rw null-address
|   +---rw address?    empty
+---:(afi-list)
|   +---rw afi-list
|   +---rw address-list*
|             simple-address
+---:(instance-id)
|   +---rw instance-id
|   +---rw iid?
|   |       instance-id-type
|   +---rw mask-length?    uint8
|   +---rw address?        simple-address
+---:(as-number-lcaf)
|   +---rw as-number-lcaf
|   +---rw as?             inet:as-number
|   +---rw address?        simple-address
+---:(application-data)
|   +---rw application-data
|   +---rw address?
|   |       simple-address
|   +---rw protocol?          uint8
|   +---rw ip-tos?            int32
|   +---rw local-port-low?
|   |       inet:port-number
|   +---rw local-port-high?
|   |       inet:port-number
|   +---rw remote-port-low?
|   |       inet:port-number

```

```

    +--rw remote-port-high?
        inet:port-number
+--:(geo-coordinates)
    +--rw geo-coordinates
        +--rw latitude?          bits
        +--rw latitude-degrees? uint8
        +--rw latitude-minutes? uint8
        +--rw latitude-seconds? uint8
        +--rw longitude?        bits
        +--rw longitude-degrees? uint16
        +--rw longitude-minutes? uint8
        +--rw longitude-seconds? uint8
        +--rw altitude?         int32
        +--rw address?
            simple-address
+--:(nat-traversal)
    +--rw nat-traversal
        +--rw ms-udp-port?      uint16
        +--rw etr-udp-port?    uint16
        +--rw global-etr-rloc?
            |
            | simple-address
        +--rw ms-rloc?
            |
            | simple-address
        +--rw private-etr-rloc?
            |
            | simple-address
        +--rw rtr-rlocs*
            simple-address
+--:(explicit-locator-path)
    +--rw explicit-locator-path
        +--rw hop* [hop-id]
            +--rw hop-id      string
            +--rw address?    simple-address
            +--rw lrs-bits?   bits
+--:(source-dest-key)
    +--rw source-dest-key
        +--rw source?        simple-address
        +--rw dest?          simple-address
+--:(key-value-address)
    +--rw key-value-address
        +--rw key?           simple-address
        +--rw value?        simple-address
+--:(service-path)
    +--rw service-path
        +--rw service-path-id?
            |
            | service-path-id-type
        +--rw service-index?  uint8
+--rw priority?              uint8
+--rw weight?                uint8

```

```

|           +--rw multicast-priority?  uint8
|           +--rw multicast-weight?    uint8
+--rw lisp-role* [lisp-role-type]
|   +--rw lisp-role-type    lisp-role-ref
+--rw lisp-router-id
|   +--rw site-id?         uint64
|   +--rw xtr-id?         lisp:xtr-id-type
+--rw virtual-networks
|   +--rw virtual-network* [vni]
|       +--rw vni          lcaf:instance-id-type
|       +--rw ni-name?
|           -> /ni:network-instances/network-instance/name

```

2.2. Module Definition

```

<CODE BEGINS> file "ietf-lisp@2018-06-29.yang"
module ietf-lisp {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp";

  prefix lisp;

  import ietf-interfaces {
    prefix if;
    reference
      "RFC 8343: A YANG Data Model for Interface Management";
  }
  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note
  import ietf-lisp-address-types {
    prefix lcaf;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-yang-types {
    prefix yang;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
      (NMDA version)";
  }
  import ietf-network-instance {
    prefix "ni";
    // RFC Ed.: replace occurrences of YYYY with actual RFC number
    // of draft-ietf-rtgwg-ni-model and remove this note

```

```
reference
  "RFC YYYY: YANG Model for Network Instances";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web: <http://tools.ietf.org/wg/lisp/>
  WG List: <mailto:lisp@ietf.org>

  Editor: Vina Ermagan
         <mailto:vermagan@cisco.com>

  Editor: Alberto Rodriguez-Natal
         <mailto:natal@cisco.com>

  Editor: Reshad Rahman
         <mailto:rrahman@cisco.com>";
description
  "This YANG module defines the generic parameters for LISP.
  The module can be extended by vendors to define vendor-specific
  LISP parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.
  ";

reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6830";
}

/*
 * Identity definitions
```

```
*/
identity lisp {
  base "rt:control-plane-protocol";
  description "LISP protocol.";
  reference
    "RFC 6830: The Locator/ID Separation Protocol (LISP).";
}

identity lisp-role {
  description
    "LISP router role.";
}
identity itr {
  base lisp-role;
  description
    "LISP ITR.";
}
identity pitr {
  base lisp-role;
  description
    "LISP PITR.";
}
identity etr {
  base lisp-role;
  description
    "LISP ETR.";
}
identity petr {
  base lisp-role;
  description
    "LISP PETR.";
}
}
identity mapping-system {
  description
    "Mapping System interface";
}
identity single-node-mapping-system {
  base mapping-system;
  description
    "logically singular Map Server";
}
}
typedef mapping-system-ref {
  type identityref {
    base mapping-system;
  }
  description
    "Mapping System reference";
}
}
```

```
typedef lisp-role-ref {
  type identityref {
    base lisp-role;
  }
  description
    "LISP role reference";
}
typedef map-reply-action {
  type enumeration {
    enum no-action {
      value 0;
      description
        "Mapping is kept alive and no encapsulation occurs.";
    }
    enum natively-forward {
      value 1;
      description
        "Matching packets are not encapsulated or dropped but
        natively forwarded.";
    }
    enum send-map-request {
      value 2;
      description
        "Matching packets invoke Map-Requests.";
    }
    enum drop {
      value 3;
      description
        "Matching packets are dropped.";
    }
  }
  description
    "Defines the lisp map-cache ACT type";
  reference "https://tools.ietf.org/html/rfc6830#section-6.1.4";
}
typedef eid-id {
  type string;
  description
    "Type encoding of lisp-addresses to be generally used in EID
    keyed lists.";
}
typedef auth-key-type {
  type enumeration {
    enum none {
      value 0;
      description
        "No authentication.";
    }
  }
}
```

```
    enum hmac-sha-1-96 {
      value 1;
      description
        "HMAC-SHA-1-96 (RFC2404) authentication is used.";
    }
    enum hmac-sha-256-128 {
      value 2;
      description
        "HMAC-SHA-256-128 (RFC4868) authentication is used.";
    }
  }
  description
    "Enumeration of the authentication mechanisms supported by
    LISP.";
  reference
    "https://tools.ietf.org/html/rfc6830#section-6.1.6";
}
typedef xtr-id-type {
  type binary {
    length "16";
  }
  description
    "128 bit xTR identifier.";
}

grouping locator-properties {
  description
    "Properties of a RLOC";
  leaf priority {
    type uint8;
    description
      "Locator priority.";
  }
  leaf weight {
    type uint8;
    description
      "Locator weight.";
  }
  leaf multicast-priority {
    type uint8;
    description
      "Locator's multicast priority";
  }
  leaf multicast-weight {
    type uint8;
    description
      "Locator's multicast weight";
  }
}
```

```
    }

    grouping locators-grouping {
      description
        "Group that defines a list of LISP locators.";
      list locator {
        key "id";
        description
          "List of routing locators";
        leaf id {
          type string {
            length "1..64";
          }
          description
            "Locator id";
        }
        container locator-address {
          uses lcaf:lisp-address;
          description
            "The locator address provided in LISP canonical
            address format.";
        }
        uses locator-properties;
      }
    }

    grouping local-locators-grouping {
      description
        "Group that defines a list of LISP locators.";
      list interface {
        key "interface-ref";
        description
          "The address type of the locator";
        leaf interface-ref {
          type if:interface-ref;
          description
            "The name of the interface supporting the locator.";
        }
        uses locator-properties;
      }
    }

    grouping mapping {
      description
        "Group that defines a LISP mapping.";
      container eid {
        uses lcaf:lisp-address;
      }
    }
  }
}
```

```
    description
      "End-host Identifier (EID) to be mapped to a list of
       locators";
  }
  leaf time-to-live {
    type uint32;
    units minutes;
    description
      "Mapping validity period in minutes.";
  }
  leaf creation-time {
    type yang:date-and-time;
    config false;
    description
      "Time when the mapping was created.";
  }
  leaf authoritative {
    type bits {
      bit A {
        description
          "Authoritative bit.";
      }
    }
    description
      "Bit that indicates if mapping comes from an
       authoritative source.";
  }
  leaf static {
    type boolean;
    default "false";
    description
      "This leaf should be true if the mapping is static.";
  }
  choice locator-list {
    description
      "list of locartors are either negative, or positive.";
    case negative-mapping {
      leaf map-reply-action {
        type map-reply-action;
        description
          "Forwarding action for a negative mapping.";
      }
    }
    case positive-mapping {
      container rlocs {
        uses locators-grouping;
        description
          "List of locators for a positive mapping.";
      }
    }
  }
}
```

```

    }
  }
}

grouping mappings {
  description
    "Group that defines a list of LISP mappings.";
  list virtual-network {
    key "vni";
    description
      "Virtual network to which the mappings belong.";
    leaf vni {
      type lcaf:instance-id-type;
      description
        "Virtual network identifier.";
    }
    container mappings {
      description
        "Mappings within the virtual network.";
      list mapping {
        key "id";
        description
          "List of EID to RLOCs mappings.";
        leaf id {
          type eid-id;
          description
            "Id that uniquely identifies a mapping.";
        }
        uses mapping;
      }
    }
  }
}

augment "/rt:routing/rt:control-plane-protocols"
  + "/rt:control-plane-protocol" {
  when "derived-from-or-self(rt:type, 'lisp:lisp')" {
    description
      "This augmentation is only valid for a control-plane protocol
      instance of LISP.";
  }
  description "LISP protocol ietf-routing module
  control-plane-protocol augmentation.";

  container lisp {
    description
      "Parameters for the LISP subsystem.";
  }
}

```

```
container locator-sets {
  description
    "Container that defines a named locator set which can be
    referenced elsewhere.";
  list locator-set {
    key "locator-set-name";
    description
      "Multiple locator sets can be defined.";
    leaf locator-set-name {
      type string {
        length "1..64";
      }
      description
        "Locator set name";
    }
    choice locator-type {
      description
        "Locator sets can be based on local interfaces, or
        general locators.";
      case local-interface {
        uses local-locators-grouping;
        description
          "List of locators in this set based on local
          interfaces.";
      }
      case general-locator {
        uses locators-grouping;
        description
          "List of locators in this set based on lisp-address.";
      }
    }
  }
}

list lisp-role {
  key lisp-role-type;
  description
    "List of lisp device roles such as MS, MR, ITR,
    PITR, ETR or PETR.";
  leaf lisp-role-type {
    type lisp-role-ref;
    description
      "The type of LISP device - identity derived from the
      'lisp-device' base identity.";
  }
}

container lisp-router-id {
```



```

}
<CODE ENDS>

```

3. LISP-ITR Module

This module captures the configuration data model of a LISP ITR. The model also captures some operational data elements.

3.1. Module Structure

```

module: ietf-lisp-itr
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol/lisp:lisp:
      +--rw itr!
        +--rw rloc-probing!
          | +--rw interval?          uint16
          | +--rw retries?          uint8
          | +--rw retries-interval? uint16
        +--rw itr-rlocs?          leafref
        +--rw map-resolvers
          | +--rw map-resolver*      inet:ip-address
        +--rw proxy-etr
          | +--rw proxy-etr-address* inet:ip-address
        +--rw map-cache
          +--ro size?              uint32
          +--ro limit?            uint32
          +--rw virtual-network* [vni]
            +--rw vni              lcaf:instance-id-type
            +--rw mappings
              +--rw mapping* [id]
                +--rw id              eid-id
                +--rw eid
                  +--rw address-type
                  | +--rw lisp-address-family-ref
                +--rw virtual-network-id?
                  | +--rw instance-id-type
                +--rw (address)?
                  +--:(no-address)
                  | +--rw no-address?          empty
                  +--:(ipv4)
                  | +--rw ipv4?
                  | +--rw inet:ipv4-address
                  +--:(ipv4-prefix)
                  | +--rw ipv4-prefix?
                  | +--rw inet:ipv4-prefix
                  +--:(ipv6)
                  | +--rw ipv6?
                  | +--rw inet:ipv6-address

```

```

+--:(ipv6-prefix)
|  +--rw ipv6-prefix?
|     inet:ipv6-prefix
+--:(mac)
|  +--rw mac?
|     yang:mac-address
+--:(distinguished-name)
|  +--rw distinguished-name?
|     distinguished-name-type
+--:(as-number)
|  +--rw as-number?
|     inet:as-number
+--:(null-address)
|  +--rw null-address
|     +--rw address?    empty
+--:(afi-list)
|  +--rw afi-list
|     +--rw address-list*  simple-address
+--:(instance-id)
|  +--rw instance-id
|     +--rw iid?           instance-id-type
|     +--rw mask-length?  uint8
|     +--rw address?      simple-address
+--:(as-number-lcaf)
|  +--rw as-number-lcaf
|     +--rw as?           inet:as-number
|     +--rw address?     simple-address
+--:(application-data)
|  +--rw application-data
|     +--rw address?
|         simple-address
|     +--rw protocol?    uint8
|     +--rw ip-tos?      int32
|     +--rw local-port-low?
|         inet:port-number
|     +--rw local-port-high?
|         inet:port-number
|     +--rw remote-port-low?
|         inet:port-number
|     +--rw remote-port-high?
|         inet:port-number
+--:(geo-coordinates)
|  +--rw geo-coordinates
|     +--rw latitude?    bits
|     +--rw latitude-degrees?  uint8
|     +--rw latitude-minutes?  uint8
|     +--rw latitude-seconds?  uint8
|     +--rw longitude?    bits

```

```

    +--rw longitude-degrees?  uint16
    +--rw longitude-minutes?  uint8
    +--rw longitude-seconds?  uint8
    +--rw altitude?           int32
    +--rw address?
        simple-address
+--:(nat-traversal)
    +--rw nat-traversal
        +--rw ms-udp-port?    uint16
        +--rw etr-udp-port?  uint16
        +--rw global-etr-rloc?
            |
            | simple-address
        +--rw ms-rloc?
            |
            | simple-address
        +--rw private-etr-rloc?
            |
            | simple-address
        +--rw rtr-rlocs*
            simple-address
+--:(explicit-locator-path)
    +--rw explicit-locator-path
        +--rw hop* [hop-id]
            +--rw hop-id      string
            +--rw address?    simple-address
            +--rw lrs-bits?   bits
+--:(source-dest-key)
    +--rw source-dest-key
        +--rw source?        simple-address
        +--rw dest?          simple-address
+--:(key-value-address)
    +--rw key-value-address
        +--rw key?           simple-address
        +--rw value?         simple-address
+--:(service-path)
    +--rw service-path
        +--rw service-path-id?
            |
            | service-path-id-type
        +--rw service-index?  uint8
+--rw time-to-live?          uint32
+--ro creation-time?         yang:date-and-time
+--rw authoritative?        bits
+--rw static?                boolean
+--rw (locator-list)?
    +--:(negative-mapping)
    | +--rw map-reply-action?  map-reply-action
    +--:(positive-mapping)
    +--rw rlocs
        +--rw locator* [id]
            +--rw id           string

```

```

+--rw locator-address
|
+--rw address-type
|   lisp-address-family-ref
+--rw virtual-network-id?
|   instance-id-type
+--rw (address)?
|   +--:(no-address)
|   |   +--rw no-address?
|   |   |   empty
|   +--:(ipv4)
|   |   +--rw ipv4?
|   |   |   inet:ipv4-address
|   +--:(ipv4-prefix)
|   |   +--rw ipv4-prefix?
|   |   |   inet:ipv4-prefix
|   +--:(ipv6)
|   |   +--rw ipv6?
|   |   |   inet:ipv6-address
|   +--:(ipv6-prefix)
|   |   +--rw ipv6-prefix?
|   |   |   inet:ipv6-prefix
|   +--:(mac)
|   |   +--rw mac?
|   |   |   yang:mac-address
|   +--:(distinguished-name)
|   |   +--rw distinguished-name?
|   |   |   distinguished-name-type
|   +--:(as-number)
|   |   +--rw as-number?
|   |   |   inet:as-number
|   +--:(null-address)
|   |   +--rw null-address
|   |   |   +--rw address?   empty
|   +--:(afi-list)
|   |   +--rw afi-list
|   |   |   +--rw address-list*
|   |   |   |   simple-address
|   +--:(instance-id)
|   |   +--rw instance-id
|   |   |   +--rw iid?
|   |   |   |   instance-id-type
|   |   +--rw mask-length?   uint8
|   |   +--rw address?
|   |   |   simple-address
|   +--:(as-number-lcaf)
|   |   +--rw as-number-lcaf
|   |   |   +--rw as?
|   |   |   |   inet:as-number

```

```

    +--rw address?
       simple-address
+--:(application-data)
  +--rw application-data
     +--rw address?
        |   simple-address
     +--rw protocol?
        |   uint8
     +--rw ip-tos?
        |   int32
     +--rw local-port-low?
        |   inet:port-number
     +--rw local-port-high?
        |   inet:port-number
     +--rw remote-port-low?
        |   inet:port-number
     +--rw remote-port-high?
        |   inet:port-number
+--:(geo-coordinates)
  +--rw geo-coordinates
     +--rw latitude?
        |   bits
     +--rw latitude-degrees?
        |   uint8
     +--rw latitude-minutes?
        |   uint8
     +--rw latitude-seconds?
        |   uint8
     +--rw longitude?
        |   bits
     +--rw longitude-degrees?
        |   uint16
     +--rw longitude-minutes?
        |   uint8
     +--rw longitude-seconds?
        |   uint8
     +--rw altitude?
        |   int32
     +--rw address?
        |   simple-address
+--:(nat-traversal)
  +--rw nat-traversal
     +--rw ms-udp-port?
        |   uint16
     +--rw etr-udp-port?
        |   uint16
     +--rw global-etr-rloc?
        |   simple-address

```



```
// and remove this note
import ietf-lisp {
  prefix lisp;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-inet-types {
  prefix inet;
  reference "RFC 6991: Common YANG Data Types";
}
import ietf-routing {
  prefix "rt";
  reference
    "RFC 8349: A YANG Data Model for Routing Management
    (NMDA version)";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web: <http://tools.ietf.org/wg/lisp/>
  WG List: <mailto:lisp@ietf.org>

  Editor: Vina Ermagan
         <mailto:vermagan@cisco.com>

  Editor: Alberto Rodriguez-Natal
         <mailto:natal@cisco.com>

  Editor: Reshad Rahman
         <mailto:rrahman@cisco.com>";
description
  "This YANG module defines the generic parameters for a LISP
  ITR. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.
  ";
```

```
reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6830";
}
augment "/rt:routing/rt:control-plane-protocols"
+ "/rt:control-plane-protocol/lisp:lisp" {
  when "lisp:lisp-role/lisp:lisp-role-type = 'lisp:itr' or
    lisp:lisp-role/lisp:lisp-role-type = 'lisp:pitr'" {
    description
      "Augment is valid when LISP role type is ITR or PITR.";
  }
  description
    "This augments LISP devices list with (P)ITR specific
    parameters.";
  container itr {
    presence "LISP (P)ITR operation enabled";
    description
      "ITR parameters";
    container rloc-probing {
      presence "RLOC probing active";
      description
        "RLOC-probing parameters";
      leaf interval {
        type uint16;
        units "seconds";
        description
          "Interval in seconds for resending the probes";
      }
    }
    leaf retries {
      type uint8;
      description
        "Number of retries for sending the probes";
    }
    leaf retries-interval {
      type uint16;
      units "seconds";
      description
        "Interval in seconds between retries when sending probes.
        The action taken if all retries fail to receive is
        impementation specific.";
    }
  }
  leaf itr-rlocs {
    type leafref {
```

```
    path "/rt:routing/rt:control-plane-protocols"
      + "/rt:control-plane-protocol/lisp:lisp"
      + "/lisp:locator-sets/lisp:locator-set"
      + "/lisp:locator-set-name";
  }
  description
    "Reference to a locator set that the (P)ITR includes in
    Map-Requests";
}
container map-resolvers {
  description
    "Map-Resolvers that the (P)ITR uses.";
  leaf-list map-resolver {
    type inet:ip-address;
    description
      "Each Map-Resolver within the list of Map-Resolvers.";
  }
}
container proxy-etr {
  when "../..//lisp:lisp-role/lisp:lisp-role-type = 'lisp:itr'" {
    description
      "Container exists only when LISP role type is ITR";
  }
  description
    "Proxy ETRs that the ITR uses.";
  leaf-list proxy-etr-address {
    type inet:ip-address;
    description
      "Proxy ETR RLOC address.";
  }
}
container map-cache {
  leaf size {
    type uint32;
    config false;
    description
      "Current number of entries in the EID-to-RLOC map-cache";
  }
  leaf limit {
    type uint32;
    config false;
    description
      "Maximum permissible number of entries in the EID-to-RLOC
      map-cache";
  }
}

uses lisp:mappings;
description
```

```

        "EID to RLOCs mappings cache.";
    }
}
}
}
<CODE ENDS>

```

4. LISP-ETR Module

This module captures the configuration data model of a LISP ETR. The model also captures some operational data elements.

4.1. Module Structure

```

module: ietf-lisp-etr
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol/lisp:lisp:
  +--rw etr!
    +--rw map-servers
      | +--rw map-server* [ms-address]
      |   +--rw ms-address      inet:ip-address
      |   +--rw auth-key?      string
      |   +--rw auth-key-type? lisp:auth-key-type
    +--rw local-eids
      +--rw virtual-network* [vni]
      +--rw vni      lcaf:instance-id-type
      +--rw eids
        +--rw local-eid* [id]
          +--rw id      lisp:eid-id
          +--rw eid-address
            +--rw address-type
            |   lisp-address-family-ref
            +--rw virtual-network-id?
            |   instance-id-type
            +--rw (address)?
              +--:(no-address)
              | +--rw no-address?      empty
              +--:(ipv4)
              | +--rw ipv4?
              |   inet:ipv4-address
              +--:(ipv4-prefix)
              | +--rw ipv4-prefix?
              |   inet:ipv4-prefix
              +--:(ipv6)
              | +--rw ipv6?
              |   inet:ipv6-address
              +--:(ipv6-prefix)
              | +--rw ipv6-prefix?

```

```

|           inet:ipv6-prefix
+--: (mac)
|   +--rw mac?
|       yang:mac-address
+--: (distinguished-name)
|   +--rw distinguished-name?
|       distinguished-name-type
+--: (as-number)
|   +--rw as-number?
|       inet:as-number
+--: (null-address)
|   +--rw null-address
|       +--rw address?    empty
+--: (afi-list)
|   +--rw afi-list
|       +--rw address-list*    simple-address
+--: (instance-id)
|   +--rw instance-id
|       +--rw iid?             instance-id-type
|       +--rw mask-length?    uint8
|       +--rw address?        simple-address
+--: (as-number-lcaf)
|   +--rw as-number-lcaf
|       +--rw as?             inet:as-number
|       +--rw address?        simple-address
+--: (application-data)
|   +--rw application-data
|       +--rw address?
|           |           simple-address
|       +--rw protocol?          uint8
|       +--rw ip-tos?            int32
|       +--rw local-port-low?
|           |           inet:port-number
|       +--rw local-port-high?
|           |           inet:port-number
|       +--rw remote-port-low?
|           |           inet:port-number
|       +--rw remote-port-high?
|           |           inet:port-number
+--: (geo-coordinates)
|   +--rw geo-coordinates
|       +--rw latitude?          bits
|       +--rw latitude-degrees?  uint8
|       +--rw latitude-minutes?  uint8
|       +--rw latitude-seconds?  uint8
|       +--rw longitude?         bits
|       +--rw longitude-degrees? uint16
|       +--rw longitude-minutes? uint8

```

```

    +--rw longitude-seconds?  uint8
    +--rw altitude?          int32
    +--rw address?
        simple-address
+--:(nat-traversal)
  +--rw nat-traversal
    +--rw ms-udp-port?      uint16
    +--rw etr-udp-port?    uint16
    +--rw global-etr-rloc?
        |
        | simple-address
    +--rw ms-rloc?
        |
        | simple-address
    +--rw private-etr-rloc?
        |
        | simple-address
    +--rw rtr-rlocs*
        simple-address
+--:(explicit-locator-path)
  +--rw explicit-locator-path
    +--rw hop* [hop-id]
        +--rw hop-id      string
        +--rw address?    simple-address
        +--rw lrs-bits?   bits
+--:(source-dest-key)
  +--rw source-dest-key
    +--rw source?         simple-address
    +--rw dest?           simple-address
+--:(key-value-address)
  +--rw key-value-address
    +--rw key?            simple-address
    +--rw value?         simple-address
+--:(service-path)
  +--rw service-path
    +--rw service-path-id?
        |
        | service-path-id-type
    +--rw service-index?  uint8
+--rw rlocs?              leafref
+--rw record-ttl?        uint32
+--rw want-map-notify?   boolean
+--rw proxy-reply?       boolean
+--rw registration-interval?  uint16

```

4.2. Module Definition

```

<CODE BEGINS> file "ietf-lisp-etr@2018-06-29.yang"
module ietf-lisp-etr {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-etr";

```

```
prefix lisp-etr;

// RFC Ed.: replace occurrences of XXXX with actual RFC number
// and remove this note
import ietf-lisp {
  prefix lisp;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-lisp-address-types {
  prefix lcaf;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-inet-types {
  prefix inet;
  reference "RFC 6991: Common YANG Data Types";
}
import ietf-routing {
  prefix "rt";
  reference
    "RFC 8349: A YANG Data Model for Routing Management
    (NMDA version)";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web:    <http://tools.ietf.org/wg/lisp/>
  WG List:   <mailto:lisp@ietf.org>

  Editor:    Vina Ermagan
             <mailto:vermagan@cisco.com>

  Editor:    Alberto Rodriguez-Natal
             <mailto:natal@cisco.com>

  Editor:    Reshad Rahman
             <mailto:rrahman@cisco.com>";
description
  "This YANG module defines the generic parameters for a LISP
  ETR. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
```

set forth in Section 4.c of the IETF Trust's Legal Provisions
Relating to IETF Documents
(<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see
the RFC itself for full legal notices.

";

```
reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6830";
}
augment "/rt:routing/rt:control-plane-protocols"
  + "/rt:control-plane-protocol/lisp:lisp" {
  when "lisp:lisp-role/lisp:lisp-role-type = 'lisp:etr' or
    lisp:lisp-role/lisp:lisp-role-type = 'lisp:petr'" {
    description
      "Augment is valid when LISP device type is (P)ETR.";
  }
  description
    "This augments LISP devices list with (P)ETR specific
    parameters.";
  container etr {
    presence "LISP (P)ETR operation enabled";
    description
      "(P)ETR parameters.";

    container map-servers {
      when "../lisp:lisp-role/lisp:lisp-role-type = 'lisp:etr'" {
        description
          "Container exists only when LISP device type is ETR.";
      }
      description
        "Map-Servers that the ETR uses.";
      list map-server {
        key "ms-address";
        description
          "Each Map-Server within the list of Map-Servers.";
        leaf ms-address {
          type inet:ip-address;
          description
            "Map-Server address.";
        }
        leaf auth-key {
```

```
        type string;
        description
            "Map-Server authentication key.";
    }
    leaf auth-key-type {
        type lisp:auth-key-type;
        description
            "Map-Server authentication type.";
    }
}
}

container local-eids {
    when "../..//lisp:lisp-role/lisp:lisp-role-type = 'lisp:etr'" {
        description
            "Container exists only when LISP device type is ETR.";
    }
    description
        "Virtual networks served by the ETR.";
    list virtual-network {
        key "vni";
        description
            "Virtual network for local-EIDs.";
        leaf vni {
            type lcaf:instance-id-type;
            description
                "Virtual network identifier.";
        }
    }
    container eids {
        description
            "EIDs served by the ETR.";
        list local-eid {
            key "id";
            min-elements 1;
            description
                "List of local EIDs.";
            leaf id {
                type lisp:eid-id;
                description
                    "Unique id of local EID.";
            }
        }
        container eid-address {
            uses lcaf:lisp-address;
            description
                "EID address in generic LISP address format.";
        }
        leaf rlocs {
            type leafref {

```


5. LISP-Map-Server Module

This module captures the configuration data model of a LISP Map Server [RFC6833]. The model also captures some operational data elements.

5.1. Module Structure

```

module: ietf-lisp-mapserver
augment /rt:routing/rt:control-plane-protocols
  /rt:control-plane-protocol/lisp:lisp:
  +--rw map-server!
    +--rw sites
      |
      |   +--rw site* [site-id]
      |   |   +--rw site-id      uint64
      |   |   +--rw auth-key
      |   |   |   +--rw auth-key-value?  string
      |   |   |   +--rw auth-key-type*   lisp:auth-key-type
      |   +--rw virtual-network-ids
      |   |   +--rw virtual-network-identifier* [vni]
      |   |   |   +--rw vni          lcaf:instance-id-type
      |   |   |   +--rw mappings
      |   |   |   |   +--rw mapping* [eid-id]
      |   |   |   |   |   +--rw eid-id          lisp:eid-id
      |   |   |   |   |   +--rw eid-address
      |   |   |   |   |   |   +--rw address-type
      |   |   |   |   |   |   |   lisp-address-family-ref
      |   |   |   |   |   |   +--rw virtual-network-id?
      |   |   |   |   |   |   |   instance-id-type
      |   |   |   |   |   |   +--rw (address)?
      |   |   |   |   |   |   |   +---:(no-address)
      |   |   |   |   |   |   |   |   +--rw no-address?          empty
      |   |   |   |   |   |   |   +---:(ipv4)
      |   |   |   |   |   |   |   |   +--rw ipv4?
      |   |   |   |   |   |   |   |   |   inet:ipv4-address
      |   |   |   |   |   |   |   +---:(ipv4-prefix)
      |   |   |   |   |   |   |   |   +--rw ipv4-prefix?
      |   |   |   |   |   |   |   |   |   inet:ipv4-prefix
      |   |   |   |   |   |   |   +---:(ipv6)
      |   |   |   |   |   |   |   |   +--rw ipv6?
      |   |   |   |   |   |   |   |   |   inet:ipv6-address
      |   |   |   |   |   |   |   +---:(ipv6-prefix)
      |   |   |   |   |   |   |   |   +--rw ipv6-prefix?
      |   |   |   |   |   |   |   |   |   inet:ipv6-prefix
      |   |   |   |   |   |   |   +---:(mac)
      |   |   |   |   |   |   |   |   +--rw mac?
      |   |   |   |   |   |   |   |   |   yang:mac-address
      |   |   |   |   |   |   |   +---:(distinguished-name)

```

```

    |--rw distinguished-name?
        distinguished-name-type
+--:(as-number)
    |--rw as-number?
        inet:as-number
+--:(null-address)
    |--rw null-address
        |--rw address?    empty
+--:(afi-list)
    |--rw afi-list
        |--rw address-list*  simple-address
+--:(instance-id)
    |--rw instance-id
        |--rw iid?          instance-id-type
        |--rw mask-length?  uint8
        |--rw address?      simple-address
+--:(as-number-lcaf)
    |--rw as-number-lcaf
        |--rw as?          inet:as-number
        |--rw address?     simple-address
+--:(application-data)
    |--rw application-data
        |--rw address?
            |   simple-address
        |--rw protocol?    uint8
        |--rw ip-tos?      int32
        |--rw local-port-low?
            |   inet:port-number
        |--rw local-port-high?
            |   inet:port-number
        |--rw remote-port-low?
            |   inet:port-number
        |--rw remote-port-high?
            |   inet:port-number
+--:(geo-coordinates)
    |--rw geo-coordinates
        |--rw latitude?    bits
        |--rw latitude-degrees?  uint8
        |--rw latitude-minutes?  uint8
        |--rw latitude-seconds?  uint8
        |--rw longitude?    bits
        |--rw longitude-degrees?  uint16
        |--rw longitude-minutes?  uint8
        |--rw longitude-seconds?  uint8
        |--rw altitude?    int32
        |--rw address?
            |   simple-address
+--:(nat-traversal)

```

```

+--rw nat-traversal
  +--rw ms-udp-port?          uint16
  +--rw etr-udp-port?        uint16
  +--rw global-etr-rloc?
  |   simple-address
  +--rw ms-rloc?
  |   simple-address
  +--rw private-etr-rloc?
  |   simple-address
  +--rw rtr-rlocs*
  |   simple-address
+---:(explicit-locator-path)
  +--rw explicit-locator-path
  |   +--rw hop* [hop-id]
  |   |   +--rw hop-id      string
  |   |   +--rw address?    simple-address
  |   |   +--rw lrs-bits?   bits
+---:(source-dest-key)
  |   +--rw source-dest-key
  |   |   +--rw source?     simple-address
  |   |   +--rw dest?      simple-address
+---:(key-value-address)
  |   +--rw key-value-address
  |   |   +--rw key?        simple-address
  |   |   +--rw value?     simple-address
+---:(service-path)
  |   +--rw service-path
  |   |   +--rw service-path-id?
  |   |   |   service-path-id-type
  |   |   +--rw service-index?  uint8
+--rw site-id*                uint64
+--rw more-specifics-accepted? boolean
+--rw mapping-expiration-timeout? int16
+--ro first-registration-time?
  |   yang:date-and-time
+--ro last-registration-time?
  |   yang:date-and-time
+--rw mapping-records
  +--rw mapping-record* [xtr-id]
  |   +--rw xtr-id
  |   |   lisp:xtr-id-type
  +--rw site-id?              uint64
+--rw eid
  |   +--rw address-type
  |   |   lisp-address-family-ref
  +--rw virtual-network-id?
  |   instance-id-type
+--rw (address)?

```

```

+--:(no-address)
|   +--rw no-address?
|       empty
+--:(ipv4)
|   +--rw ipv4?
|       inet:ipv4-address
+--:(ipv4-prefix)
|   +--rw ipv4-prefix?
|       inet:ipv4-prefix
+--:(ipv6)
|   +--rw ipv6?
|       inet:ipv6-address
+--:(ipv6-prefix)
|   +--rw ipv6-prefix?
|       inet:ipv6-prefix
+--:(mac)
|   +--rw mac?
|       yang:mac-address
+--:(distinguished-name)
|   +--rw distinguished-name?
|       distinguished-name-type
+--:(as-number)
|   +--rw as-number?
|       inet:as-number
+--:(null-address)
|   +--rw null-address
|       +--rw address?    empty
+--:(afi-list)
|   +--rw afi-list
|       +--rw address-list*
|           simple-address
+--:(instance-id)
|   +--rw instance-id
|       +--rw iid?
|           |   instance-id-type
|       +--rw mask-length?  uint8
|       +--rw address?
|           simple-address
+--:(as-number-lcaf)
|   +--rw as-number-lcaf
|       +--rw as?          inet:as-number
|       +--rw address?    simple-address
+--:(application-data)
|   +--rw application-data
|       +--rw address?
|           |   simple-address
|       +--rw protocol?   uint8
|       +--rw ip-tos?     int32

```

```

    +--rw local-port-low?
    |   inet:port-number
    +--rw local-port-high?
    |   inet:port-number
    +--rw remote-port-low?
    |   inet:port-number
    +--rw remote-port-high?
    |   inet:port-number
+---:(geo-coordinates)
+--rw geo-coordinates
+--rw latitude?          bits
+--rw latitude-degrees? uint8
+--rw latitude-minutes? uint8
+--rw latitude-seconds? uint8
+--rw longitude?        bits
+--rw longitude-degrees?
|   uint16
+--rw longitude-minutes? uint8
+--rw longitude-seconds? uint8
+--rw altitude?         int32
+--rw address?
    simple-address
+---:(nat-traversal)
+--rw nat-traversal
+--rw ms-udp-port?      uint16
+--rw etr-udp-port?    uint16
+--rw global-etr-rloc?
|   simple-address
+--rw ms-rloc?
|   simple-address
+--rw private-etr-rloc?
|   simple-address
+--rw rtr-rlocs*
    simple-address
+---:(explicit-locator-path)
+--rw explicit-locator-path
+--rw hop* [hop-id]
    +--rw hop-id      string
    +--rw address?
    |   simple-address
    +--rw lrs-bits?  bits
+---:(source-dest-key)
+--rw source-dest-key
+--rw source?   simple-address
+--rw dest?    simple-address
+---:(key-value-address)
+--rw key-value-address
+--rw key?      simple-address

```

```

|         +--rw value?   simple-address
+---:(service-path)
|         +--rw service-path
|         |         +--rw service-path-id?
|         |         |         service-path-id-type
|         |         +--rw service-index?   uint8
+--rw time-to-live?           uint32
+--ro creation-time?
|         yang:date-and-time
+--rw authoritative?         bits
+--rw static?                 boolean
+--rw (locator-list)?
|   +---:(negative-mapping)
|   |   +--rw map-reply-action?
|   |   |   map-reply-action
|   +---:(positive-mapping)
|   |   +--rw rlocs
|   |   |   +--rw locator* [id]
|   |   |   |   +--rw id
|   |   |   |   |   string
|   |   |   +--rw locator-address
|   |   |   |   +--rw address-type
|   |   |   |   |   lisp-address-family-ref
|   |   |   +--rw virtual-network-id?
|   |   |   |   instance-id-type
|   |   +--rw (address)?
|   |   |   +---:(no-address)
|   |   |   |   +--rw no-address?
|   |   |   |   |   empty
|   |   |   +---:(ipv4)
|   |   |   |   +--rw ipv4?
|   |   |   |   |   inet:ipv4-address
|   |   |   +---:(ipv4-prefix)
|   |   |   |   +--rw ipv4-prefix?
|   |   |   |   |   inet:ipv4-prefix
|   |   |   +---:(ipv6)
|   |   |   |   +--rw ipv6?
|   |   |   |   |   inet:ipv6-address
|   |   |   +---:(ipv6-prefix)
|   |   |   |   +--rw ipv6-prefix?
|   |   |   |   |   inet:ipv6-prefix
|   |   |   +---:(mac)
|   |   |   |   +--rw mac?
|   |   |   |   |   yang:mac-address
|   |   |   +---:(distinguished-name)
|   |   |   |   +--rw distinguished-name?
|   |   |   |   |   distinguished-name-type
|   |   +---:(as-number)

```

```

+--rw as-number?
   |
   | inet:as-number
+--:(null-address)
   |
   | +--rw null-address
   |   |
   |   | +--rw address?
   |   |   |
   |   |   | empty
+--:(afi-list)
   |
   | +--rw afi-list
   |   |
   |   | +--rw address-list*
   |   |   |
   |   |   | simple-address
+--:(instance-id)
   |
   | +--rw instance-id
   |   |
   |   | +--rw iid?
   |   |   |
   |   |   | instance-id-type
+--rw mask-length?
   |
   | uint8
+--rw address?
   |
   | simple-address
+--:(as-number-lcaf)
   |
   | +--rw as-number-lcaf
   |   |
   |   | +--rw as?
   |   |   |
   |   |   | inet:as-number
+--rw address?
   |
   | simple-address
+--:(application-data)
   |
   | +--rw application-data
   |   |
   |   | +--rw address?
   |   |   |
   |   |   | simple-address
+--rw protocol?
   |
   | uint8
+--rw ip-tos?
   |
   | int32
+--rw local-port-low?
   |
   | inet:port-number
+--rw local-port-high?
   |
   | inet:port-number
+--rw remote-port-low?
   |
   | inet:port-number
+--rw remote-port-high?
   |
   | inet:port-number
+--:(geo-coordinates)
   |
   | +--rw geo-coordinates
   |   |
   |   | +--rw latitude?
   |   |   |
   |   |   | bits
+--rw latitude-degrees?
   |
   | uint8
+--rw latitude-minutes?
   |
   | uint8

```

```

+--rw latitude-seconds?
|   uint8
+--rw longitude?
|   bits
+--rw longitude-degrees?
|   uint16
+--rw longitude-minutes?
|   uint8
+--rw longitude-seconds?
|   uint8
+--rw altitude?
|   int32
+--rw address?
|   simple-address
+--:(nat-traversal)
+--rw nat-traversal
+--rw ms-udp-port?
|   uint16
+--rw etr-udp-port?
|   uint16
+--rw global-etr-rloc?
|   simple-address
+--rw ms-rloc?
|   simple-address
+--rw private-etr-rloc?
|   simple-address
+--rw rtr-rlocs*
|   simple-address
+--:(explicit-locator-path)
+--rw explicit-locator-path
+--rw hop* [hop-id]
|   +--rw hop-id
|   |   string
|   +--rw address?
|   |   simple-address
+--rw lrs-bits?
|   bits
+--:(source-dest-key)
+--rw source-dest-key
+--rw source?
|   simple-address
+--rw dest?
|   simple-address
+--:(key-value-address)
+--rw key-value-address
+--rw key?
|   simple-address
+--rw value?

```



```
// RFC Ed.: replace occurrences of XXXX with actual RFC number
// and remove this note
import ietf-lisp {
  prefix lisp;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-lisp-address-types {
  prefix lcaf;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-yang-types {
  prefix yang;
  reference "RFC 6991: Common YANG Data Types";
}
import ietf-routing {
  prefix "rt";
  reference
    "RFC 8349: A YANG Data Model for Routing Management
    (NMDA version)";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web: <http://tools.ietf.org/wg/lisp/>
  WG List: <mailto:lisp@ietf.org>

  Editor: Vina Ermagan
  <mailto:vermagan@cisco.com>

  Editor: Alberto Rodriguez-Natal
  <mailto:natal@cisco.com>

  Editor: Reshad Rahman
  <mailto:rrahman@cisco.com>";
description
  "This YANG module defines the generic parameters for a LISP
  Map-Server. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
```

(<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.

```
";

reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6833";
}

identity ms {
  base lisp:lisp-role;
  description
    "LISP Map-Server.";
}

grouping ms-counters {
  description "Group that defines map-server counters.";
  container counters {
    config false;
    description "Container for the counters";

    leaf map-registers-in {
      type yang:counter64;
      description "Number of incoming Map-Register messages";
    }

    leaf map-registers-in-auth-failed {
      type yang:counter64;
      description
        "Number of incoming Map-Register messages failed
        authentication";
    }

    leaf map-notify-records-out {
      type yang:counter64;
      description
        "Number of outgoing Map-Notify records";
    }

    leaf proxy-reply-records-out {
      type yang:counter64;
      description

```

```
        "Number of outgoing proxy Map-Reply records";
    }

    leaf map-requests-forwarded-out {
        type yang:counter64;
        description
            "Number of outgoing Map-Requests forwarded to ETR";
    }
}

augment "/rt:routing/rt:control-plane-protocols"
+ "/rt:control-plane-protocol/lisp:lisp" {
when "lisp:lisp-role/lisp:lisp-role-type = 'lisp-ms:ms'" {
    description
        "Augment is valid when LISP device type is Map-Server.";
}
description
    "This augments LISP devices list with Map-Server specific
    parameters.";
container map-server {
    presence "LISP Map-Server operation enabled";
    description
        "Map-Server parameters.";
    container sites{
        description
            "Sites to accept registrations from.";
        list site {
            key site-id;
            description
                "Site that can send registrations.";
            leaf site-id {
                type uint64;
                description "Site ID";
            }
            container auth-key {
                description
                    "Site authentication key.";
                leaf auth-key-value {
                    type string;
                    description
                        "Clear text authentication key";
                }
            }
            leaf-list auth-key-type {
                type lisp:auth-key-type;
                description
                    "Authentication key type.";
            }
        }
    }
}
```

```
    }
  }
}
container virtual-network-ids {
  description
    "Sites for which the Map-Server accepts registrations.";
  list virtual-network-identifier {
    key "vni";
    description
      "Virtual network instances in the Map-Server.";
    leaf vni {
      type lcaf:instance-id-type;
      description
        "Virtual network identifier.";
    }
  }
  container mappings {
    description
      "EIDs registered by device.";
    list mapping {
      key "eid-id";
      description
        "List of EIDs registered by device.";
      leaf eid-id {
        type lisp:eid-id;
        description
          "Id of the EID registered.";
      }
    }
    container eid-address {
      uses lcaf:lisp-address;
      description
        "EID in generic LISP address format registered
        with the Map-Server.";
    }
    leaf-list site-id {
      type uint64;
      description "Site ID";
    }
    leaf more-specifics-accepted {
      type boolean;
      default "false";
      description
        "Flag indicating if more specific prefixes
        can be registered.";
    }
    leaf mapping-expiration-timeout {
      type int16;
      units "seconds";
      default "180"; //3 times the mapregister int
```

```
        description
            "Time before mapping is expired if no new
            registrations are received.";
    }
    leaf first-registration-time {
        type yang:date-and-time;
        config false;
        description
            "Time at which the first registration for this EID
            was received";
    }
    leaf last-registration-time {
        type yang:date-and-time;
        config false;
        description
            "Time at which the last registration for this EID
            was received";
    }
    container mapping-records {
        description
            "Datastore of registered mappings.";
        list mapping-record {
            key xtr-id;
            description
                "Registered mapping.";
            leaf xtr-id {
                type lisp:xtr-id-type;
                description "xTR ID";
            }
            leaf site-id {
                type uint64;
                description "Site ID";
            }
            uses lisp:mapping;
        }
    }
}
uses ms-counters;
}
leaf mapping-system-type {
    type lisp:mapping-system-ref;
    description
        "A reference to the mapping system";
}

container summary {
```

```
    config false;
    description "Summary state information";

    leaf number-configured-sites {
        type uint32;
        description "Number of configured LISP sites";
    }
    leaf number-registered-sites {
        type uint32;
        description "Number of registered LISP sites";
    }
    container af-datum {
        description "Number of configured EIDs per each AF";

        list af-data {
            key "address-type";
            description "Number of configured EIDs for this AF";
            leaf address-type {
                type lcaf:lisp-address-family-ref;
                description "AF type";
            }
            leaf number-configured-eids {
                type uint32;
                description "Number of configured EIDs for this AF";
            }
            leaf number-registered-eids {
                type uint32;
                description "Number of registered EIDs for this AF";
            }
        }
    }
    uses ms-counters;
}
}
<CODE ENDS>
```

6. LISP-Map-Resolver Module

This module captures the configuration data model of a LISP Map Resolver [RFC6833]. The model also captures some operational data elements.

6.1. Module Structure

```
module: ietf-lisp-mapresolver
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol/lisp:lisp:
    +--rw map-resolver!
      +--rw mapping-system-type?  lisp:mapping-system-ref
      +--rw ms-address?           inet:ip-address
```

6.2. Module Definition

```
<CODE BEGINS> file "ietf-lisp-mapresolver@2018-06-29.yang"
module ietf-lisp-mapresolver {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-mapresolver";

  prefix lisp-mr;

  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note
  import ietf-lisp {
    prefix lisp;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-inet-types {
    prefix inet;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
      (NMDA version)";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web:  <http://tools.ietf.org/wg/lisp/>
    WG List:  <mailto:lisp@ietf.org>

    Editor:   Vina Ermagan
              <mailto:vermagan@cisco.com>

    Editor:   Alberto Rodriguez-Natal
              <mailto:natal@cisco.com>
```

```
    Editor:   Reshad Rahman
             <mailto:rrahman@cisco.com>;
description
  "This YANG module defines the generic parameters for a LISP
  Map-Resolver. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.
  ";

reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6833";
}
identity mr {
  base lisp:lisp-role;
  description
    "LISP Map-Resolver.";
}

augment "/rt:routing/rt:control-plane-protocols"
  + "/rt:control-plane-protocol/lisp:lisp" {
  when "lisp:lisp-role/lisp:lisp-role-type = 'lisp-mr:mr'" {
    description
      "Augment is valid when LISP device type is Map-Resolver.";
  }
  description
    "This augments LISP devices list with Map-Resolver specific
    parameters.";
  container map-resolver {
    presence "LISP Map-Resolver operation enabled";
    description
      "Map-Resolver parameters.";
  }
}
```

```
    leaf mapping-system-type {
      type lisp:mapping-system-ref;
      description
        "A reference to the mapping system";
    }
    leaf ms-address {
      when "../mapping-system-type='lisp:single-node-mapping-system'";
      type inet:ip-address;
      description
        "address to reach the Map Server when "
        + "lisp-mr:single-node-mapping-system is being used.";
    }
  }
}
<CODE ENDS>
```

7. LISP-Address-Types Module

This module captures the various LISP address types, and is an essential building block used in other LISP modules.

7.1. Module Definition

```
<CODE BEGINS> file "ietf-lisp-address-types@2018-06-29.yang"
module ietf-lisp-address-types {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types";

  prefix laddr;

  import ietf-inet-types {
    prefix inet;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-yang-types {
    prefix yang;
    reference "RFC 6991: Common YANG Data Types";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web: <http://tools.ietf.org/wg/lisp/>
    WG List: <mailto:lisp@ietf.org>

    Editor: Vina Ermagan
```

<mailto:vermagan@cisco.com>

Editor: Alberto Rodriguez-Natal
<mailto:natal@cisco.com>

Editor: Reshad Rahman
<mailto:rrahman@cisco.com>;

description

"This YANG module defines the LISP Canonical Address Formats (LCAF) for LISP. The module can be extended by vendors to define vendor-specific parameters.

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.

```
";
// RFC Ed.: replace XXXX with actual RFC number and remove
// this note
reference "RFC XXXX";

revision 2018-06-29 {
  description
    "Initial revision.";
  reference
    "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10";
}
identity lisp-address-family {
  description
    "Base identity from which identities describing LISP address
    families are derived.";
}
identity no-address-afi {
  base lisp-address-family;
  description
    "IANA Reserved.";
}
identity ipv4-afi {
  base lisp-address-family;
```

```
    description
      "IANA IPv4 address family.";
  }
  identity ipv4-prefix-afi {
    base lisp-address-family;
    description
      "IANA IPv4 address family prefix.";
  }
  identity ipv6-afi {
    base lisp-address-family;
    description
      "IANA IPv6 address family.";
  }
  identity ipv6-prefix-afi {
    base lisp-address-family;
    description
      "IANA IPv6 address family prefix.";
  }
  identity mac-afi {
    base lisp-address-family;
    description
      "IANA MAC address family.";
  }
  identity distinguished-name-afi {
    base lisp-address-family;
    description
      "IANA Distinguished Name address family.";
  }
  identity as-number-afi {
    base lisp-address-family;
    description
      "IANA AS Number address family.";
  }
  identity lcaf {
    base lisp-address-family;
    description
      "IANA LISP Canonical Address Format address family.";
  }
  identity null-address-lcaf {
    base lcaf;
    description
      "Null body LCAF type.";
  }
  identity afi-list-lcaf {
    base lcaf;
    description
      "AFI-List LCAF type.";
  }
}
```

```
identity instance-id-lcaf {
  base lcaf;
  description
    "Instance-ID LCAF type.";
}
identity as-number-lcaf {
  base lcaf;
  description
    "AS Number LCAF type.";
}
identity application-data-lcaf {
  base lcaf;
  description
    "Application Data LCAF type.";
}
identity geo-coordinates-lcaf {
  base lcaf;
  description
    "Geo-coordinates LCAF type.";
}
identity opaque-key-lcaf {
  base lcaf;
  description
    "Opaque Key LCAF type.";
}
identity nat-traversal-lcaf {
  base lcaf;
  description
    "NAT-Traversal LCAF type.";
}
identity nonce-locator-lcaf {
  base lcaf;
  description
    "Nonce-Locator LCAF type.";
}
identity multicast-info-lcaf {
  base lcaf;
  description
    "Multicast Info LCAF type.";
}
identity explicit-locator-path-lcaf {
  base lcaf;
  description
    "Explicit Locator Path LCAF type.";
}
identity security-key-lcaf {
  base lcaf;
  description
```

```
        "Security Key LCAF type.";
    }
    identity source-dest-key-lcaf {
        base lcaf;
        description
            "Source/Dest LCAF type.";
    }
    identity replication-list-lcaf {
        base lcaf;
        description
            "Replication-List LCAF type.";
    }
    identity json-data-model-lcaf {
        base lcaf;
        description
            "JSON Data Model LCAF type.";
    }
    identity key-value-address-lcaf {
        base lcaf;
        description
            "Key/Value Address LCAF type.";
    }
    identity encapsulation-format-lcaf {
        base lcaf;
        description
            "Encapsulation Format LCAF type.";
    }
    identity service-path-lcaf {
        base lcaf;
        description
            "Service Path LCAF type.";
    }
    typedef instance-id-type {
        type uint32 {
            range "0..16777215";
        }
        description
            "Defines the range of values for an Instance ID.";
    }
    typedef service-path-id-type {
        type uint32 {
            range "0..16777215";
        }
        description
            "Defines the range of values for a Service Path ID.";
    }
    typedef distinguished-name-type {
        type string;
```

```
    description
      "Distinguished Name address.";
    reference
      "http://www.iana.org/assignments/address-family-numbers/
      address-family-numbers.xhtml";
  }
  typedef simple-address {
    type union {
      type inet:ip-address;
      type inet:ip-prefix;
      type yang:mac-address;
      type distinguished-name-type;
      type inet:as-number;
    }
    description
      "Union of address types that can be part of LCAFs.";
  }

  typedef lisp-address-family-ref {
    type identityref {
      base lisp-address-family;
    }
    description
      "LISP address family reference.";
  }

  typedef lcaf-ref {
    type identityref {
      base lcaf;
    }
    description
      "LCAF types reference.";
  }

  grouping lisp-address {
    description
      "Generic LISP address.";
    leaf address-type {
      type lisp-address-family-ref;
      mandatory true;
      description
        "Type of the LISP address.";
    }
    leaf virtual-network-id {
      type instance-id-type;
      description
        "Virtual Network Identifier (instance-id) of the address.";
    }
    choice address {
```

```
description
  "Various LISP address types, including IP, MAC, and LCAF.";

leaf no-address {
  when "../address-type = 'laddr:no-address-afi'" {
    description
      "When AFI is 0.";
  }
  type empty;
  description
    "No address.";
}
leaf ipv4 {
  when "../address-type = 'laddr:ipv4-afi'" {
    description
      "When AFI is IPv4.";
  }
  type inet:ipv4-address;
  description
    "IPv4 address.";
}
leaf ipv4-prefix {
  when "../address-type = 'laddr:ipv4-prefix-afi'" {
    description
      "When AFI is IPv4.";
  }
  type inet:ipv4-prefix;
  description
    "IPv4 prefix.";
}
leaf ipv6 {
  when "../address-type = 'laddr:ipv6-afi'" {
    description
      "When AFI is IPv6.";
  }
  type inet:ipv6-address;
  description
    "IPv6 address.";
}
leaf ipv6-prefix {
  when "../address-type = 'laddr:ipv6-prefix-afi'" {
    description
      "When AFI is IPv6.";
  }
  type inet:ipv6-prefix;
  description
    "IPv6 address.";
}
```

```
leaf mac {
  when "../address-type = 'laddr:mac-afi'" {
    description
      "When AFI is MAC.";
  }
  type yang:mac-address;
  description
    "MAC address.";
}
leaf distinguished-name {
  when "../address-type = 'laddr:distinguished-name-afi'" {
    description
      "When AFI is distinguished-name.";
  }
  type distinguished-name-type;
  description
    "Distinguished Name address.";
}
leaf as-number {
  when "../address-type = 'laddr:as-number-afi'" {
    description
      "When AFI is as-number.";
  }
  type inet:as-number;
  description
    "AS Number.";
}
container null-address {
  when "../address-type = 'laddr:null-address-lcaf'" {
    description
      "When LCAF type is null.";
  }
  description
    "Null body LCAF type";
  leaf address {
    type empty;
    description
      "AFI address.";
  }
}
container afi-list {
  when "../address-type = 'laddr:afi-list-lcaf'" {
    description
      "When LCAF type is AFI-List.";
  }
  description
    "AFI-List LCAF type.";
  reference
```

```
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.16.1";
    leaf-list address-list {
        type simple-address;
        description
            "List of AFI addresses.";
    }
}
container instance-id {
    when "../address-type = 'laddr:instance-id-lcaf'" {
        description
            "When LCAF type is Instance-ID";
    }
    description
        "Instance ID LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.2";
    leaf iid {
        type instance-id-type;
        description
            "Instance ID value.";
    }
    leaf mask-length {
        type uint8;
        description
            "Mask length.";
    }
    leaf address {
        type simple-address;
        description
            "AFI address.";
    }
}
container as-number-lcaf {
    when "../address-type = 'laddr:as-number-lcaf'" {
        description
            "When LCAF type is AS-Number.";
    }
    description
        "AS Number LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.3";
    leaf as {
        type inet:as-number;
        description
            "AS number.";
    }
}
```

```
    }
    leaf address {
      type simple-address;
      description
        "AFI address.";
    }
  }
  container application-data {
    when "../address-type = 'laddr:application-data-lcaf'" {
      description
        "When LCAF type is Application Data.";
    }
    description
      "Application Data LCAF type.";
    reference
      "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
      #section-4.4";
    leaf address {
      type simple-address;
      description
        "AFI address.";
    }
    leaf protocol {
      type uint8;
      description
        "Protocol number.";
    }
    leaf ip-tos {
      type int32;
      description
        "Type of service field.";
    }
    leaf local-port-low {
      type inet:port-number;
      description
        "Low end of local port range.";
    }
    leaf local-port-high {
      type inet:port-number;
      description
        "High end of local port range.";
    }
    leaf remote-port-low {
      type inet:port-number;
      description
        "Low end of remote port range.";
    }
    leaf remote-port-high {
```

```
        type inet:port-number;
        description
            "High end of remote port range.";
    }
}
container geo-coordinates {
    when "../address-type = 'laddr:geo-coordinates-lcaf'" {
        description
            "When LCAF type is Geo-coordinates.";
    }
    description
        "Geo-coordinates LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.5";
    leaf latitude {
        type bits {
            bit N {
                description
                    "Latitude bit.";
            }
        }
        description
            "Bit that selects between North and South latitude.";
    }
    leaf latitude-degrees {
        type uint8 {
            range "0 .. 90";
        }
        description
            "Degrees of latitude.";
    }
    leaf latitude-minutes {
        type uint8 {
            range "0..59";
        }
        description
            "Minutes of latitude.";
    }
    leaf latitude-seconds {
        type uint8 {
            range "0..59";
        }
        description
            "Seconds of latitude.";
    }
    leaf longitude {
        type bits {
```

```
        bit E {
            description
                "Longitude bit.";
        }
    }
    description
        "Bit that selects between East and West longitude.";
}
leaf longitude-degrees {
    type uint16 {
        range "0 .. 180";
    }
    description
        "Degrees of longitude.";
}
leaf longitude-minutes {
    type uint8 {
        range "0..59";
    }
    description
        "Minutes of longitude.";
}
leaf longitude-seconds {
    type uint8 {
        range "0..59";
    }
    description
        "Seconds of longitude.";
}
leaf altitude {
    type int32;
    description
        "Height relative to sea level in meters.";
}
leaf address {
    type simple-address;
    description
        "AFI address.";
}
}
container nat-traversal {
    when "../address-type = 'laddr:nat-traversal-lcaf'" {
        description
            "When LCAF type is NAT-Traversal.";
    }
    description
        "NAT-Traversal LCAF type.";
    reference
```

```
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.6";
leaf ms-udp-port {
  type uint16;
  description
    "Map-Server UDP port (set to 4342).";
}
leaf etr-udp-port {
  type uint16;
  description
    "ETR UDP port.";
}
leaf global-etr-rloc {
  type simple-address;
  description
    "Global ETR RLOC address.";
}
leaf ms-rloc {
  type simple-address;
  description
    "Map-Server RLOC address.";
}
leaf private-etr-rloc {
  type simple-address;
  description
    "Private ETR RLOC address.";
}
leaf-list rtr-rlocs {
  type simple-address;
  description
    "List of RTR RLOC addresses.";
}
}
container explicit-locator-path {
  when "../address-type = 'laddr:explicit-locator-path-lcaf'" {
    description
      "When LCAF type type is Explicit Locator Path.";
  }
  description
    "Explicit Locator Path LCAF type.";
  reference
    "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
    #section-4.9";
  list hop {
    key "hop-id";
    ordered-by user;
    description
      "List of locator hops forming the explicit path.";
  }
}
```

```
leaf hop-id {
  type string {
    length "1..64";
  }
  description
    "Unique identifier for the hop.";
}
leaf address {
  type simple-address;
  description
    "AFI address.";
}
leaf lrs-bits {
  type bits{
    bit lookup {
      description
        "Lookup bit.";
    }
    bit rloc-probe {
      description
        "RLOC-probe bit.";
    }
    bit strict {
      description
        "Strict bit.";
    }
  }
  description
    "Flag bits per hop.";
}
}
container source-dest-key {
  when "../address-type = 'laddr:source-dest-key-lcaf'" {
    description
      "When LCAF type type is Source/Dest.";
  }
  description
    "Source/Dest LCAF type.";
  reference
    "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10#section-4.11";
  leaf source {
    type simple-address;
    description
      "Source address.";
  }
  leaf dest {
```

```
        type simple-address;
        description
            "Destination address.";
    }
}
container key-value-address {
    when "../address-type = 'laddr:key-value-address-lcaf'" {
        description
            "When LCAF type type is Key/Value Address.";
    }
    description
        "Key/Value Address LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.11";
    leaf key {
        type simple-address;
        description
            "Address as Key.";
    }
    leaf value {
        type simple-address;
        description
            "Address as Value.";
    }
}
container service-path {
    when "../address-type = 'laddr:service-path-lcaf'" {
        description
            "When LCAF type service path identifier.";
    }
    description
        "Service Path LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ermagan-lisp-nsh-00";
    leaf service-path-id {
        type service-path-id-type;
        description
            "Service path identifier for the path for NSH header";
    }
    leaf service-index {
        type uint8;
        description
            "Service path index for NSH header";
    }
}
}
```

```
}  
<CODE ENDS>
```

7.2. Data Model examples

This section presents some simple and illustrative examples on how to configure LISP.

7.2.1. LISP protocol instance

The following is an example configuration for a LISP protocol instance with the name "LISP1". There are also 2 VNIs configured.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <network-instances xmlns="urn:ietf:params:xml:ns:yang:ietf-network-instance">
    <network-instance>
      <name>VRF-BLUE</name>
      <vrf-root/>
      <enabled>true</enabled>
    </network-instance>
    <network-instance>
      <name>VRF-RED</name>
      <vrf-root/>
      <enabled>true</enabled>
    </network-instance>
  </network-instances>
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <virtual-networks>
            <virtual-network>
              <vni>1000</vni>
              <ni-name>VRF-BLUE</ni-name>
            </virtual-network>
            <virtual-network>
              <vni>2000</vni>
              <ni-name>VRF-RED</ni-name>
            </virtual-network>
          </virtual-networks>
        </lisp>
      </control-plane-protocol>
    </control-plane-protocols>
  </routing>
</config>
```

7.2.2. LISP ITR

The following is an example configuration for ITR functionality under "LISP1". There are 2 Map-Resolvers configured.

```

<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type>itr</lisp-role-type>
          </lisp-role>
          <itr xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp-itr">
            <map-resolvers>
              <map-resolver>2001:db8:203:0:113::1</map-resolver>
              <map-resolver>2001:db8:204:0:113::1</map-resolver>
            </map-resolvers>
          </itr>
        </lisp>
      </control-plane-protocol>
    </control-plane-protocols>
  </routing>
</config>

```

7.2.3. LISP ETR

The following is an example configuration for ETR functionality under "LISP1". There are 2 Map-Servers and 2 local EIDs configured.

```

<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type>etr</lisp-role-type>
          </lisp-role>
          <lisp-router-id>
            <site-id>1</site-id>
          </lisp-router-id>
        </lisp>
      </control-plane-protocol>
    </control-plane-protocols>
  </routing>
</config>

```

```
</lisp-router-id>
<etr xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp-etr">
  <map-servers>
    <map-server>
      <ms-address>2001:db8:203:0:113::1</ms-address>
      <auth-key>*Kye^$$1#gb91U04zpa!</auth-key>
      <auth-key-type>hmac-sha-256-128</auth-key-type>
    </map-server>
    <map-server>
      <ms-address>2001:db8:204:0:113::1</ms-address>
      <auth-key>*Kye^$$1#gb91U04zpa!</auth-key>
      <auth-key-type>hmac-sha-256-128</auth-key-type>
    </map-server>
  </map-servers>
  <local-eids>
    <virtual-network>
      <vni>1000</vni>
      <eids>
        <local-eid>
          <id>2001:db8:400:0:100::0</id>
          <eid-address>
            <address-type xmlns:laddr=
              "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
              laddr:ipv6-prefix-afi
            </address-type>
            <ipv6-prefix>2001:db8:400:0:100::/80</ipv6-prefix>
          </eid-address>
        </local-eid>
      </eids>
    </virtual-network>
    <virtual-network>
      <vni>2000</vni>
      <eids>
        <local-eid>
          <id>2001:db8:800:0:200::0</id>
          <eid-address>
            <address-type xmlns:laddr=
              "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
              laddr:ipv6-prefix-afi
            </address-type>
            <ipv6-prefix>2001:db8:800:0:200::/80</ipv6-prefix>
          </eid-address>
        </local-eid>
      </eids>
    </virtual-network>
  </local-eids>
</etr>
</lisp>
```

```

    </control-plane-protocol>
  </control-plane-protocols>
</routing>
</config>

```

7.2.4. LISP Map-Server

The following is an example configuration for Map-Server functionality under "LISP1". There are 2 mappings configured.

```

<config xmlns="http://tail-f.com/ns/config/1.0">
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type xmlns:lisp-ms=
              "urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver">
              lisp-ms:ms
            </lisp-role-type>
          </lisp-role>
          <map-server xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver">
            <sites>
              <site>
                <site-id>1</site-id>
                <auth-key>
                  <auth-key-value>*Kye^$$1#gb91U04zpa!</auth-key-value>
                  <auth-key-type>hmac-sha-256-128</auth-key-type>
                </auth-key>
              </site>
            </sites>
            <virtual-network-ids>
              <virtual-network-identifier>
                <vni>1000</vni>
                <mappings>
                  <mapping>
                    <eid-id>1</eid-id>
                    <eid-address>
                      <address-type xmlns:laddr=
                        "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
                        laddr:ipv6-prefix-afi

```

```

        </address-type>
        <ipv6-prefix>2001:db8:400:0:100::/80</ipv6-prefix>
    </eid-address>
</mapping>
</mappings>
</virtual-network-identifier>
<virtual-network-identifier>
    <vni>2000</vni>
    <mappings>
        <mapping>
            <eid-id>1</eid-id>
            <eid-address>
                <address-type xmlns:laddr=
                    "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
                    laddr:ipv6-prefix-afi
                </address-type>
                <ipv6-prefix>2001:db8:800:0:200::/80</ipv6-prefix>
            </eid-address>
        </mapping>
    </mappings>
</virtual-network-identifier>
</virtual-network-ids>
</map-server>
</lisp>
</control-plane-protocol>
</control-plane-protocols>
</routing>
</config>

```

8. Acknowledgments

The tree view and the YANG model shown in this document have been formatted with the 'pyang' tool.

9. IANA Considerations

The IANA is requested to assign a new namespace URI from the IETF XML registry.

This document registers the following namespace URIs in the IETF XML registry [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-lisp

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-itr

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-etr

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-mapresolver

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-address-types

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

10. Security Considerations

The YANG modules specified in this document define a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a pre-configured subset of all available NETCONF or RESTCONF protocol operations and content.

The security considerations of LISP control-plane [RFC6833] and LISP data-plane [RFC6830] as well as the LISP threat analysis [RFC7835] apply to this YANG model.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/  
lisp:lisp/
```

Access to the locator-sets node may modify which interfaces are used for data and/or control traffic as well as affect the load balancing of data-plane traffic. Access to the lisp-role node may prevent the device from perform its intended data-plane and/or control-plane operation. Access to the router-id node allows to modify the unique identifier of the device, which may result in disruption of its LISP control-plane operation. Access to the virtual-networks node may allow to redirect data-plane traffic to erroneous local or remote network instances.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-server
```

Access to the sites node can prevent authorized devices from registering mappings in the Map-Server and/or allow unauthorized devices to do so. Access to the virtual-network-ids node can result in corrupted mapping state that may propagate across the LISP network, potentially resulting in forwarding of data-plane traffic to arbitrary destinations and general disruption of the data-plane operation. Access to mapping-system-type and/or ddt-mapping-system nodes may prevent the device to connect to the Mapping System infrastructure and consequentially to attract Map-Request messages.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-resolver
```

Access to mapping-system-type, ms-address and/or ddt-mapping-system nodes may prevent the device to connect to the Mapping System infrastructure and forward Map-Request messages.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:itr
```

Access to the rloc-probing node can increase the control-plane overhead in the device or affect the capability of the device to detect failures on the underlay. Access to the itr-rlocs node may prevent the device from getting Map-Reply messages. Access to the map-resolvers node can prevent the device from sending its Map-Request messages to valid Map-Resolvers. Access to the proxy-etr nodes can affect the capability of the device to send data-plane traffic towards non-LISP destinations. Access to the map-cache node can result in forwarding of data-plane traffic to arbitrary destinations and general disruption of data-plane operation.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:etr
```

Access to the map-servers node can prevent the device from registering its local mappings into the Mapping System. Access to the local-eids node can disrupt data-plane operation on the device and/or result in the device registering corrupted mappings into the Mapping System.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/  
lisp:lisp
```

Access to the locator-sets node can expose the locators the device is using for its control and/or data operation. Access to the lisp-role node can disclose the LISP roles instantiated at the device which facilitates mounting attacks against the device. Access to the router-id node can expose the unique identifier of device which may allow a third party to track its control-plane operation and/or impersonate the device. Access to the virtual-networks node can leak the local mapping between LISP Instance IDs and local network instances.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-server
```

Access to the sites node can expose the credentials used to register mappings and allow unauthorized devices to do so. Access to the virtual-network-ids node can expose the mappings currently registered in the device, which has privacy implications. Access to the mapping-system-type node may reveal the Mapping System in use which can be used to mount attacks against the device and/or the Mapping System. Access to the summary and counters nodes may expose operational statistics of the device.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-resolver
```

Access to the mapping-system-type node may reveal the Mapping System in use which can be used to mount attacks against the device and/or the Mapping System. Access to the ms-address and/or ddt-mapping-system nodes can leak the information about the Mapping System infrastructure used by the device, which can be used to block communication and/or mount attacks against it.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:itr
```

Access to the rloc-probing node can expose if and how the device is using control-plane signaling to probe underlay locators. Access to the itr-rlocs node may disclose the addresses the device is using to receive Map-Reply messages. Access to the map-resolvers node can expose the Map-Resolvers used by the device, which can be used to mount attacks against the device and/or the Mapping System. Access to the proxy-etrns node can disclose the PЕТRs used by the device, which can be used to mount attacks against the device and/or PЕТRs. Access to the map-cache node can expose the mappings currently cached in the device, which has privacy implications.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:etr
```

Access to the map-servers node can expose the credentials used by the device to register mappings into the Mapping System allowing an unauthorized device to impersonate and register mappings on behalf the authorized device. Access to the local-eids node can expose the local EIDs currently being served by the device, which has privacy implications.

11. 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>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.
- [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>>.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, DOI 10.17487/RFC6833, January 2013, <<https://www.rfc-editor.org/info/rfc6833>>.

- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, DOI 10.17487/RFC6836, January 2013, <<https://www.rfc-editor.org/info/rfc6836>>.
- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/info/rfc7835>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [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>>.
- [RFC8111] Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "Locator/ID Separation Protocol Delegated Database Tree (LISP-DDT)", RFC 8111, DOI 10.17487/RFC8111, May 2017, <<https://www.rfc-editor.org/info/rfc8111>>.
- [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>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC8349] Lhotka, L., Lindem, A., and Y. Qu, "A YANG Data Model for Routing Management (NMDA Version)", RFC 8349, DOI 10.17487/RFC8349, March 2018, <<https://www.rfc-editor.org/info/rfc8349>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

Authors' Addresses

Vina Ermagan
Cisco Systems
San Jose, CA
USA

Email: vermagan@cisco.com

Alberto Rodriguez-Natal
Cisco Systems
San Jose, CA
USA

Email: natal@cisco.com

Florin Coras
Cisco Systems
San Jose, CA
USA

Email: fcoras@cisco.com

Carl Moberg
Cisco Systems
San Jose, CA
USA

Email: camoberg@cisco.com

Reshad Rahman
Cisco Systems
Canada

Email: rrahman@cisco.com

Albert Cabellos-Aparicio
Technical University of Catalonia
Barcelona
Spain

Email: acabello@ac.upc.edu

Fabio Maino
Cisco Systems
San Jose, CA
USA

Email: fmaino@cisco.com

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: 17 October 2024

V. Ermagan
Google
A. Rodriguez-Natal
F. Coras
Cisco
C. Moberg
Avassa
R. Rahman
Equinix
A. Cabellos-Aparicio
Technical University of Catalonia
F. Maino
Cisco
15 April 2024

LISP YANG Model
draft-ietf-lisp-yang-21

Abstract

This document describes a YANG data model to use with the Locator/ID Separation Protocol (LISP).

The YANG modules in this document conform to the Network Management Datastore Architecture (NMDA).

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 17 October 2024.

Copyright Notice

Copyright (c) 2024 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	3
1.1.	Requirements Language	3
1.2.	Tree Diagrams	3
1.3.	Prefixes	3
2.	LISP Module	4
2.1.	Module Structure	4
2.2.	Module Definition	7
3.	LISP-ITR Module	19
3.1.	Module Structure	19
3.2.	Module Definition	24
4.	LISP-ETR Module	28
4.1.	Module Structure	28
4.2.	Module Definition	30
5.	LISP-Map-Server Module	34
5.1.	Module Structure	34
5.2.	Module Definition	43
6.	LISP-Map-Resolver Module	49
6.1.	Module Structure	50
6.2.	Module Definition	50
7.	LISP-Address-Types Module	52
7.1.	Module Definition	52
7.2.	Data Model examples	67
7.2.1.	LISP protocol instance	67
7.2.2.	LISP ITR	69
7.2.3.	LISP ETR	69
7.2.4.	LISP Map-Server	72
8.	Acknowledgments	73
9.	IANA Considerations	73
10.	Security Considerations	76
11.	Normative References	79
	Authors' Addresses	82

1. Introduction

The Locator/ID Separation Protocol (LISP) defines several network elements subject to be configured. This document presents the YANG data models required for basic configuration of all major LISP [RFC6830] elements. The models also capture some essential operational data elements as well.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Tree Diagrams

This document uses the graphical representation of data models defined in [RFC8340].

1.3. Prefixes

The table below provides a summary of the prefixes used by this document.

Prefix	YANG module	Reference
lisp	ietf-lisp	Section 2
if	ietf-interfaces	[RFC8343]
lisp-at	ietf-lisp-address-types	Section 7
yang	ietf-yang-types	[RFC6991]
rt	ietf-routing	[RFC8349]
ni	ietf-network-instance	[RFC8529]
lisp-itr	ietf-lisp-itr	Section 3
inet	ietf-inet-types	[RFC6991]
lisp-etr	ietf-lisp-etr	Section 4
lisp-ms	ietf-lisp-mapserver	Section 5
lisp-mr	ietf-lisp-mapresolver	Section 6

Table 1: Prefixes and corresponding YANG modules

2. LISP Module

This is the base LISP module. It is further augmented by the LISP device role specific modules defined elsewhere in this document.

2.1. Module Structure

```

module: ietf-lisp
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol:
      +--rw lisp
        +--rw locator-sets
          +--rw locator-set* [locator-set-name]
            +--rw locator-set-name    string
            +--rw (locator-type)?
              +--:(local-interface)
                +--rw interface* [interface-ref]
                  +--rw interface-ref    if:interface-ref
                  +--rw priority?        uint8
                  +--rw weight?         uint8

```

```

|      +---rw multicast-priority?  uint8
|      +---rw multicast-weight?    uint8
+---:(general-locator)
  +---rw locator* [locator-id]
  +---rw locator-id                string
  +---rw locator-address
  |   +---rw address-type
  |   |       lisp-address-family-ref
  +---rw (address)?
  |   +---:(no-address)
  |   |   +---rw no-address?          empty
  +---:(ipv4)
  |   +---rw ipv4?
  |   |       inet:ipv4-address
  +---:(ipv4-prefix)
  |   +---rw ipv4-prefix?
  |   |       inet:ipv4-prefix
  +---:(ipv6)
  |   +---rw ipv6?
  |   |       inet:ipv6-address
  +---:(ipv6-prefix)
  |   +---rw ipv6-prefix?
  |   |       inet:ipv6-prefix
  +---:(mac)
  |   +---rw mac?
  |   |       yang:mac-address
  +---:(distinguished-name)
  |   +---rw distinguished-name?
  |   |       distinguished-name-type
  +---:(as-number)
  |   +---rw as-number?
  |   |       inet:as-number
  +---:(null-address)
  |   +---rw null-address
  |   |       +---rw address?  empty
  +---:(afi-list)
  |   +---rw afi-list
  |   |       +---rw address-list*
  |   |       |       simple-address
  +---:(instance-id)
  |   +---rw instance-id
  |   |       +---rw instance-id?
  |   |       |       instance-id-type
  |   +---rw mask-length?  uint8
  |   +---rw address?      simple-address
  +---:(as-number-lcaf)
  |   +---rw as-number-lcaf
  |   |       +---rw as?          inet:as-number

```

```

|         +---rw address?   simple-address
+---:(application-data)
|         +---rw application-data
|         |         +---rw address?
|         |         |         simple-address
|         |         +---rw protocol?           uint8
|         |         +---rw ip-tos?             int32
|         |         +---rw local-port-low?
|         |         |         inet:port-number
|         |         +---rw local-port-high?
|         |         |         inet:port-number
|         |         +---rw remote-port-low?
|         |         |         inet:port-number
|         |         +---rw remote-port-high?
|         |         |         inet:port-number
+---:(geo-coordinates)
|         +---rw geo-coordinates
|         |         +---rw latitude?           bits
|         |         +---rw latitude-degrees?  uint8
|         |         +---rw latitude-minutes?  uint8
|         |         +---rw latitude-seconds?  uint8
|         |         +---rw longitude?         bits
|         |         +---rw longitude-degrees? uint16
|         |         +---rw longitude-minutes? uint8
|         |         +---rw longitude-seconds? uint8
|         |         +---rw altitude?          int32
|         |         +---rw address?
|         |         |         simple-address
+---:(nat-traversal)
|         +---rw nat-traversal
|         |         +---rw ms-udp-port?       uint16
|         |         +---rw etr-udp-port?      uint16
|         |         +---rw global-etr-rloc?
|         |         |         simple-address
|         |         +---rw ms-rloc?
|         |         |         simple-address
|         |         +---rw private-etr-rloc?
|         |         |         simple-address
|         |         +---rw rtr-rlocs*
|         |         |         simple-address
+---:(explicit-locator-path)
|         +---rw explicit-locator-path
|         |         +---rw hop* [hop-id]
|         |         |         +---rw hop-id   string
|         |         |         +---rw address? simple-address
|         |         |         +---rw lrs-bits? bits
+---:(source-dest-key)
|         +---rw source-dest-key

```



```
}
import ietf-yang-types {
  prefix yang;
  reference "RFC 6991: Common YANG Data Types";
}
import ietf-routing {
  prefix "rt";
  reference
    "RFC 8349: A YANG Data Model for Routing Management
    (NMDA version)";
}
import ietf-network-instance {
  prefix "ni";
  reference
    "RFC 8529: YANG Model for Network Instances";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web: <http://tools.ietf.org/wg/lisp/>
  WG List: <mailto:lisp@ietf.org>

  Editor: Vina Ermagan
         <mailto:ermagan@gmail.com>

  Editor: Alberto Rodriguez-Natal
         <mailto:natal@cisco.com>

  Editor: Reshad Rahman
         <mailto:reshad@yahoo.com>;
description
  "This YANG module defines the generic parameters for LISP.
  The module can be extended by vendors to define vendor-specific
  LISP parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).
```

```
    ";

    reference "RFC XXXX";

    revision 2021-02-22 {
      description
        "Initial revision.";
      reference
        "https://tools.ietf.org/html/rfc6830";
    }

    /*
     * Identity definitions
     */
    identity lisp {
      base "rt:control-plane-protocol";
      description "LISP protocol.";
      reference
        "RFC 6830: The Locator/ID Separation Protocol (LISP).";
    }

    identity lisp-role {
      description
        "LISP router role.";
    }
    identity itr {
      base lisp-role;
      description
        "LISP ITR.";
    }
    identity pitr {
      base lisp-role;
      description
        "LISP PITR.";
    }
    identity etr {
      base lisp-role;
      description
        "LISP ETR.";
    }
    identity petr {
      base lisp-role;
      description
        "LISP PETR.";
    }

    identity mapping-system {
      description
```

```
    "Mapping System interface";
  }
  identity single-node-mapping-system {
    base mapping-system;
    description
      "logically singular Map Server";
  }

  identity map-reply-act {
    description
      "Defines the lisp map-cache ACT type";
    reference
      "https://www.iana.org/assignments/lisp-parameters"
      + "/lisp-parameters.xhtml#lisp-act-value";
  }
  identity no-action {
    base map-reply-act;
    description
      "Mapping is kept alive and no encapsulation
      occurs.";
  }
  identity natively-forward {
    base map-reply-act;
    description
      "Matching packets are not encapsulated or
      dropped but natively forwarded.";
  }
  identity send-map-request {
    base map-reply-act;
    description
      "Matching packets invoke Map-Requests.";
  }
  identity drop-no-reason {
    base map-reply-act;
    description
      "Matching packets are dropped.";
  }
  identity drop-policy-denied {
    base map-reply-act;
    description
      "Matching packets are dropped (due to policy).";
  }
  identity drop-auth-failure {
    base map-reply-act;
    description
      "Matching packets are dropped (due to authentication
      failure).";
  }
}
```

```
identity auth-algorithm {
  description
    "Base identity for the authentication mechanisms supported by
    LISP.";
  reference
    "https://www.iana.org/assignments/lisp-parameters"
    + "/lisp-parameters.xhtml#lisp-key-id-numbers";
}
identity no-auth-algorithm {
  base auth-algorithm;
  description
    "No authentication.";
}
identity hmac-sha-1-96-none {
  base auth-algorithm;
  description
    "MAC = HMAC-SHA-1-96 (RFC2404), KDF = none";
}
identity hmac-sha-256-128-none {
  base auth-algorithm;
  description
    "MAC = HMAC-SHA-256-128 (RFC4868), KDF = none";
}
identity hmac-sha-256-128-HKDF-SHA2562 {
  base auth-algorithm;
  description
    "MAC = HMAC-SHA-256-128, KDF = HKDF-SHA2562 (RFC4868)";
}

typedef mapping-system-ref {
  type identityref {
    base mapping-system;
  }
  description
    "Mapping System reference";
}

typedef lisp-role-ref {
  type identityref {
    base lisp-role;
  }
  description
    "LISP role reference";
}

typedef map-reply-action {
  type identityref {
    base map-reply-act;
  }
}
```

```
    description
      "Map-Reply action reference";
  }
  typedef eid-id {
    type string {
      pattern '[a-zA-Z0-9\-\_\.]*';
    }
    description
      "Type encoding of lisp-addresses to be generally used in EID
      keyed lists.";
  }
  typedef auth-algorithm-type {
    type identityref {
      base auth-algorithm;
    }
    description
      "Authentication algorithm reference";
  }
  typedef xtr-id-type {
    type binary {
      length "16";
    }
    description
      "128-bit xTR identifier.";
  }

  grouping locator-properties {
    description
      "Properties of a RLOC";
    leaf priority {
      type uint8;
      description
        "Locator priority.";
    }
    leaf weight {
      type uint8;
      description
        "Locator weight.";
    }
    leaf multicast-priority {
      type uint8;
      description
        "Locator's multicast priority";
    }
    leaf multicast-weight {
      type uint8;
      description
        "Locator's multicast weight";
    }
  }
}
```

```
    }
  }

  grouping locators-grouping {
    description
      "Grouping that defines a list of LISP locators.";
    list locator {
      key "locator-id";
      description
        "List of routing locators";
      leaf locator-id {
        type string {
          length "1..64";
          pattern '[a-zA-Z0-9\-\_\.]*';
        }
        description
          "Locator id";
      }
      container locator-address {
        uses lisp-at:lisp-address;
        description
          "The locator address provided in LISP canonical
          address format.";
      }
      uses locator-properties;
    }
  }

  grouping local-locators-grouping {
    description
      "Grouping that defines a list of LISP locators.";
    list interface {
      key "interface-ref";
      description
        "The address type of the locator";
      leaf interface-ref {
        type if:interface-ref;
        description
          "The name of the interface supporting the locator.";
      }
      uses locator-properties;
    }
  }

  grouping mapping {
    description
      "Grouping that defines a LISP mapping.";
```

```
container eid {
  uses lisp-at:lisp-address;
  description
    "End-host Identifier (EID) to be mapped to a list of
    locators";
}
leaf time-to-live {
  type uint32;
  units minutes;
  description
    "Mapping validity period in minutes (as per RF6830).";
}
leaf creation-time {
  type yang:date-and-time;
  config false;
  description
    "Time when the mapping was created.";
}
leaf authoritative {
  type bits {
    bit A {
      description
        "Authoritative bit.";
    }
  }
  description
    "Bit that indicates if mapping comes from an
    authoritative source.";
}
leaf static {
  type boolean;
  default "false";
  description
    "This leaf should be true if the mapping is static.";
}
choice locator-list {
  description
    "list of locartors are either negative, or positive.";
  case negative-mapping {
    leaf map-reply-action {
      type map-reply-action;
      description
        "Forwarding action for a negative mapping.";
    }
  }
  case positive-mapping {
    container rlocs {
      uses locators-grouping;
    }
  }
}
```

```
        description
            "List of locators for a positive mapping.";
    }
}
}

grouping mappings {
    description
        "Grouping that defines a list of LISP mappings.";
    list vpn {
        key "instance-id";
        description
            "VPN to which the mappings belong.";
        leaf instance-id {
            type leafref {
                path "/rt:routing/rt:control-plane-protocols"
                    + "/rt:control-plane-protocol/lisp:lisp"
                    + "/lisp:vpns/lisp:vpn"
                    + "/lisp:instance-id";
            }
            description
                "VPN identifier.";
        }
        container mappings {
            description
                "Mappings within the VPN.";
            list mapping {
                key "eid-id";
                description
                    "List of EID to RLOCs mappings.";
                leaf eid-id {
                    type eid-id;
                    description
                        "Id that uniquely identifies a mapping.";
                }
                uses mapping;
            }
        }
    }
}

grouping auth-key {
    description "Grouping that defines authentication keys.";
    container authentication-keys {
        description "Multiple authentication keys can be defined.";
        list authentication-key {
            key "auth-key-id";
        }
    }
}
```

```
    description
    "Authentication key parameters.";
    leaf auth-key-id {
      type string {
        pattern '[a-zA-Z0-9\-\_\.]*';
      }
      description
      "Identifier of the authentication key.";
    }
    leaf-list auth-algorithm-id {
      type lisp:auth-algorithm-type;
      description
      "Authentication algorithm used with the key.";
    }
    leaf auth-key-value {
      type string;
      description
      "Clear text authentication key.";
    }
  }
}

augment "/rt:routing/rt:control-plane-protocols"
+ "/rt:control-plane-protocol" {
  when "derived-from-or-self(rt:type, 'lisp:lisp')" {
    description
    "This augmentation is only valid for a control-plane protocol
    instance of LISP.";
  }
  description "LISP protocol ietf-routing module
  control-plane-protocol augmentation.";

  container lisp {
    description
    "Parameters for the LISP subsystem.";

    container locator-sets {
      description
      "Container that defines a named locator set which can be
      referenced elsewhere.";
      list locator-set {
        key "locator-set-name";
        description
        "Multiple locator sets can be defined.";
        leaf locator-set-name {
          type string {
            length "1..64";
          }
        }
      }
    }
  }
}
```

```
        pattern '[a-zA-Z0-9\-\_\.]*';
    }
    description
        "Locator set name";
}
choice locator-type {
    description
        "Locator sets can be based on local interfaces, or
        general locators.";
    case local-interface {
        uses local-locators-grouping;
        description
            "List of locators in this set based on local
            interfaces.";
    }
    case general-locator {
        uses locators-grouping;
        description
            "List of locators in this set based on
            lisp-address.";
    }
}
}
}

list lisp-role {
    key lisp-role-type;
    description
        "List of lisp device roles such as MS, MR, ITR,
        PITR, ETR or PETR.";
    leaf lisp-role-type {
        type lisp-role-ref;
        description
            "The type of LISP device - identity derived from the
            'lisp-device' base identity.";
    }
}

container lisp-router-id {
    when "../lisp-role/lisp-role-type = 'lisp:itr' or
        ../lisp-role/lisp-role-type = 'lisp:pitr' or
        ../lisp-role/lisp-role-type = 'lisp:etr' or
        ../lisp-role/lisp-role-type = 'lisp:petr'" {
        description "Only when ITR, PITR, ETR or PETR.";
    }
    description
        "Site-ID and xTR-ID of the device.";
    leaf site-id {
```


3. LISP-ITR Module

This module captures the configuration data model of a LISP ITR. The model also captures some operational data elements.

3.1. Module Structure

```

module: ietf-lisp-itr
augment /rt:routing/rt:control-plane-protocols
  /rt:control-plane-protocol/lisp:lisp:
  +--rw itr!
    +--rw rloc-probing!
      | +--rw interval?          uint16
      | +--rw retries?          uint8
      | +--rw retries-interval? uint16
    +--rw itr-rlocs?          leafref
    +--rw map-resolvers
      | +--rw map-resolver*     inet:ip-address
    +--rw proxy-etrs
      | +--rw proxy-etr-address* inet:ip-address
    +--rw map-cache
      +--ro size?             uint32
      +--ro limit?            uint32
      +--rw vpn* [instance-id]
        +--rw instance-id     leafref
        +--rw mappings
          +--rw mapping* [eid-id]
            +--rw eid-id          eid-id
            +--rw eid
              +--rw address-type
              |   lisp-address-family-ref
            +--rw (address)?
              +--:(no-address)
              | +--rw no-address?          empty
              +--:(ipv4)
              | +--rw ipv4?
              |   inet:ipv4-address
            +--:(ipv4-prefix)
              | +--rw ipv4-prefix?
              |   inet:ipv4-prefix
            +--:(ipv6)
              | +--rw ipv6?
              |   inet:ipv6-address
            +--:(ipv6-prefix)
              | +--rw ipv6-prefix?
              |   inet:ipv6-prefix
            +--:(mac)
              | +--rw mac?

```

```

|           yang:mac-address
+---: (distinguished-name)
|   +---rw distinguished-name?
|           distinguished-name-type
+---: (as-number)
|   +---rw as-number?
|           inet:as-number
+---: (null-address)
|   +---rw null-address
|           +---rw address?   empty
+---: (afi-list)
|   +---rw afi-list
|           +---rw address-list*   simple-address
+---: (instance-id)
|   +---rw instance-id
|           +---rw instance-id?   instance-id-type
|           +---rw mask-length?   uint8
|           +---rw address?       simple-address
+---: (as-number-lcaf)
|   +---rw as-number-lcaf
|           +---rw as?           inet:as-number
|           +---rw address?     simple-address
+---: (application-data)
|   +---rw application-data
|           +---rw address?
|               |           simple-address
|           +---rw protocol?   uint8
|           +---rw ip-tos?     int32
|           +---rw local-port-low?
|               |           inet:port-number
|           +---rw local-port-high?
|               |           inet:port-number
|           +---rw remote-port-low?
|               |           inet:port-number
|           +---rw remote-port-high?
|                   inet:port-number
+---: (geo-coordinates)
|   +---rw geo-coordinates
|           +---rw latitude?   bits
|           +---rw latitude-degrees?   uint8
|           +---rw latitude-minutes?   uint8
|           +---rw latitude-seconds?   uint8
|           +---rw longitude?   bits
|           +---rw longitude-degrees?   uint16
|           +---rw longitude-minutes?   uint8
|           +---rw longitude-seconds?   uint8
|           +---rw altitude?   int32
|           +---rw address?

```

```

|           simple-address
+--:(nat-traversal)
|   +--rw nat-traversal
|       +--rw ms-udp-port?          uint16
|       +--rw etr-udp-port?        uint16
|       +--rw global-etr-rloc?
|           |           simple-address
|       +--rw ms-rloc?
|           |           simple-address
|       +--rw private-etr-rloc?
|           |           simple-address
|       +--rw rtr-rlocs*
|           simple-address
+--:(explicit-locator-path)
|   +--rw explicit-locator-path
|       +--rw hop* [hop-id]
|           +--rw hop-id          string
|           +--rw address?        simple-address
|       +--rw lrs-bits?          bits
+--:(source-dest-key)
|   +--rw source-dest-key
|       +--rw source?            simple-address
|       +--rw dest?              simple-address
+--:(key-value-address)
|   +--rw key-value-address
|       +--rw key?                simple-address
|       +--rw value?              simple-address
+--:(service-path)
|   +--rw service-path
|       +--rw service-path-id?
|           |           service-path-id-type
|       +--rw service-index?      uint8
+--rw time-to-live?                uint32
+--ro creation-time?                yang:date-and-time
+--rw authoritative?                bits
+--rw static?                        boolean
+--rw (locator-list)?
|   +--:(negative-mapping)
|       |   +--rw map-reply-action?  map-reply-action
+--:(positive-mapping)
+--rw rlocs
|   +--rw locator* [locator-id]
|       +--rw locator-id            string
|       +--rw locator-address
|           |   +--rw address-type
|               |   lisp-address-family-ref
|           +--rw (address)?
|               +--:(no-address)

```

```

|   +--rw no-address?
|       empty
+---:(ipv4)
|   +--rw ipv4?
|       inet:ipv4-address
+---:(ipv4-prefix)
|   +--rw ipv4-prefix?
|       inet:ipv4-prefix
+---:(ipv6)
|   +--rw ipv6?
|       inet:ipv6-address
+---:(ipv6-prefix)
|   +--rw ipv6-prefix?
|       inet:ipv6-prefix
+---:(mac)
|   +--rw mac?
|       yang:mac-address
+---:(distinguished-name)
|   +--rw distinguished-name?
|       distinguished-name-type
+---:(as-number)
|   +--rw as-number?
|       inet:as-number
+---:(null-address)
|   +--rw null-address
|       +--rw address?    empty
+---:(afi-list)
|   +--rw afi-list
|       +--rw address-list*
|           simple-address
+---:(instance-id)
|   +--rw instance-id
|       +--rw instance-id?
|           |   instance-id-type
|       +--rw mask-length?  uint8
|       +--rw address?
|           simple-address
+---:(as-number-lcaf)
|   +--rw as-number-lcaf
|       +--rw as?
|           |   inet:as-number
|       +--rw address?
|           simple-address
+---:(application-data)
|   +--rw application-data
|       +--rw address?
|           |   simple-address
|       +--rw protocol?

```

```

|         uint8
+--rw ip-tos?
|         int32
+--rw local-port-low?
|         inet:port-number
+--rw local-port-high?
|         inet:port-number
+--rw remote-port-low?
|         inet:port-number
+--rw remote-port-high?
|         inet:port-number
+---:(geo-coordinates)
+--rw geo-coordinates
+--rw latitude?
|         bits
+--rw latitude-degrees?
|         uint8
+--rw latitude-minutes?
|         uint8
+--rw latitude-seconds?
|         uint8
+--rw longitude?
|         bits
+--rw longitude-degrees?
|         uint16
+--rw longitude-minutes?
|         uint8
+--rw longitude-seconds?
|         uint8
+--rw altitude?
|         int32
+--rw address?
|         simple-address
+---:(nat-traversal)
+--rw nat-traversal
+--rw ms-udp-port?
|         uint16
+--rw etr-udp-port?
|         uint16
+--rw global-etr-rloc?
|         simple-address
+--rw ms-rloc?
|         simple-address
+--rw private-etr-rloc?
|         simple-address
+--rw rtr-rlocs*
|         simple-address
+---:(explicit-locator-path)

```

```

+--rw explicit-locator-path
  +--rw hop* [hop-id]
    +--rw hop-id
      |   string
    +--rw address?
      |   simple-address
    +--rw lrs-bits?  bits
+--:(source-dest-key)
  +--rw source-dest-key
    +--rw source?
      |   simple-address
    +--rw dest?
      |   simple-address
+--:(key-value-address)
  +--rw key-value-address
    +--rw key?
      |   simple-address
    +--rw value?
      |   simple-address
+--:(service-path)
  +--rw service-path
    +--rw service-path-id?
      |   service-path-id-type
    +--rw service-index?
      |   uint8
+--rw priority?          uint8
+--rw weight?           uint8
+--rw multicast-priority? uint8
+--rw multicast-weight?  uint8

```

3.2. Module Definition

```

<CODE BEGINS> file "ietf-lisp-itr@2019-02-23.yang"
module ietf-lisp-itr {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-itr";

  prefix lisp-itr;

  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note
  import ietf-lisp {
    prefix lisp;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-inet-types {
    prefix inet;
  }

```

```
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
      (NMDA version)";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web: <http://tools.ietf.org/wg/lisp/>
    WG List: <mailto:lisp@ietf.org>

    Editor: Vina Ermagan
           <mailto:ermagan@gmail.com>

    Editor: Alberto Rodriguez-Natal
           <mailto:natal@cisco.com>

    Editor: Reshad Rahman
           <mailto:reshad@yahoo.com>";
  description
    "This YANG module defines the generic parameters for a LISP
    ITR. The module can be extended by vendors to define
    vendor-specific parameters and policies.

    Copyright (c) 2018 IETF Trust and the persons identified as
    authors of the code. All rights reserved.

    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject
    to the license terms contained in, the Simplified BSD License
    set forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (http://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX; see
    the RFC itself for full legal notices.
    ";

  reference "RFC XXXX";

  revision 2019-02-23 {
    description
      "Initial revision.";
    reference
```

```
    "https://tools.ietf.org/html/rfc6830";
  }
  augment "/rt:routing/rt:control-plane-protocols"
    + "/rt:control-plane-protocol/lisp:lisp" {
    when "lisp:lisp-role/lisp:lisp-role-type = 'lisp:itr' or
        lisp:lisp-role/lisp:lisp-role-type = 'lisp:pitr'" {
      description
        "Augment is valid when LISP role type is ITR or PITR.";
    }
    description
      "This augments the LISP devices list with (P)ITR specific
      parameters.";
    container itr {
      presence "LISP (P)ITR operation enabled";
      description
        "ITR parameters";
      container rloc-probing {
        presence "RLOC probing active";
        description
          "RLOC-probing parameters";
        leaf interval {
          type uint16;
          units "seconds";
          description
            "Interval in seconds for resending the probes";
        }
        leaf retries {
          type uint8;
          description
            "Number of retries for sending the probes";
        }
        leaf retries-interval {
          type uint16;
          units "seconds";
          description
            "Interval in seconds between retries when sending probes.
            The action taken if all retries fail to receive is
            implementation specific.";
        }
      }
    }
    leaf itr-rlocs {
      type leafref {
        path "/rt:routing/rt:control-plane-protocols"
          + "/rt:control-plane-protocol/lisp:lisp"
          + "/lisp:locator-sets/lisp:locator-set"
          + "/lisp:locator-set-name";
      }
      description

```


4. LISP-ETR Module

This module captures the configuration data model of a LISP ETR. The model also captures some operational data elements.

4.1. Module Structure

```

module: ietf-lisp-etr
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol/lisp:lisp:
      +---rw etr!
        +---rw map-servers
          |
          | +---rw map-server* [ms-address]
          |   +---rw ms-address          inet:ip-address
          |   +---rw authentication-keys
          |     +---rw authentication-key* [auth-key-id]
          |       +---rw auth-key-id      string
          |       +---rw auth-algorithm-id*
          |         | lisp:auth-algorithm-type
          |         +---rw auth-key-value? string
          +---rw local-eids
            +---rw vpn* [instance-id]
            +---rw instance-id leafref
            +---rw eids
              +---rw local-eid* [eid-id]
              +---rw eid-id          lisp:eid-id
              +---rw eid-address
                +---rw address-type
                |   lisp-address-family-ref
                +---rw (address)?
                +---:(no-address)
                |   +---rw no-address?          empty
                +---:(ipv4)
                |   +---rw ipv4?
                |     inet:ipv4-address
                +---:(ipv4-prefix)
                |   +---rw ipv4-prefix?
                |     inet:ipv4-prefix
                +---:(ipv6)
                |   +---rw ipv6?
                |     inet:ipv6-address
                +---:(ipv6-prefix)
                |   +---rw ipv6-prefix?
                |     inet:ipv6-prefix
                +---:(mac)
                |   +---rw mac?
                |     yang:mac-address
                +---:(distinguished-name)

```

```

| +--rw distinguished-name?
|     distinguished-name-type
+--:(as-number)
| +--rw as-number?
|     inet:as-number
+--:(null-address)
| +--rw null-address
|     +--rw address?    empty
+--:(afi-list)
| +--rw afi-list
|     +--rw address-list*  simple-address
+--:(instance-id)
| +--rw instance-id
|     +--rw instance-id?  instance-id-type
|     +--rw mask-length?  uint8
|     +--rw address?      simple-address
+--:(as-number-lcaf)
| +--rw as-number-lcaf
|     +--rw as?          inet:as-number
|     +--rw address?    simple-address
+--:(application-data)
| +--rw application-data
|     +--rw address?
|         | simple-address
|     +--rw protocol?          uint8
|     +--rw ip-tos?            int32
|     +--rw local-port-low?
|         | inet:port-number
|     +--rw local-port-high?
|         | inet:port-number
|     +--rw remote-port-low?
|         | inet:port-number
|     +--rw remote-port-high?
|         | inet:port-number
+--:(geo-coordinates)
| +--rw geo-coordinates
|     +--rw latitude?          bits
|     +--rw latitude-degrees?  uint8
|     +--rw latitude-minutes?  uint8
|     +--rw latitude-seconds?  uint8
|     +--rw longitude?         bits
|     +--rw longitude-degrees?  uint16
|     +--rw longitude-minutes?  uint8
|     +--rw longitude-seconds?  uint8
|     +--rw altitude?          int32
|     +--rw address?
|         | simple-address
+--:(nat-traversal)

```

```

+--rw nat-traversal
  +--rw ms-udp-port?          uint16
  +--rw etr-udp-port?        uint16
  +--rw global-etr-rloc?
  |   simple-address
  +--rw ms-rloc?
  |   simple-address
  +--rw private-etr-rloc?
  |   simple-address
  +--rw rtr-rlocs*
  |   simple-address
+--:(explicit-locator-path)
  +--rw explicit-locator-path
  |   +--rw hop* [hop-id]
  |   |   +--rw hop-id      string
  |   |   +--rw address?    simple-address
  |   |   +--rw lrs-bits?   bits
+--:(source-dest-key)
  +--rw source-dest-key
  |   +--rw source?         simple-address
  |   +--rw dest?          simple-address
+--:(key-value-address)
  +--rw key-value-address
  |   +--rw key?            simple-address
  |   +--rw value?         simple-address
+--:(service-path)
  +--rw service-path
  |   +--rw service-path-id?
  |   |   service-path-id-type
  |   +--rw service-index?  uint8
+--rw rlocs?                leafref
+--rw record-ttl?           uint32
+--rw want-map-notify?     boolean
+--rw proxy-reply?         boolean
+--rw registration-interval? uint16

```

4.2. Module Definition

```

<CODE BEGINS> file "ietf-lisp-etr@2021-02-22.yang"
module ietf-lisp-etr {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-etr";

  prefix lisp-etr;

  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note

```

```
import ietf-lisp {
  prefix lisp;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-lisp-address-types {
  prefix lisp-at;
  reference "RFC XXXX: LISP YANG model";
}
import ietf-inet-types {
  prefix inet;
  reference "RFC 6991: Common YANG Data Types";
}
import ietf-routing {
  prefix "rt";
  reference
    "RFC 8349: A YANG Data Model for Routing Management
    (NMDA version)";
}

organization
  "IETF LISP (Locator/ID Separation Protocol) Working Group";
contact
  "WG Web: <http://tools.ietf.org/wg/lisp/>
  WG List: <mailto:lisp@ietf.org>

  Editor: Vina Ermagan
         <mailto:ermagan@gmail.com>

  Editor: Alberto Rodriguez-Natal
         <mailto:natal@cisco.com>

  Editor: Reshad Rahman
         <mailto:reshad@yahoo.com>";
description
  "This YANG module defines the generic parameters for a LISP
  ETR. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).
```

```
    This version of this YANG module is part of RFC XXXX; see
    the RFC itself for full legal notices.
";

reference "RFC XXXX";

revision 2021-02-22 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6830";
}
augment "/rt:routing/rt:control-plane-protocols"
  + "/rt:control-plane-protocol/lisp:lisp" {
  when "lisp:lisp-role/lisp:lisp-role-type = 'lisp:etr' or
    lisp:lisp-role/lisp:lisp-role-type = 'lisp:petr'" {
    description
      "Augment is valid when LISP device type is (P)ETR.";
  }
  description
    "This augments the LISP devices list with (P)ETR specific
    parameters.";
  container etr {
    presence "LISP (P)ETR operation enabled";
    description
      "(P)ETR parameters.";

    container map-servers {
      when "../..//lisp:lisp-role/lisp:lisp-role-type='lisp:etr'" {
        description
          "Container exists only when LISP device type is ETR.";
      }
      description
        "Map-Servers that the ETR uses.";
      list map-server {
        key "ms-address";
        description
          "Each Map-Server within the list of Map-Servers.";
        leaf ms-address {
          type inet:ip-address;
          description
            "Map-Server address.";
        }
        uses lisp:auth-key;
      }
    }
  }

  container local-eids {
```

```
when "../../../lisp:lisp-role/lisp:lisp-role-type='lisp:etr'" {
  description
    "Container exists only when LISP device type is ETR.";
}
description
  "VPNs served by the ETR.";
list vpn {
  key "instance-id";
  description
    "VPN for local-EIDs.";
  leaf instance-id {
    type leafref {
      path "/rt:routing/rt:control-plane-protocols"
        + "/rt:control-plane-protocol/lisp:lisp"
        + "/lisp:vpns/lisp:vpn"
        + "/lisp:instance-id";
    }
    description
      "VPN identifier.";
  }
  container eids {
    description
      "EIDs served by the ETR.";
    list local-eid {
      key "eid-id";
      description
        "List of local EIDs.";
      leaf eid-id {
        type lisp:eid-id;
        description
          "Unique id of local EID.";
      }
    }
    container eid-address {
      uses lisp-at:lisp-address;
      description
        "EID address in generic LISP address format.";
    }
  }
  leaf rlocs {
    type leafref {
      path "/rt:routing/rt:control-plane-protocols"
        + "/rt:control-plane-protocol/lisp:lisp"
        + "/lisp:locator-sets/lisp:locator-set"
        + "/lisp:locator-set-name";
    }
    description
      "Locator set mapped to this local EID.";
  }
  leaf record-ttl {
```



```

module: ietf-lisp-mapserver
augment /rt:routing/rt:control-plane-protocols
  /rt:control-plane-protocol/lisp:lisp:
  +--rw map-server!
    +--rw sites
      +--rw site* [site-id]
        +--rw site-id          uint64
        +--rw authentication-keys
          | +--rw authentication-key* [auth-key-id]
          |   +--rw auth-key-id      string
          |   +--rw auth-algorithm-id*
          |     | lisp:auth-algorithm-type
          |   +--rw auth-key-value?  string
        +--rw xtr-ids* [xtr-id]
          +--rw xtr-id            uint64
          +--rw authentication-keys
            +--rw authentication-key* [auth-key-id]
              +--rw auth-key-id      string
              +--rw auth-algorithm-id*
                | lisp:auth-algorithm-type
              +--rw auth-key-value?  string
    +--rw vpns
      +--rw vpn* [instance-id]
        +--rw instance-id      lisp-at:instance-id-type
        +--rw mappings
          +--rw mapping* [eid-id]
            +--rw eid-id          lisp:eid-id
            +--rw eid-address
              +--rw address-type
                | lisp-address-family-ref
              +--rw (address)?
                +--:(no-address)
                | +--rw no-address?          empty
                +--:(ipv4)
                | +--rw ipv4?
                |   inet:ipv4-address
                +--:(ipv4-prefix)
                | +--rw ipv4-prefix?
                |   inet:ipv4-prefix
                +--:(ipv6)
                | +--rw ipv6?
                |   inet:ipv6-address
                +--:(ipv6-prefix)
                | +--rw ipv6-prefix?
                |   inet:ipv6-prefix
                +--:(mac)
                | +--rw mac?
                |   yang:mac-address

```

```

+--:(distinguished-name)
|   +--rw distinguished-name?
|       distinguished-name-type
+--:(as-number)
|   +--rw as-number?
|       inet:as-number
+--:(null-address)
|   +--rw null-address
|       +--rw address?    empty
+--:(afi-list)
|   +--rw afi-list
|       +--rw address-list*  simple-address
+--:(instance-id)
|   +--rw instance-id
|       +--rw instance-id?  instance-id-type
|       +--rw mask-length?  uint8
|       +--rw address?      simple-address
+--:(as-number-lcaf)
|   +--rw as-number-lcaf
|       +--rw as?           inet:as-number
|       +--rw address?     simple-address
+--:(application-data)
|   +--rw application-data
|       +--rw address?
|           |   simple-address
|       +--rw protocol?    uint8
|       +--rw ip-tos?      int32
|       +--rw local-port-low?
|           |   inet:port-number
|       +--rw local-port-high?
|           |   inet:port-number
|       +--rw remote-port-low?
|           |   inet:port-number
|       +--rw remote-port-high?
|           |   inet:port-number
+--:(geo-coordinates)
|   +--rw geo-coordinates
|       +--rw latitude?    bits
|       +--rw latitude-degrees?  uint8
|       +--rw latitude-minutes?  uint8
|       +--rw latitude-seconds?  uint8
|       +--rw longitude?    bits
|       +--rw longitude-degrees?  uint16
|       +--rw longitude-minutes?  uint8
|       +--rw longitude-seconds?  uint8
|       +--rw altitude?    int32
|       +--rw address?
|           |   simple-address

```

```

+---:(nat-traversal)
|   +---rw nat-traversal
|       +---rw ms-udp-port?          uint16
|       +---rw etr-udp-port?        uint16
|       +---rw global-etr-rloc?
|           |
|           |   simple-address
|       +---rw ms-rloc?
|           |
|           |   simple-address
|       +---rw private-etr-rloc?
|           |
|           |   simple-address
|       +---rw rtr-rlocs*
|           |
|           |   simple-address
+---:(explicit-locator-path)
|   +---rw explicit-locator-path
|       +---rw hop* [hop-id]
|           +---rw hop-id          string
|           +---rw address?        simple-address
|           +---rw lrs-bits?       bits
+---:(source-dest-key)
|   +---rw source-dest-key
|       +---rw source?             simple-address
|       +---rw dest?              simple-address
+---:(key-value-address)
|   +---rw key-value-address
|       +---rw key?                simple-address
|       +---rw value?             simple-address
+---:(service-path)
|   +---rw service-path
|       +---rw service-path-id?
|           |
|           |   service-path-id-type
|       +---rw service-index?      uint8
+---rw site-id*                    uint64
+---rw more-specifics-accepted?    boolean
+---rw mapping-expiration-timeout? int16
+---ro first-registration-time?
|   |
|   |   yang:date-and-time
+---ro last-registration-time?
|   |
|   |   yang:date-and-time
+---rw mapping-records
|   +---rw mapping-record* [xtr-id]
|       +---rw xtr-id
|           |
|           |   lisp:xtr-id-type
|       +---rw site-id?            uint64
+---rw eid
|   +---rw address-type
|       |
|       |   lisp-address-family-ref
|   +---rw (address)?
|       +---:(no-address)

```

```

| +--rw no-address?
|     empty
+--:(ipv4)
| +--rw ipv4?
|     inet:ipv4-address
+--:(ipv4-prefix)
| +--rw ipv4-prefix?
|     inet:ipv4-prefix
+--:(ipv6)
| +--rw ipv6?
|     inet:ipv6-address
+--:(ipv6-prefix)
| +--rw ipv6-prefix?
|     inet:ipv6-prefix
+--:(mac)
| +--rw mac?
|     yang:mac-address
+--:(distinguished-name)
| +--rw distinguished-name?
|     distinguished-name-type
+--:(as-number)
| +--rw as-number?
|     inet:as-number
+--:(null-address)
| +--rw null-address
|     +--rw address?    empty
+--:(afi-list)
| +--rw afi-list
|     +--rw address-list*
|         simple-address
+--:(instance-id)
| +--rw instance-id
|     +--rw instance-id?
|         | instance-id-type
|     +--rw mask-length?  uint8
|     +--rw address?
|         simple-address
+--:(as-number-lcaf)
| +--rw as-number-lcaf
|     +--rw as?          inet:as-number
|     +--rw address?    simple-address
+--:(application-data)
| +--rw application-data
|     +--rw address?
|         | simple-address
|     +--rw protocol?   uint8
|     +--rw ip-tos?     int32
|     +--rw local-port-low?

```

```

|         inet:port-number
+---rw local-port-high?
|         inet:port-number
+---rw remote-port-low?
|         inet:port-number
+---rw remote-port-high?
|         inet:port-number
+---:(geo-coordinates)
+---rw geo-coordinates
+---rw latitude?           bits
+---rw latitude-degrees?
|         uint8
+---rw latitude-minutes?
|         uint8
+---rw latitude-seconds?
|         uint8
+---rw longitude?         bits
+---rw longitude-degrees?
|         uint16
+---rw longitude-minutes?
|         uint8
+---rw longitude-seconds?
|         uint8
+---rw altitude?
|         int32
+---rw address?
|         simple-address
+---:(nat-traversal)
+---rw nat-traversal
+---rw ms-udp-port?
|         uint16
+---rw etr-udp-port?
|         uint16
+---rw global-etr-rloc?
|         simple-address
+---rw ms-rloc?
|         simple-address
+---rw private-etr-rloc?
|         simple-address
+---rw rtr-rlocs*
|         simple-address
+---:(explicit-locator-path)
+---rw explicit-locator-path
+---rw hop* [hop-id]
|         +---rw hop-id      string
|         +---rw address?
|         |         simple-address
+---rw lrs-bits?      bits

```

```

+--:(source-dest-key)
|   +--rw source-dest-key
|       +--rw source?    simple-address
|       +--rw dest?     simple-address
+--:(key-value-address)
|   +--rw key-value-address
|       +--rw key?      simple-address
|       +--rw value?   simple-address
+--:(service-path)
|   +--rw service-path
|       +--rw service-path-id?
|           | service-path-id-type
|       +--rw service-index?    uint8
+--rw time-to-live?            uint32
+--ro creation-time?
|   yang:date-and-time
+--rw authoritative?         bits
+--rw static?                boolean
+--rw (locator-list)?
+--:(negative-mapping)
|   +--rw map-reply-action?
|       | map-reply-action
+--:(positive-mapping)
+--rw rlocs
|   +--rw locator* [locator-id]
|       +--rw locator-id
|           | string
|       +--rw locator-address
|           +--rw address-type
|               | lisp-address-family-ref
|           +--rw (address)?
|               +--:(no-address)
|                   | +--rw no-address?
|                   | empty
|               +--:(ipv4)
|                   | +--rw ipv4?
|                   |     inet:ipv4-address
|               +--:(ipv4-prefix)
|                   | +--rw ipv4-prefix?
|                   |     inet:ipv4-prefix
|               +--:(ipv6)
|                   | +--rw ipv6?
|                   |     inet:ipv6-address
|               +--:(ipv6-prefix)
|                   | +--rw ipv6-prefix?
|                   |     inet:ipv6-prefix
|               +--:(mac)
|                   | +--rw mac?

```

```

|           yang:mac-address
+---:(distinguished-name)
|   +---rw distinguished-name?
|       distinguished-name-type
+---:(as-number)
|   +---rw as-number?
|       inet:as-number
+---:(null-address)
|   +---rw null-address
|       +---rw address?
|           empty
+---:(afi-list)
|   +---rw afi-list
|       +---rw address-list*
|           simple-address
+---:(instance-id)
|   +---rw instance-id
|       +---rw instance-id?
|           instance-id-type
|   +---rw mask-length?
|       uint8
|   +---rw address?
|       simple-address
+---:(as-number-lcaf)
|   +---rw as-number-lcaf
|       +---rw as?
|           inet:as-number
|   +---rw address?
|       simple-address
+---:(application-data)
|   +---rw application-data
|       +---rw address?
|           simple-address
|   +---rw protocol?
|       uint8
|   +---rw ip-tos?
|       int32
|   +---rw local-port-low?
|       inet:port-number
|   +---rw local-port-high?
|       inet:port-number
|   +---rw remote-port-low?
|       inet:port-number
|   +---rw remote-port-high?
|       inet:port-number
+---:(geo-coordinates)
|   +---rw geo-coordinates
|       +---rw latitude?

```

```

|         bits
+--rw latitude-degrees?
|         uint8
+--rw latitude-minutes?
|         uint8
+--rw latitude-seconds?
|         uint8
+--rw longitude?
|         bits
+--rw longitude-degrees?
|         uint16
+--rw longitude-minutes?
|         uint8
+--rw longitude-seconds?
|         uint8
+--rw altitude?
|         int32
+--rw address?
|         simple-address
+--:(nat-traversal)
+--rw nat-traversal
+--rw ms-udp-port?
|         uint16
+--rw etr-udp-port?
|         uint16
+--rw global-etr-rloc?
|         simple-address
+--rw ms-rloc?
|         simple-address
+--rw private-etr-rloc?
|         simple-address
+--rw rtr-rlocs*
|         simple-address
+--:(explicit-locator-path)
+--rw explicit-locator-path
+--rw hop* [hop-id]
|         hop-id
|         string
+--rw address?
|         simple-address
+--rw lrs-bits?
|         bits
+--:(source-dest-key)
+--rw source-dest-key
+--rw source?
|         simple-address
+--rw dest?
|         simple-address

```



```
<CODE BEGINS> file "ietf-lisp-mapserver@2021-02-22.yang"
module ietf-lisp-mapserver {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver";

  prefix lisp-ms;

  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note
  import ietf-lisp {
    prefix lisp;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-lisp-address-types {
    prefix lisp-at;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-yang-types {
    prefix yang;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
      (NMDA version)";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web: <http://tools.ietf.org/wg/lisp/>
    WG List: <mailto:lisp@ietf.org>

    Editor: Vina Ermagan
           <mailto:ermagan@gmail.com>

    Editor: Alberto Rodriguez-Natal
           <mailto:natal@cisco.com>

    Editor: Reshad Rahman
           <mailto:reshad@yahoo.com>";
  description
    "This YANG module defines the generic parameters for a LISP
    Map-Server. The module can be extended by vendors to define
    vendor-specific parameters and policies.
```

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.

```
";

reference "RFC XXXX";

revision 2021-02-22 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6833";
}

identity ms {
  base lisp:lisp-role;
  description
    "LISP Map-Server.";
}

grouping ms-counters {
  description "Grouping that defines map-server counters.";
  container counters {
    config false;
    description "Container for the counters";

    leaf map-registers-in {
      type yang:counter64;
      description "Number of incoming Map-Register messages";
    }

    leaf map-registers-in-auth-failed {
      type yang:counter64;
      description
        "Number of incoming Map-Register messages failed
        authentication";
    }

    leaf map-notify-records-out {
```

```
        type yang:counter64;
        description
            "Number of outgoing Map-Notify records";
    }

    leaf proxy-reply-records-out {
        type yang:counter64;
        description
            "Number of outgoing proxy Map-Reply records";
    }

    leaf map-requests-forwarded-out {
        type yang:counter64;
        description
            "Number of outgoing Map-Requests forwarded to ETR";
    }
}

augment "/rt:routing/rt:control-plane-protocols"
+ "/rt:control-plane-protocol/lisp:lisp" {
    when "lisp:lisp-role/lisp:lisp-role-type = 'lisp-ms:ms'" {
        description
            "Augment is valid when LISP device type is Map-Server.";
    }
    description
        "This augments the LISP devices list with Map-Server
        specific parameters.";
    container map-server {
        presence "LISP Map-Server operation enabled";
        description
            "Map-Server parameters.";
        container sites{
            description
                "Sites to accept registrations from.";
            list site {
                key site-id;
                description
                    "Site that can send registrations.";
                leaf site-id {
                    type uint64;
                    description "Site ID";
                }
            }
            uses lisp:auth-key;
            list xtr-ids {
                key xtr-id;
                description "xTR-ID specific configuration.";
                leaf xtr-id {
```

```
        type uint64;
        description "xTR ID";
    }
    uses lisp:auth-key;
}
}
}
container vpns {
    description
        "VPNs for which the Map-Server accepts registrations.";
    list vpn {
        key "instance-id";
        description
            "VPN instances in the Map-Server.";
        leaf instance-id {
            type lisp-at:instance-id-type;
            description
                "VPN identifier.";
        }
        container mappings {
            description
                "EIDs registered by device.";
            list mapping {
                key "eid-id";
                description
                    "List of EIDs registered by device.";
                leaf eid-id {
                    type lisp:eid-id;
                    description
                        "Id of the EID registered.";
                }
            }
            container eid-address {
                uses lisp-at:lisp-address;
                description
                    "EID in generic LISP address format registered
                    with the Map-Server.";
            }
            leaf-list site-id {
                type uint64;
                description "Site ID";
            }
            leaf more-specifics-accepted {
                type boolean;
                default "false";
                description
                    "Flag indicating if more specific prefixes
                    can be registered.";
            }
        }
    }
}
```

```
leaf mapping-expiration-timeout {
  type int16;
  units "seconds";
  default "180"; //3 times the mapregister int
  description
    "Time before mapping is expired if no new
    registrations are received.";
}
leaf first-registration-time {
  type yang:date-and-time;
  config false;
  description
    "Time at which the first registration for this
    EID was received";
}
leaf last-registration-time {
  type yang:date-and-time;
  config false;
  description
    "Time at which the last registration for this EID
    was received";
}
container mapping-records {
  description
    "Datastore of registered mappings.";
  list mapping-record {
    key xtr-id;
    description
      "Registered mapping.";
    leaf xtr-id {
      type lisp:xtr-id-type;
      description "xTR ID";
    }
    leaf site-id {
      type uint64;
      description "Site ID";
    }
    uses lisp:mapping;
  }
}
}
}
uses ms-counters;
}
}
leaf mapping-system-type {
  type lisp:mapping-system-ref;
  description
```

```
        "A reference to the mapping system";
    }

    container summary {
        config false;
        description "Summary state information";

        leaf number-configured-sites {
            type uint32;
            description "Number of configured LISP sites";
        }
        leaf number-registered-sites {
            type uint32;
            description "Number of registered LISP sites";
        }
    }
    container af-datum {
        description "Number of configured EIDs per each AF";

        list af-data {
            key "address-type";
            description "Number of configured EIDs for this AF";
            leaf address-type {
                type lisp-at:lisp-address-family-ref;
                description "AF type";
            }
            leaf number-configured-eids {
                type uint32;
                description "Number of configured EIDs for this AF";
            }
            leaf number-registered-eids {
                type uint32;
                description "Number of registered EIDs for this AF";
            }
        }
    }
}
uses ms-counters;
}
}
<CODE ENDS>
```

6. LISP-Map-Resolver Module

This module captures the configuration data model of a LISP Map Resolver [RFC6833]. The model also captures some operational data elements.

6.1. Module Structure

```
module: ietf-lisp-mapresolver
  augment /rt:routing/rt:control-plane-protocols
    /rt:control-plane-protocol/lisp:lisp:
    +--rw map-resolver!
      +--rw mapping-system-type?   lisp:mapping-system-ref
      +--rw ms-address?            inet:ip-address
```

6.2. Module Definition

```
<CODE BEGINS> file "ietf-lisp-mapresolver@2019-02-23.yang"
module ietf-lisp-mapresolver {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-mapresolver";

  prefix lisp-mr;

  // RFC Ed.: replace occurrences of XXXX with actual RFC number
  // and remove this note
  import ietf-lisp {
    prefix lisp;
    reference "RFC XXXX: LISP YANG model";
  }
  import ietf-inet-types {
    prefix inet;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-routing {
    prefix "rt";
    reference
      "RFC 8349: A YANG Data Model for Routing Management
      (NMDA version)";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web:  <http://tools.ietf.org/wg/lisp/>
    WG List:  <mailto:lisp@ietf.org>

    Editor:   Vina Ermagan
              <mailto:ermagan@gmail.com>

    Editor:   Alberto Rodriguez-Natal
              <mailto:natal@cisco.com>
```

```
    Editor:   Reshad Rahman
             <mailto:reshad@yahoo.com>;
description
  "This YANG module defines the generic parameters for a LISP
  Map-Resolver. The module can be extended by vendors to define
  vendor-specific parameters and policies.

  Copyright (c) 2018 IETF Trust and the persons identified as
  authors of the code. All rights reserved.

  Redistribution and use in source and binary forms, with or
  without modification, is permitted pursuant to, and subject
  to the license terms contained in, the Simplified BSD License
  set forth in Section 4.c of the IETF Trust's Legal Provisions
  Relating to IETF Documents
  (http://trustee.ietf.org/license-info).

  This version of this YANG module is part of RFC XXXX; see
  the RFC itself for full legal notices.
";

reference "RFC XXXX";

revision 2019-02-23 {
  description
    "Initial revision.";
  reference
    "https://tools.ietf.org/html/rfc6833";
}
identity mr {
  base lisp:lisp-role;
  description
    "LISP Map-Resolver.";
}

augment "/rt:routing/rt:control-plane-protocols"
  + "/rt:control-plane-protocol/lisp:lisp" {
  when "lisp:lisp-role/lisp:lisp-role-type = 'lisp-mr:mr'" {
  description
    "Augment is valid when LISP device type is Map-Resolver.";
  }
  description
    "This augments the LISP devices list with Map-Resolver
    specific parameters.";
  container map-resolver {
  presence "LISP Map-Resolver operation enabled";
  description
    "Map-Resolver parameters.";
  }
}
```

```
    leaf mapping-system-type {
      type lisp:mapping-system-ref;
      description
        "A reference to the mapping system";
    }
    leaf ms-address {
      when "../mapping-system-type="
        + "'lisp:single-node-mapping-system'";
      type inet:ip-address;
      description
        "address to reach the Map Server when "
        + "lisp-mr:single-node-mapping-system is being used.";
    }
  }
}
}
}
<CODE ENDS>
```

7. LISP-Address-Types Module

This module captures the various LISP address types, and is an essential building block used in other LISP modules.

7.1. Module Definition

```
<CODE BEGINS> file "ietf-lisp-address-types@2021-02-22.yang"
module ietf-lisp-address-types {
  yang-version 1.1;

  namespace "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types";

  prefix lisp-at;

  import ietf-inet-types {
    prefix inet;
    reference "RFC 6991: Common YANG Data Types";
  }
  import ietf-yang-types {
    prefix yang;
    reference "RFC 6991: Common YANG Data Types";
  }

  organization
    "IETF LISP (Locator/ID Separation Protocol) Working Group";
  contact
    "WG Web: <http://tools.ietf.org/wg/lisp/>
    WG List: <mailto:lisp@ietf.org>
```

Editor: Vina Ermagan
<mailto:ermagan@gmail.com>

Editor: Alberto Rodriguez-Natal
<mailto:natal@cisco.com>

Editor: Reshad Rahman
<mailto:reshad@yahoo.com>";

description

"This YANG module defines the LISP Canonical Address Formats (LCAF) for LISP. The module can be extended by vendors to define vendor-specific parameters.

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.

```
";
// RFC Ed.: replace XXXX with actual RFC number and remove
// this note
reference "RFC XXXX";

revision 2021-02-22 {
  description
    "Initial revision.";
  reference
    "RC8060: LISP Canonical Address Format (LCAF)";
}
identity lisp-address-family {
  description
    "Base identity from which identities describing LISP address
    families are derived.";
}
identity no-address-afi {
  base lisp-address-family;
  description
    "IANA Reserved.";
}
identity ipv4-afi {
```

```
    base lisp-address-family;
    description
      "IANA IPv4 address family.";
  }
  identity ipv4-prefix-afi {
    base lisp-address-family;
    description
      "IANA IPv4 address family prefix.";
  }
  identity ipv6-afi {
    base lisp-address-family;
    description
      "IANA IPv6 address family.";
  }
  identity ipv6-prefix-afi {
    base lisp-address-family;
    description
      "IANA IPv6 address family prefix.";
  }
  identity mac-afi {
    base lisp-address-family;
    description
      "IANA MAC address family.";
  }
  identity distinguished-name-afi {
    base lisp-address-family;
    description
      "IANA Distinguished Name address family.";
  }
  identity as-number-afi {
    base lisp-address-family;
    description
      "IANA AS Number address family.";
  }
  identity lcaf {
    base lisp-address-family;
    description
      "IANA LISP Canonical Address Format address family.";
  }
  identity null-address-lcaf {
    base lcaf;
    description
      "Null body LCAF type.";
  }
  identity afi-list-lcaf {
    base lcaf;
    description
      "AFI-List LCAF type.";
```

```
    }
    identity instance-id-lcaf {
      base lcaf;
      description
        "Instance-ID LCAF type.";
    }
    identity as-number-lcaf {
      base lcaf;
      description
        "AS Number LCAF type.";
    }
    identity application-data-lcaf {
      base lcaf;
      description
        "Application Data LCAF type.";
    }
    identity geo-coordinates-lcaf {
      base lcaf;
      description
        "Geo-coordinates LCAF type.";
    }
    identity opaque-key-lcaf {
      base lcaf;
      description
        "Opaque Key LCAF type.";
    }
    identity nat-traversal-lcaf {
      base lcaf;
      description
        "NAT-Traversal LCAF type.";
    }
    identity nonce-locator-lcaf {
      base lcaf;
      description
        "Nonce-Locator LCAF type.";
    }
    identity multicast-info-lcaf {
      base lcaf;
      description
        "Multicast Info LCAF type.";
    }
    identity explicit-locator-path-lcaf {
      base lcaf;
      description
        "Explicit Locator Path LCAF type.";
    }
    identity security-key-lcaf {
      base lcaf;
```

```
    description
      "Security Key LCAF type.";
  }
  identity source-dest-key-lcaf {
    base lcaf;
    description
      "Source/Dest LCAF type.";
  }
  identity replication-list-lcaf {
    base lcaf;
    description
      "Replication-List LCAF type.";
  }
  identity json-data-model-lcaf {
    base lcaf;
    description
      "JSON Data Model LCAF type.";
  }
  identity key-value-address-lcaf {
    base lcaf;
    description
      "Key/Value Address LCAF type.";
  }
  identity encapsulation-format-lcaf {
    base lcaf;
    description
      "Encapsulation Format LCAF type.";
  }
  identity service-path-lcaf {
    base lcaf;
    description
      "Service Path LCAF type.";
  }
  typedef instance-id-type {
    type uint32 {
      range "0..16777215";
    }
    description
      "Defines the range of values for an Instance ID.";
  }
  typedef service-path-id-type {
    type uint32 {
      range "0..16777215";
    }
    description
      "Defines the range of values for a Service Path ID.";
  }
  typedef distinguished-name-type {
```

```
    type string;
    description
      "Distinguished Name address.";
    reference
      "http://www.iana.org/assignments/address-family-numbers/
      address-family-numbers.xhtml";
  }
  typedef simple-address {
    type union {
      type inet:ip-address;
      type inet:ip-prefix;
      type yang:mac-address;
      type distinguished-name-type;
      type inet:as-number;
    }
    description
      "Union of address types that can be part of LCAFs.";
  }
  typedef lisp-address-family-ref {
    type identityref {
      base lisp-address-family;
    }
    description
      "LISP address family reference.";
  }
  typedef lcaf-ref {
    type identityref {
      base lcaf;
    }
    description
      "LCAF types reference.";
  }

  grouping lisp-address {
    description
      "Generic LISP address.";
    leaf address-type {
      type lisp-address-family-ref;
      mandatory true;
      description
        "Type of the LISP address.";
    }
    choice address {
      description
        "Various LISP address types, including IP, MAC, and LCAF.";

      leaf no-address {
        when "../address-type = 'lisp-at:no-address-afi'" {

```

```
        description
            "When AFI is 0.";
    }
    type empty;
    description
        "No address.";
}
leaf ipv4 {
    when "../address-type = 'lisp-at:ipv4-afi'" {
        description
            "When AFI is IPv4.";
    }
    type inet:ipv4-address;
    description
        "IPv4 address.";
}
leaf ipv4-prefix {
    when "../address-type = 'lisp-at:ipv4-prefix-afi'" {
        description
            "When AFI is IPv4.";
    }
    type inet:ipv4-prefix;
    description
        "IPv4 prefix.";
}
leaf ipv6 {
    when "../address-type = 'lisp-at:ipv6-afi'" {
        description
            "When AFI is IPv6.";
    }
    type inet:ipv6-address;
    description
        "IPv6 address.";
}
leaf ipv6-prefix {
    when "../address-type = 'lisp-at:ipv6-prefix-afi'" {
        description
            "When AFI is IPv6.";
    }
    type inet:ipv6-prefix;
    description
        "IPv6 address.";
}
leaf mac {
    when "../address-type = 'lisp-at:mac-afi'" {
        description
            "When AFI is MAC.";
    }
}
```

```
    type yang:mac-address;
    description
      "MAC address.";
  }
  leaf distinguished-name {
    when "../address-type = 'lisp-at:distinguished-name-afi'" {
      description
        "When AFI is distinguished-name.";
    }
    type distinguished-name-type;
    description
      "Distinguished Name address.";
  }
  leaf as-number {
    when "../address-type = 'lisp-at:as-number-afi'" {
      description
        "When AFI is as-number.";
    }
    type inet:as-number;
    description
      "AS Number.";
  }
  container null-address {
    when "../address-type = 'lisp-at:null-address-lcaf'" {
      description
        "When LCAF type is null.";
    }
    description
      "Null body LCAF type";
    leaf address {
      type empty;
      description
        "AFI address.";
    }
  }
  container afi-list {
    when "../address-type = 'lisp-at:afi-list-lcaf'" {
      description
        "When LCAF type is AFI-List.";
    }
    description
      "AFI-List LCAF type.";
    reference
      "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
      #section-4.16.1";
    leaf-list address-list {
      type simple-address;
      description

```

```
        "List of AFI addresses.";
    }
}
container instance-id {
    when "../address-type = 'lisp-at:instance-id-lcaf'" {
        description
            "When LCAF type is Instance ID as per RFC8060.";
    }
    description
        "Instance ID LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.2";
    leaf instance-id {
        type instance-id-type;
        description
            "Instance ID value.";
    }
    leaf mask-length {
        type uint8;
        description
            "Mask length.";
    }
    leaf address {
        type simple-address;
        description
            "AFI address.";
    }
}
container as-number-lcaf {
    when "../address-type = 'lisp-at:as-number-lcaf'" {
        description
            "When LCAF type is AS-Number.";
    }
    description
        "AS Number LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.3";
    leaf as {
        type inet:as-number;
        description
            "AS number.";
    }
    leaf address {
        type simple-address;
        description
            "AFI address.";
    }
}
```

```
    }
  }
  container application-data {
    when "../address-type = 'lisp-at:application-data-lcaf'" {
      description
        "When LCAF type is Application Data.";
    }
    description
      "Application Data LCAF type.";
    reference
      "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
      #section-4.4";
    leaf address {
      type simple-address;
      description
        "AFI address.";
    }
    leaf protocol {
      type uint8;
      description
        "Protocol number.";
    }
    leaf ip-tos {
      type int32;
      description
        "Type of service field.";
    }
    leaf local-port-low {
      type inet:port-number;
      description
        "Low end of local port range.";
    }
    leaf local-port-high {
      type inet:port-number;
      description
        "High end of local port range.";
    }
    leaf remote-port-low {
      type inet:port-number;
      description
        "Low end of remote port range.";
    }
    leaf remote-port-high {
      type inet:port-number;
      description
        "High end of remote port range.";
    }
  }
}
```

```
container geo-coordinates {
  when "../address-type = 'lisp-at:geo-coordinates-lcaf'" {
    description
      "When LCAF type is Geo-coordinates.";
  }
  description
    "Geo-coordinates LCAF type. Coordinates are specified
    using the WGS 84 (World Geodetic System 1984) reference
    coordinate system";
  reference
    "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
    #section-4.5";
  leaf latitude {
    type bits {
      bit N {
        description
          "Latitude bit.";
      }
    }
    description
      "Bit that selects between North and South latitude.";
  }
  leaf latitude-degrees {
    type uint8 {
      range "0 .. 90";
    }
    description
      "Degrees of latitude.";
  }
  leaf latitude-minutes {
    type uint8 {
      range "0..59";
    }
    description
      "Minutes of latitude.";
  }
  leaf latitude-seconds {
    type uint8 {
      range "0..59";
    }
    description
      "Seconds of latitude.";
  }
  leaf longitude {
    type bits {
      bit E {
        description
          "Longitude bit.";
      }
    }
  }
}
```

```
    }
  }
  description
    "Bit that selects between East and West longitude.";
}
leaf longitude-degrees {
  type uint16 {
    range "0 .. 180";
  }
  description
    "Degrees of longitude.";
}
leaf longitude-minutes {
  type uint8 {
    range "0..59";
  }
  description
    "Minutes of longitude.";
}
leaf longitude-seconds {
  type uint8 {
    range "0..59";
  }
  description
    "Seconds of longitude.";
}
leaf altitude {
  type int32;
  description
    "Height relative to sea level in meters.";
}
leaf address {
  type simple-address;
  description
    "AFI address.";
}
}
container nat-traversal {
  when "../address-type = 'lisp-at:nat-traversal-lcaf'" {
    description
      "When LCAF type is NAT-Traversal.";
  }
  description
    "NAT-Traversal LCAF type.";
  reference
    "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
    #section-4.6";
  leaf ms-udp-port {
```

```
        type uint16;
        description
            "Map-Server UDP port (set to 4342).";
    }
    leaf etr-udp-port {
        type uint16;
        description
            "ETR UDP port.";
    }
    leaf global-etr-rloc {
        type simple-address;
        description
            "Global ETR RLOC address.";
    }
    leaf ms-rloc {
        type simple-address;
        description
            "Map-Server RLOC address.";
    }
    leaf private-etr-rloc {
        type simple-address;
        description
            "Private ETR RLOC address.";
    }
    leaf-list rtr-rlocs {
        type simple-address;
        description
            "List of RTR RLOC addresses.";
    }
}
container explicit-locator-path {
    when "../address-type = 'lisp-at:explicit-locator-path-lcaf'" {
        description
            "When LCAF type type is Explicit Locator Path.";
    }
    description
        "Explicit Locator Path LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10
        #section-4.9";
    list hop {
        key "hop-id";
        ordered-by user;
        description
            "List of locator hops forming the explicit path.";
        leaf hop-id {
            type string {
                length "1..64";
            }
        }
    }
}
```

```
        pattern '[a-zA-Z0-9\-\_\.]*';
    }
    description
        "Unique identifier for the hop.";
}
leaf address {
    type simple-address;
    description
        "AFI address.";
}
leaf lrs-bits {
    type bits{
        bit lookup {
            description
                "Lookup bit.";
        }
        bit rloc-probe {
            description
                "RLOC-probe bit.";
        }
        bit strict {
            description
                "Strict bit.";
        }
    }
    description
        "Flag bits per hop.";
}
}
}
container source-dest-key {
    when "../address-type = 'lisp-at:source-dest-key-lcaf'" {
        description
            "When LCAF type type is Source/Dest.";
    }
    description
        "Source/Dest LCAF type.";
    reference
        "http://tools.ietf.org/html/draft-ietf-lisp-lcaf-10#section-4.11";
    leaf source {
        type simple-address;
        description
            "Source address.";
    }
    leaf dest {
        type simple-address;
        description
```


7.2. Data Model examples

This section presents some simple and illustrative examples on how to configure LISP.

7.2.1. LISP protocol instance

The following is an example configuration for a LISP protocol instance with the name "LISP1". There are also 2 VNIs configured.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <network-instances
    xmlns="urn:ietf:params:xml:ns:yang:ietf-network-instance">
    <network-instance>
      <name>VRF-BLUE</name>
      <vrf-root/>
      <enabled>>true</enabled>
    </network-instance>
    <network-instance>
      <name>VRF-RED</name>
      <vrf-root/>
      <enabled>>true</enabled>
    </network-instance>
  </network-instances>
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type>etr</lisp-role-type>
          </lisp-role>
          <lisp-role>
            <lisp-role-type>itr</lisp-role-type>
          </lisp-role>
          <vpns>
            <vpn>
              <instance-id>1000</instance-id>
              <iid-name>VRF-BLUE</iid-name>
            </vpn>
            <vpn>
              <instance-id>2000</instance-id>
              <iid-name>VRF-RED</iid-name>
            </vpn>
          </vpns>
        </lisp>
      </control-plane-protocol>
    </control-plane-protocols>
  </routing>
</config>
```

7.2.2. LISP ITR

The following is an example configuration for ITR functionality under "LISP1". There are 2 Map-Resolvers configured.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type>itr</lisp-role-type>
          </lisp-role>
          <itr xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp-itr">
            <map-resolvers>
              <map-resolver>2001:db8:203:0:113::1</map-resolver>
              <map-resolver>2001:db8:204:0:113::1</map-resolver>
            </map-resolvers>
          </itr>
        </lisp>
      </control-plane-protocol>
    </control-plane-protocols>
  </routing>
</config>
```

7.2.3. LISP ETR

The following is an example configuration for ETR functionality under "LISP1". There are 2 Map-Servers and 2 local EIDs configured.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <network-instances
    xmlns="urn:ietf:params:xml:ns:yang:ietf-network-instance">
    <network-instance>
      <name>VRF-BLUE</name>
      <vrf-root/>
      <enabled>true</enabled>
    </network-instance>
```

```
<network-instance>
  <name>VRF-RED</name>
  <vrf-root/>
  <enabled>>true</enabled>
</network-instance>
</network-instances>
<routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
  <control-plane-protocols>
    <control-plane-protocol>
      <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
        lisp:lisp
      </type>
      <name>LISP1</name>
      <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
        <lisp-role>
          <lisp-role-type>etr</lisp-role-type>
        </lisp-role>
        <lisp-router-id>
          <site-id>1</site-id>
        </lisp-router-id>
        <vpns>
          <vpn>
            <instance-id>1000</instance-id>
            <iid-name>VRF-BLUE</iid-name>
          </vpn>
          <vpn>
            <instance-id>2000</instance-id>
            <iid-name>VRF-RED</iid-name>
          </vpn>
        </vpns>
        <etr xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp-etr">
          <map-servers>
            <map-server>
              <ms-address>2001:db8:203:0:113::1</ms-address>
              <authentication-keys>
                <authentication-key>
                  <auth-key-id>key1</auth-key-id>
                  <auth-algorithm-id>
                    hmac-sha-256-128
                  </auth-algorithm-id>
                  <auth-key-value>*Kye^$$1#gb91U04zpa</auth-key-value>
                </authentication-key>
              </authentication-keys>
            </map-server>
            <map-server>
              <ms-address>2001:db8:204:0:113::1</ms-address>
              <authentication-keys>
                <authentication-key>
```

```
        <auth-key-id>key1</auth-key-id>
        <auth-algorithm-id>
          hmac-sha-256-128
        </auth-algorithm-id>
        <auth-key-value>*Kye^$$1#gb91U04zpa</auth-key-value>
      </authentication-key>
    </authentication-keys>
  </map-server>
</map-servers>
<local-eids>
  <vpn>
    <instance-id>1000</instance-id>
    <eids>
      <local-eid>
        <id>2001:db8:400:0:100::0</id>
        <eid-address>
          <address-type xmlns:lisp-at=
            "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
            lisp-at:ipv6-prefix-afi
          </address-type>
          <ipv6-prefix>2001:db8:400:0:100::/80</ipv6-prefix>
        </eid-address>
      </local-eid>
    </eids>
  </vpn>
  <vpn>
    <instance-id>2000</instance-id>
    <eids>
      <local-eid>
        <id>2001:db8:800:0:200::0</id>
        <eid-address>
          <address-type xmlns:lisp-at=
            "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
            lisp-at:ipv6-prefix-afi
          </address-type>
          <ipv6-prefix>2001:db8:800:0:200::/80</ipv6-prefix>
        </eid-address>
      </local-eid>
    </eids>
  </vpn>
</local-eids>
</etr>
</lisp>
</control-plane-protocol>
</control-plane-protocols>
</routing>
</config>
```

7.2.4. LISP Map-Server

The following is an example configuration for Map-Server functionality under "LISP1". There are 2 mappings configured.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <routing xmlns="urn:ietf:params:xml:ns:yang:ietf-routing">
    <control-plane-protocols>
      <control-plane-protocol>
        <type xmlns:lisp="urn:ietf:params:xml:ns:yang:ietf-lisp">
          lisp:lisp
        </type>
        <name>LISP1</name>
        <lisp xmlns="urn:ietf:params:xml:ns:yang:ietf-lisp">
          <lisp-role>
            <lisp-role-type xmlns:lisp-ms=
              "urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver">
              lisp-ms:ms
            </lisp-role-type>
          </lisp-role>
          <map-server xmlns=
            "urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver">
            <sites>
              <site>
                <site-id>1</site-id>
                <authentication-keys>
                  <authentication-key>
                    <auth-key-id>key1</auth-key-id>
                    <auth-algorithm-id>
                      hmac-sha-256-128
                    </auth-algorithm-id>
                    <auth-key-value>*Kye^$$1#gb91U04zpa</auth-key-value>
                  </authentication-key>
                </authentication-keys>
              </site>
            </sites>
            <vpns>
              <vpn>
                <instance-id>1000</instance-id>
                <mappings>
                  <mapping>
                    <eid-id>1</eid-id>
                    <eid-address>
                      <address-type xmlns:lisp-at=
                        "urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
                        lisp-at:ipv6-prefix-afi
```

```

        </address-type>
        <ipv6-prefix>2001:db8:400:0:100::/80</ipv6-prefix>
    </eid-address>
  </mapping>
</mappings>
</vpn>
<vpn>
  <instance-id>2000</instance-id>
  <mappings>
    <mapping>
      <eid-id>1</eid-id>
      <eid-address>
        <address-type xmlns:lisp-at=
"urn:ietf:params:xml:ns:yang:ietf-lisp-address-types">
          lisp-at:ipv6-prefix-afi
        </address-type>
        <ipv6-prefix>2001:db8:800:0:200::/80</ipv6-prefix>
      </eid-address>
    </mapping>
  </mappings>
</vpn>
</vpns>
</map-server>
</lisp>
</control-plane-protocol>
</control-plane-protocols>
</routing>
</config>

```

8. Acknowledgments

The tree view and the YANG model shown in this document have been formatted with the 'pyang' tool.

9. IANA Considerations

The IANA is requested to assign a new namespace URI from the IETF XML registry.

This document registers the following namespace URIs in the IETF XML registry [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-lisp

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-itr

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-etr

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-mapresolver

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-lisp-address-types

Registrant Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

This document registers the following YANG modules in the "YANG Module Names" registry [RFC6020]:

Name: ietf-lisp

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp

Prefix: lisp

Reference: RFC XXX

Name: ietf-lisp-itr

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp-itr

Prefix: lisp-itr

Reference: RFC XXX

Name: ietf-lisp-etr

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp-etr

Prefix: lisp-etr

Reference: RFC XXX

Name: ietf-lisp-mapserver

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp-mapserver

Prefix: lisp-ms

Reference: RFC XXX

Name: ietf-lisp-mapresolver

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp-mapresolver

Prefix: lisp-mr

Reference: RFC XXX

Name: ietf-lisp-address-types

Namespace: urn:ietf:params:xml:ns:yang:ietf-lisp-address-types

Prefix: lisp-at

Reference: RFC XXX

10. Security Considerations

The YANG modules specified in this document define a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS

[RFC8446].

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a pre-configured subset of all available NETCONF or RESTCONF protocol operations and content.

The security considerations of LISP control-plane [RFC6833] and LISP data-plane [RFC6830] as well as the LISP threat analysis [RFC7835] apply to this YANG model.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/  
lisp:lisp/
```

Access to the locator-sets node may modify which interfaces are used for data and/or control traffic as well as affect the load balancing of data-plane traffic. Access to the lisp-role node may prevent the device from perform its intended data-plane and/or control-plane operation. Access to the router-id node allows to modify the unique identifier of the device, which may result in disruption of its LISP control-plane operation. Access to the vpn node may allow to redirect data-plane traffic to erroneous local or remote network instances.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-server
```

Access to the sites node can prevent authorized devices from registering mappings in the Map-Server and/or allow unauthorized devices to so. Access to the vpn node can result in corrupted mapping sate that may propagate across the LISP network, potentially resulting in forwarding of data-plane traffic to arbitrary destinations and general disruption of the data-plane operation. Access to mapping-system-type and/or ddt-mapping-system nodes may prevent the device to connect to the Mapping System infrastructure and consequentially to attract Map-Request messages.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-  
protocol/lisp:lisp/lisp:map-resolver
```

Access to mapping-system-type, ms-address and/or ddt-mapping-system nodes may prevent the device to connect to the Mapping System infrastructure and forward Map-Request messages.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:itr
```

Access to the rloc-probing node can increase the control-plane overhead in the device or affect the capability of the device to detect failures on the underlay. Access to the itr-rlocs node may prevent the device from getting Map-Reply messages. Access to the map-resolvers node can prevent the device from sending its Map-Request messages to valid Map-Resolvers. Access to the proxy-etr nodes can affect the capability of the device to send data-plane traffic towards non-LISP destinations. Access to the map-cache node can result in forwarding of data-plane traffic to arbitrary destinations and general disruption of data-plane operation.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:etr
```

Access to the map-servers node can prevent the device from registering its local mappings into the Mapping System. Access to the local-eids node can disrupt data-plane operation on the device and/or result in the device registering corrupted mappings into the Mapping System.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. These are the subtrees and data nodes and their sensitivity/vulnerability:

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp
```

Access to the locator-sets node can expose the locators the device is using for its control and/or data operation. Access to the lisp-role node can disclose the LISP roles instantiated at the device which facilitates mounting attacks against the device. Access to the router-id node can expose the unique identifier of device which may allow a third party to track its control-plane operation and/or impersonate the device. Access to the vpn node can leak the local mapping between LISP Instance IDs and local network instances.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:map-server
```

Access to the sites node can expose the credentials used to register mappings and allow unauthorized devices to do so. Access to the vpn node can expose the mappings currently registered in the device, which has privacy implications. Access to the mapping-system-type node may reveal the Mapping System in use which can be used to mount attacks against the device and/or the Mapping System. Access to the summary and counters nodes may expose operational statistics of the device.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:map-resolver
```

Access to the mapping-system-type node may reveal the Mapping System in use which can be used to mount attacks against the device and/or the Mapping System. Access to the ms-address and/or ddt-mapping-system nodes can leak the information about the Mapping System infrastructure used by the device, which can be used to block communication and/or mount attacks against it.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:itr
```

Access to the rloc-probing node can expose if and how the device is using control-plane signaling to probe underlay locators. Access to the itr-rlocs node may disclose the addresses the device is using to receive Map-Reply messages. Access to the map-resolvers node can expose the Map-Resolvers used by the device, which can be used to mount attacks against the device and/or the Mapping System. Access to the proxy-etrns node can disclose the PETRs used by the device, which can be used to mount attacks against the device and/or PETRs. Access to the map-cache node can expose the mappings currently cached in the device, which has privacy implications.

```
/rt:routing/rt:control-plane-protocols/rt:control-plane-protocol/lisp:lisp/lisp:etr
```

Access to the map-servers node can expose the credentials used by the device to register mappings into the Mapping System allowing an unauthorized device to impersonate and register mappings on behalf the authorized device. Access to the local-eids node can expose the local EIDs currently being served by the device, which has privacy implications.

11. 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>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.
- [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>>.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, DOI 10.17487/RFC6833, January 2013, <<https://www.rfc-editor.org/info/rfc6833>>.
- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, DOI 10.17487/RFC6836, January 2013, <<https://www.rfc-editor.org/info/rfc6836>>.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.

- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/info/rfc7835>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [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>>.
- [RFC8111] Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "Locator/ID Separation Protocol Delegated Database Tree (LISP-DDT)", RFC 8111, DOI 10.17487/RFC8111, May 2017, <<https://www.rfc-editor.org/info/rfc8111>>.
- [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>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 8343, DOI 10.17487/RFC8343, March 2018, <<https://www.rfc-editor.org/info/rfc8343>>.
- [RFC8349] Lhotka, L., Lindem, A., and Y. Qu, "A YANG Data Model for Routing Management (NMDA Version)", RFC 8349, DOI 10.17487/RFC8349, March 2018, <<https://www.rfc-editor.org/info/rfc8349>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8529] Berger, L., Hopps, C., Lindem, A., Bogdanovic, D., and X. Liu, "YANG Data Model for Network Instances", RFC 8529, DOI 10.17487/RFC8529, March 2019, <<https://www.rfc-editor.org/info/rfc8529>>.

Authors' Addresses

Vina Ermagan
Google
United States of America
Email: ermagan@gmail.com

Alberto Rodriguez-Natal
Cisco
San Jose, CA
United States of America
Email: natal@cisco.com

Florin Coras
Cisco
San Jose, CA
United States of America
Email: fcoras@cisco.com

Carl Moberg
Avassa
Email: calle@avassa.io

Reshad Rahman
Equinix
Canada
Email: reshad@yahoo.com

Albert Cabellos-Aparicio
Technical University of Catalonia
Barcelona
Spain
Email: acabello@ac.upc.edu

Fabio Maino
Cisco
San Jose, CA
United States of America
Email: fmaino@cisco.com