

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 29, 2017

V. Fuller
D. Farinacci
Cisco Systems
A. Cabellos (Ed.)
UPC/BarcelonaTech
March 28, 2017

**Locator/ID Separation Protocol (LISP) Control-Plane
draft-ietf-lisp-rfc6833bis-01**

Abstract

This document describes the Control-Plane and Mapping Service for the Locator/ID Separation Protocol (LISP), implemented by two new 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 infrastructure, connect directly to LISP-capable Internet end sites, and comprise the bulk of LISP-speaking devices, reducing their implementation and operational complexity should also reduce the overall cost and effort of deploying 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 <http://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 September 29, 2017.

Copyright Notice

Copyright (c) 2017 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 (<http://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](#)
- [2. Definition of Terms](#) [4](#)
- [3. Basic Overview](#) [5](#)
- [4. LISP IPv4 and IPv6 Control-Plane Packet Formats](#) [7](#)
 - [4.1. LISP Control Packet Type Allocations](#) [9](#)
 - [4.2. Map-Request Message Format](#) [10](#)
 - [4.3. EID-to-RLOC UDP Map-Request Message](#) [12](#)
 - [4.4. Map-Reply Message Format](#) [14](#)
 - [4.5. EID-to-RLOC UDP Map-Reply Message](#) [18](#)
 - [4.6. Map-Register Message Format](#) [21](#)
 - [4.7. Map-Notify/Map-Notify-Ack Message Format](#) [24](#)
 - [4.8. Encapsulated Control Message Format](#) [25](#)
- [5. Interactions with Other LISP Components](#) [27](#)
 - [5.1. ITR EID-to-RLOC Mapping Resolution](#) [27](#)
 - [5.2. EID-Prefix Configuration and ETR Registration](#) [28](#)
 - [5.3. Map-Server Processing](#) [30](#)
 - [5.4. Map-Resolver Processing](#) [30](#)
 - [5.4.1. Anycast Map-Resolver Operation](#) [31](#)
- [6. Security Considerations](#) [31](#)
- [7. References](#) [32](#)
 - [7.1. Normative References](#) [32](#)
 - [7.2. Informative References](#) [33](#)
- [Appendix A. Acknowledgments](#) [36](#)
- [Appendix B. Document Change Log](#) [36](#)
 - [B.1. Changes to \[draft-ietf-lisp-rfc6833bis-01\]\(#\)](#) [36](#)
 - [B.2. Changes to \[draft-ietf-lisp-rfc6833bis-00\]\(#\)](#) [36](#)
 - [B.3. Changes to \[draft-farinacci-lisp-rfc6833bis-00\]\(#\)](#) [36](#)
- Authors' Addresses [37](#)

1. Introduction

The Locator/ID Separation Protocol [[I-D.ietf-lisp-introduction](#)] and [[I-D.ietf-lisp-rfc6830bis](#)] specifies an architecture and mechanism for replacing the addresses currently used by IP with 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 globally. Several such databases have been proposed; among them are the Content distribution Overlay Network Service for LISP (LISP-CONS) [[LISP-CONS](#)], 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) [[I-D.ietf-lisp-ddt](#)].

The LISP Mapping Service defines two new 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.lewis-lisp-gpe](#)], VXLAN [[RFC7348](#)], and VXLAN-GPE [[I-D.quinn-vxlan-gpe](#)].

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 that while this document assumes a LISP+ALT 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.

The LISP Mapping Service is an important component of the LISP toolset. Issues and concerns about the deployment of LISP for Internet traffic are discussed in [[I-D.ietf-lisp-rfc6830bis](#)].

2. 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-Resolver: A network infrastructure component that accepts LISP Encapsulated 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.

Encapsulated Map-Request: A LISP Map-Request carried within an Encapsulated Control Message, which has an additional LISP header prepended. Sent to UDP destination port 4342. The "outer" addresses are globally 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.

Negative Map-Reply: A LISP Map-Reply that contains an empty Locator-Set. Returned in response to a Map-Request if the destination EID does not exist in the mapping database. Typically, this means that the "EID" being requested is an IP address connected to a non-LISP site.

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 be used by the Map-Server when forwarding Map-Requests (re-formatted as Encapsulated Map-Requests) received through the database mapping system. 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.

For definitions of other terms -- notably Map-Request, Map-Reply, Ingress Tunnel Router (ITR), and Egress Tunnel Router (ETR) -- please consult the LISP specification [[I-D.ietf-lisp-rfc6830bis](#)].

3. 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 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 ETR when a Map-Server needs to forward a Map-Request to it.

When LISP+ALT 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 [[I-D.ietf-lisp-ddt](#)] 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 [[I-D.ietf-lisp-ddt](#)], 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.

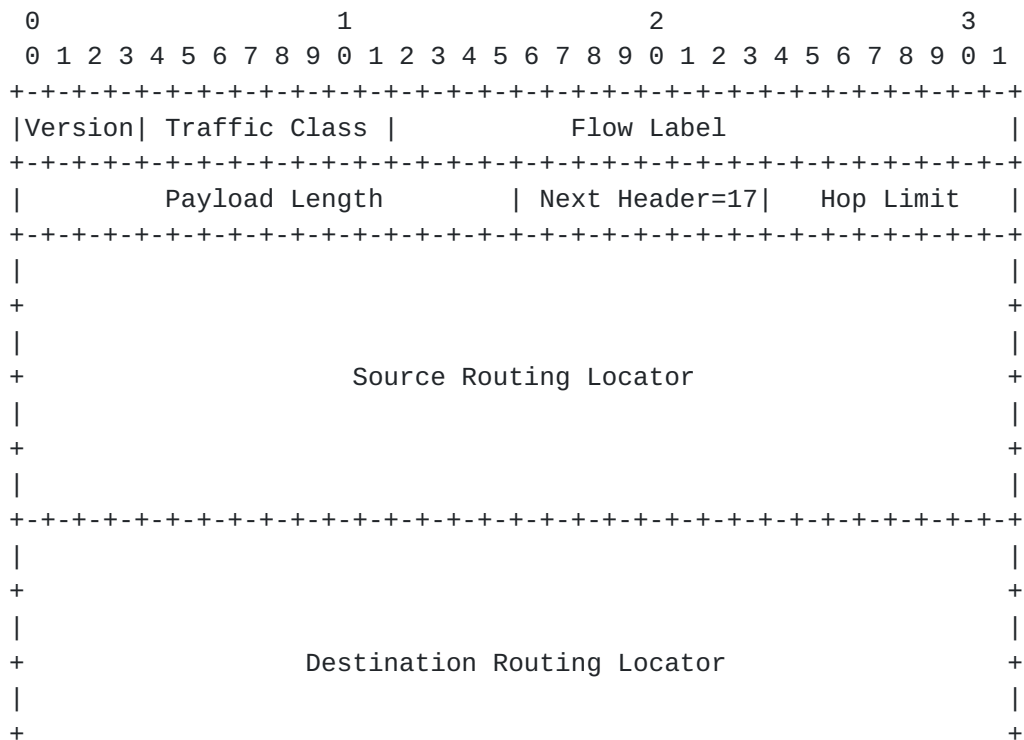
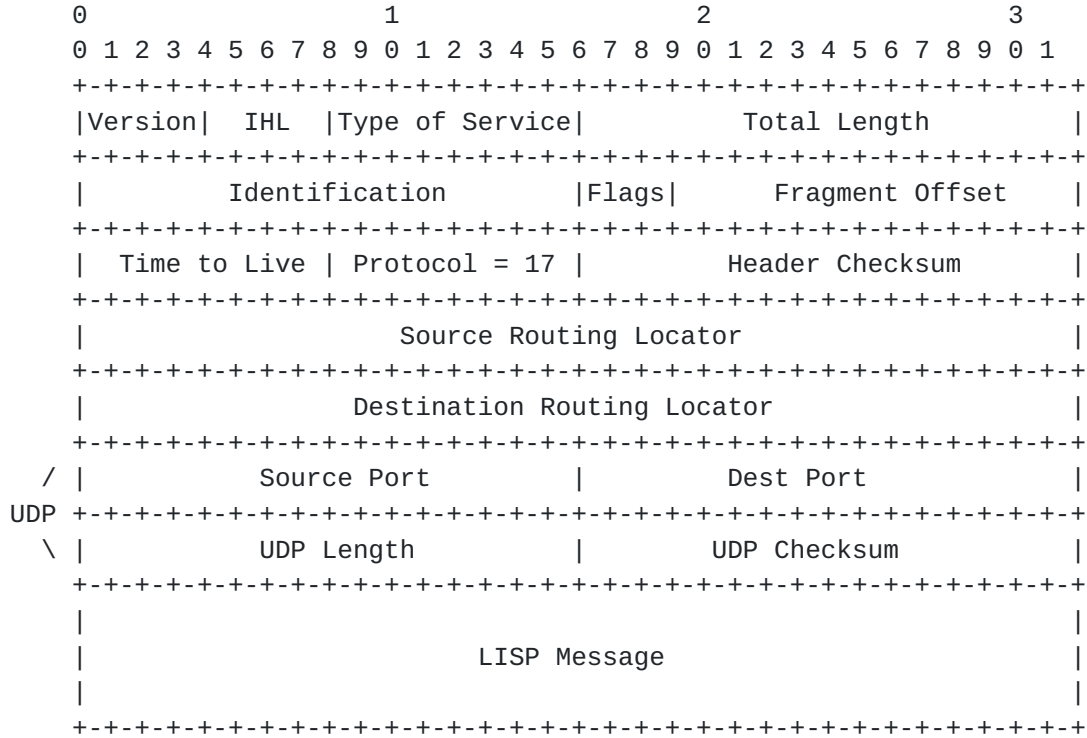
Note that while it is conceivable that a non-LISP-DDT Map-Resolver could cache responses to improve performance, issues surrounding cache management will need to be resolved so that doing so will be reliable and practical. As initially deployed, Map-Resolvers will operate only in a non-caching mode, decapsulating and forwarding Encapsulated Map Requests received from ITRs. Any specification of caching functionality is left for future work.

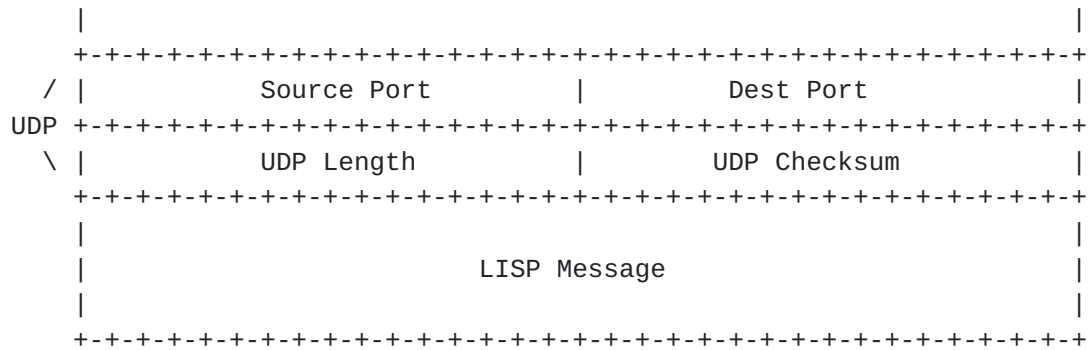
Note that a single device can implement the functions of both a Map-Server and a Map-Resolver, and in many cases the functions will be co-located in that way. Also, there can be ALT-only nodes and DDT-only nodes, when LISP+ALT and LISP-DDT are used, respectively, to connect Map-Resolvers and Map-Servers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

Detailed descriptions of the LISP packet types referenced by this document may be found in [[I-D.ietf-lisp-rfc6830bis](#)].

4. LISP IPv4 and IPv6 Control-Plane Packet Formats

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





The LISP UDP-based messages are the Map-Request and Map-Reply messages. When a UDP Map-Request is 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 is 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.

The UDP checksum is computed and set to non-zero for Map-Request, Map-Reply, Map-Register, and Encapsulated Control Message (ECM) control messages. It MUST be checked on receipt, and if the checksum fails, the packet MUST be dropped.

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

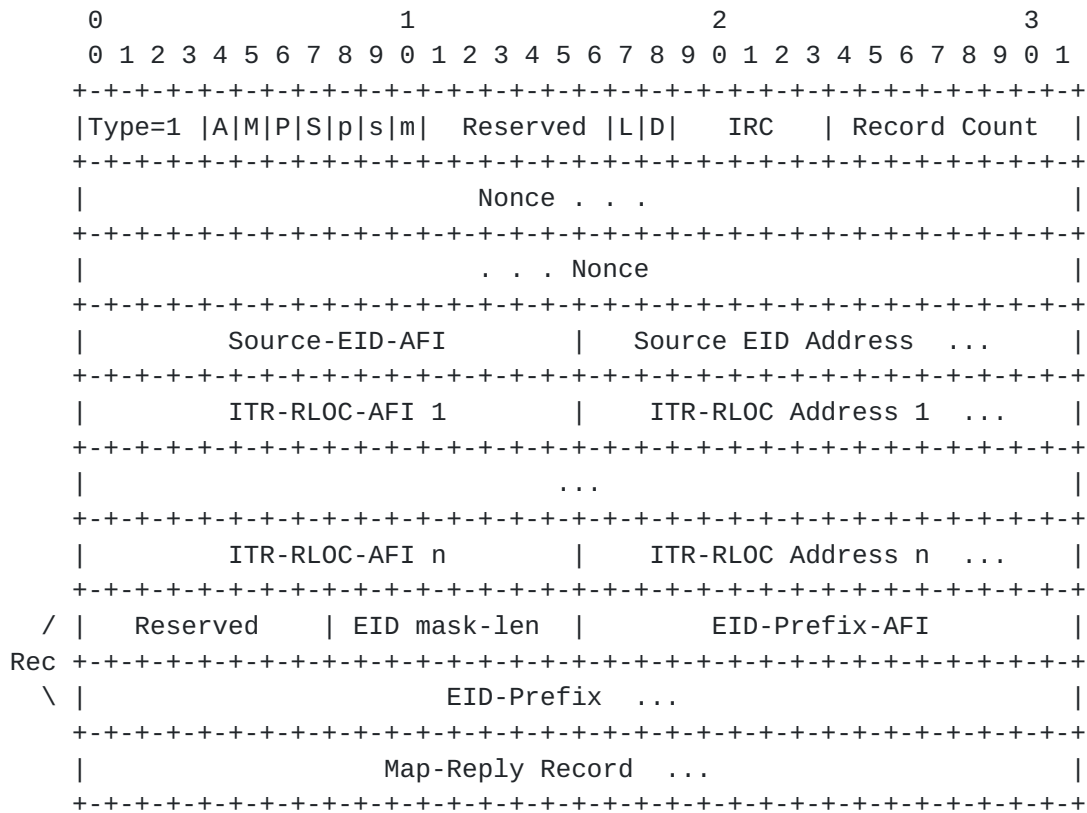
[4.1.](#) LISP Control Packet Type Allocations

This section will be the authoritative source for allocating LISP Type values and for defining LISP control message formats. Current allocations 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'
LISP Info-Request/Reply:	7	b'0111'
LISP Encapsulated Control Message:	8	b'1000'

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,

4.2. Map-Request Message Format



Packet field descriptions:

Type: 1 (Map-Request)

A: This is an authoritative bit, which is set to 0 for UDP-based Map-Requests sent by an ITR. It is set to 1 when an ITR wants the destination site to return the Map-Reply rather than the mapping database system.

M: This is the map-data-present bit. When set, it indicates that a Map-Reply Record segment is included in the Map-Request.

P: This is the probe-bit, which indicates that a Map-Request SHOULD be treated as a Locator reachability probe. The receiver SHOULD respond with a Map-Reply with the probe-bit set, indicating that the Map-Reply is a Locator reachability probe reply, with the nonce copied from the Map-Request. See RLOC-Probing [I-D.ietf-lisp-rfc6830bis] for more details.

S: This is the Solicit-Map-Request (SMR) bit. See Solicit-Map-Request (SMRs) [I-D.ietf-lisp-rfc6830bis] for details.

p: This is the Pitr bit. This bit is set to 1 when a Pitr sends a Map-Request.

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.

m: This is the LISP mobile-node m-bit. This bit is set by xTRs that operate as a mobile node as defined in [[I-D.meyer-lisp-mn](#)].

Reserved: 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 local to the site. That is, that it is part of the RLOC-set for the LISP site.

D: This is the dont-map-reply bit. It is used in the SMR procedure described in [[I-D.ietf-lisp-rtc6830bis](#)]. When an xTR sends an SMR Map-Request 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. Support for requesting multiple EIDs in a single Map-Request message will be specified in a future version of the protocol.

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 security of the LISP mapping protocol critically depends on the strength of the nonce in the Map-Request message. The nonce SHOULD be generated by a properly seeded pseudo-random (or strong

random) source. See [[RFC4086](#)] for advice on generating security-sensitive random data.

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 is 4 octets for an IPv4 address family and 16 octets for an IPv6 address family. 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.

[4.3.](#) 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.

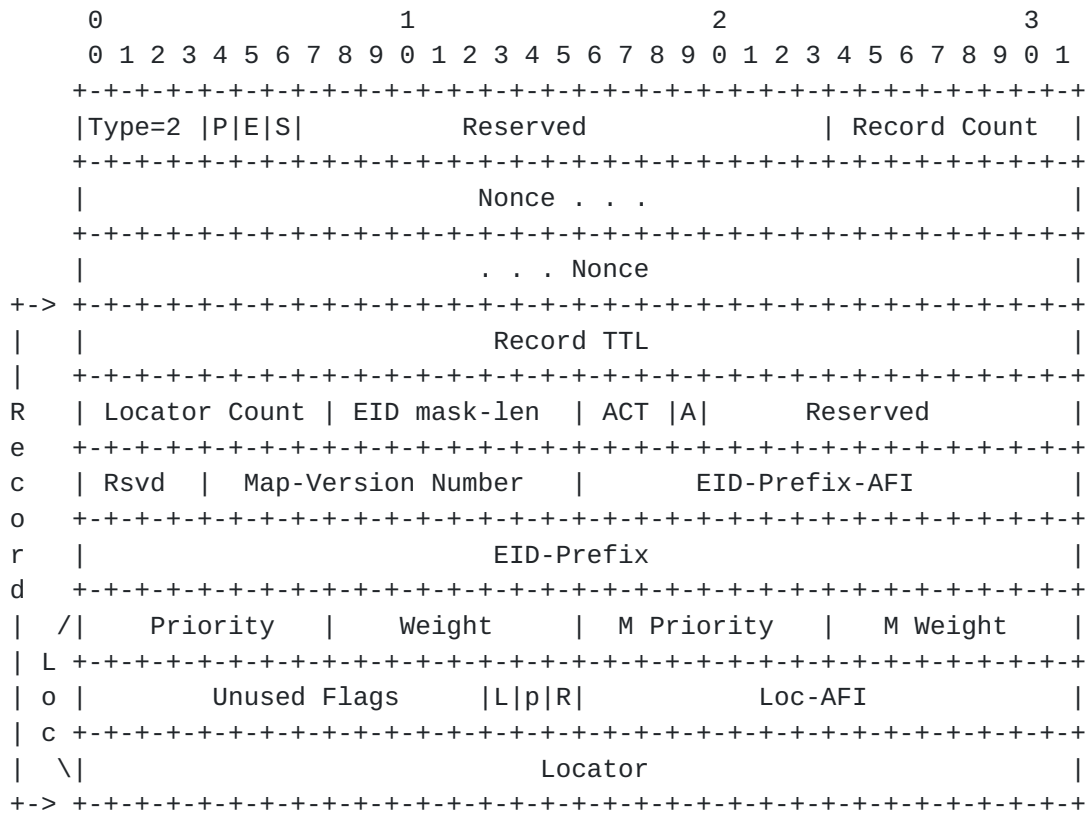
One or more Map-Request ('ITR-RLOC-AFI', 'ITR-RLOC-Address') fields MUST be filled in by the ITR. The number of fields (minus 1) encoded MUST be placed in the 'IRC' field. The ITR MAY include all locally configured Locators in this list or just provide one locator address from each address family it 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 4.8](#).

Map-Requests MUST be rate-limited. It is RECOMMENDED that a Map-Request for the same EID-Prefix be sent no more than once per second.

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 MAY originate a "verifying Map-Request", addressed to the map-requesting ITR and the ETR MAY add a Map-Cache entry. If the ETR has a Map-Cache entry that matches the "piggybacked" EID and the RLOC is in the Locator-Set for the entry, then it may send the "verifying Map-Request" directly to the originating Map-Request source. If the RLOC is not in the Locator-Set, then the ETR MUST send the "verifying Map-Request" to the "piggybacked" EID. Doing this forces the "verifying Map-Request" to go through the mapping database system to reach the authoritative source of information about that EID, guarding against RLOC-spoofing in the "piggybacked" mapping data.

4.4. Map-Reply Message Format



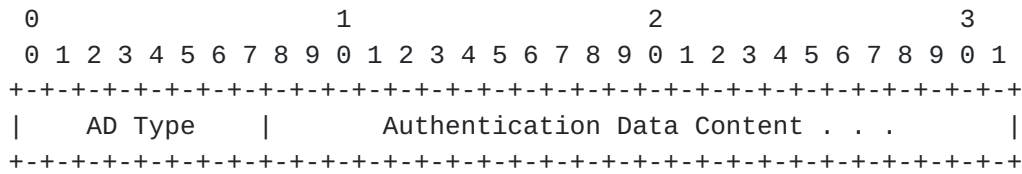
Packet field descriptions:

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 [[I-D.ietf-lisp-rfc6830bis](#)] for more details.

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-rfc6830bis](#)] 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 of signing a Map-Reply message can be found in [[I-D.ietf-lisp-sec](#)].



Reserved: This 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.

Nonce: This is a 24-bit value set in a Data-Probe packet, or a 64-bit value from the Map-Request is echoed in this 'Nonce' field of the Map-Reply. When a 24-bit value is supplied, it resides in the low-order 64 bits of the 'Nonce' field.

Record TTL: This is the time in minutes the recipient of the Map-Reply will store the mapping. If the TTL is 0, the entry SHOULD 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. 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. The actions defined are used by an ITR or PITR when a destination EID matches a negative Map-Cache entry. Unassigned values should cause a Map-Cache entry to be created, and when packets match this negative cache entry, they will be dropped. The current assigned values are:

(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 packet 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, when sent, is always set to 1 by an ETR. When a Map-Server is proxy Map-Replying for a LISP site, the Authoritative bit is set to 0. This indicates to requesting ITRs that the Map-Reply was not originated by a LISP node managed at the site that owns the EID-Prefix.

Map-Version Number: When this 12-bit value is non-zero, the Map-Reply sender is informing the ITR what the version number is for the EID record contained in the Map-Reply. The ETR can allocate this number internally but MUST coordinate this value with other ETRs for the site. When this value is 0, there is no versioning information conveyed. The Map-Version Number can be included in Map-Request and Map-Register messages. See Map-Versioning [[I-D.ietf-lisp-rfc6830bis](#)] for more 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 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 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. See also EID-Reachability [[I-D.ietf-lisp-rfc6830bis](#)] for another way the R-bit may be used.

Locator: This is an IPv4 or IPv6 address (as encoded by the 'Loc-AFI' field) assigned to an ETR. Note that the destination RLOC address MAY be an anycast address. A source RLOC 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.

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

A Map-Reply returns an EID-Prefix with a prefix 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 record of a Map-Request. The RLOCs in the Map-Reply are globally 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 [[I-D.ietf-lisp-rfc6830bis](#)] for details.

Note that a Map-Reply may contain different EID-Prefix granularity (prefix + 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 ETR is configured with overlapping EID-Prefixes, 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:

```
10.0.0.0/8
10.1.0.0/16
10.1.1.0/24
10.1.2.0/24
```

A Map-Request for EID 10.1.1.1 would cause a Map-Reply with a record count of 1 to be returned with a mapping record EID-Prefix of 10.1.1.0/24.

A Map-Request for EID 10.1.5.5 would cause a Map-Reply with a record count of 3 to be returned with mapping records for EID-Prefixes 10.1.0.0/16, 10.1.1.0/24, and 10.1.2.0/24.

Note that not all overlapping EID-Prefixes need to be returned but only the more-specific entries (note that in the second example above 10.0.0.0/8 was not returned for requesting EID 10.1.5.5) 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.

Map-Replies SHOULD be sent for an EID-Prefix no more often than once per second to the same requesting router. 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 appear in the same order for each ETR that originates a Map-Reply message. The 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-RLOC' 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.

4.6. 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 Map-Register message format is:



Packet field descriptions:

Type: 3 (Map-Register)

P: This is the proxy Map-Reply bit. When set to 1, an ETR sends a Map-Register message 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 xTR-ID bit. When this bit is set, what is appended to the Map-Register is a 128-bit xTR router-ID and then a 64-bit site-ID. See LISP NAT-Traversal procedures in [\[I-D.ermagan-lisp-nat-traversal\]](#) for details.

Reserved: This 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 the 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 [\[I-D.portoles-lisp-eid-mobility\]](#) 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.

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 [\[I-D.ietf-lisp-signal-free-multicast\]](#) for one use case example.

m: This is the mobile-node bit. When set to 1, the registering xTR supports the procedures in [\[I-D.meyer-lisp-mn\]](#).

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 set to 0 in Map-Register messages. Since the Map-Register message is authenticated, the 'Nonce' field is not currently used for any security function but may be in the future as part of an anti-replay solution.

Key ID: This is a configured ID to find the configured Message Authentication Code (MAC) algorithm and key value used for the authentication function. See Key ID Numbers in [\[I-D.ietf-lisp-rfc6830bis\]](#) 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.

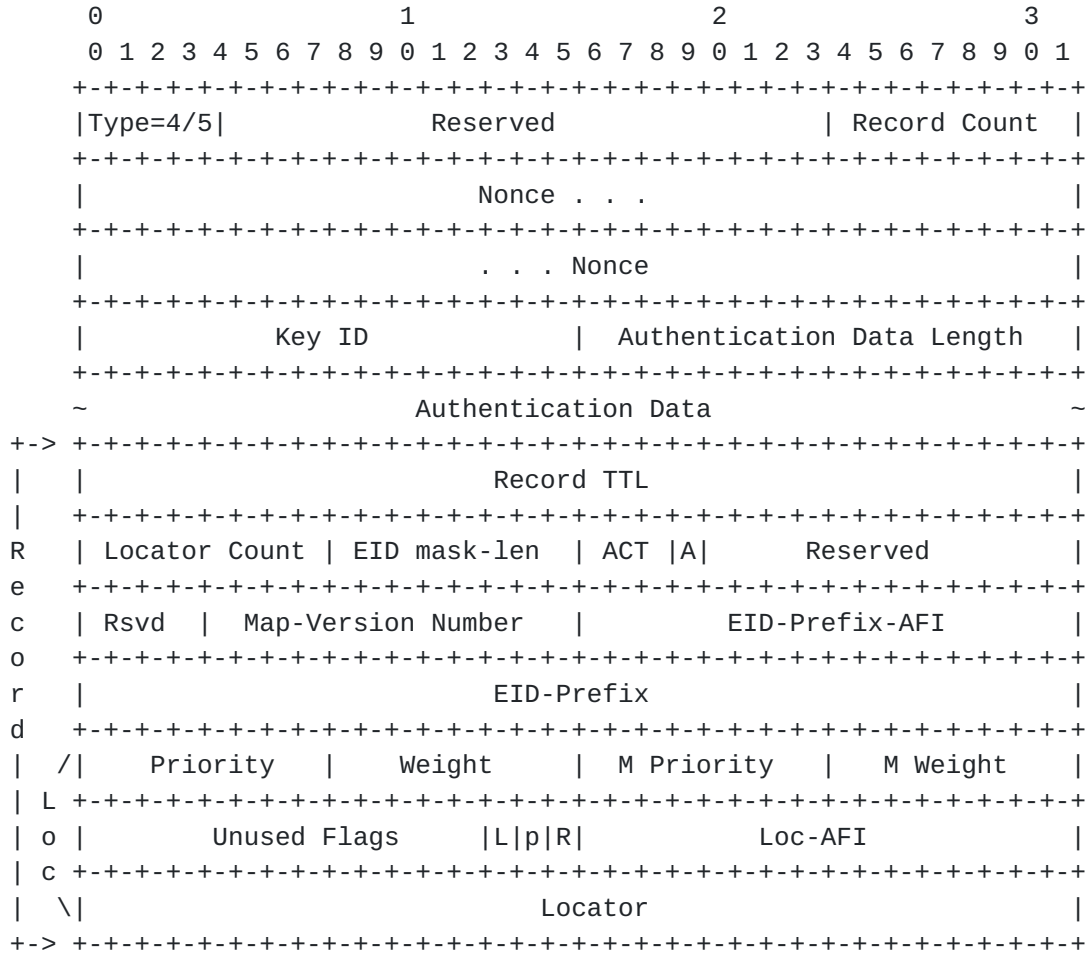
Authentication Data: This is the message digest used from the output of the MAC algorithm. The entire Map-Register payload is authenticated with this field preset to 0. After the MAC is computed, it is placed in this field. Implementations of this specification MUST include support for HMAC-SHA-1-96 [[RFC2404](#)], and support for HMAC-SHA-256-128 [[RFC4868](#)] is RECOMMENDED.

The definition of the rest of the Map-Register can be found in [Section 4.4](#).

4.7. 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:

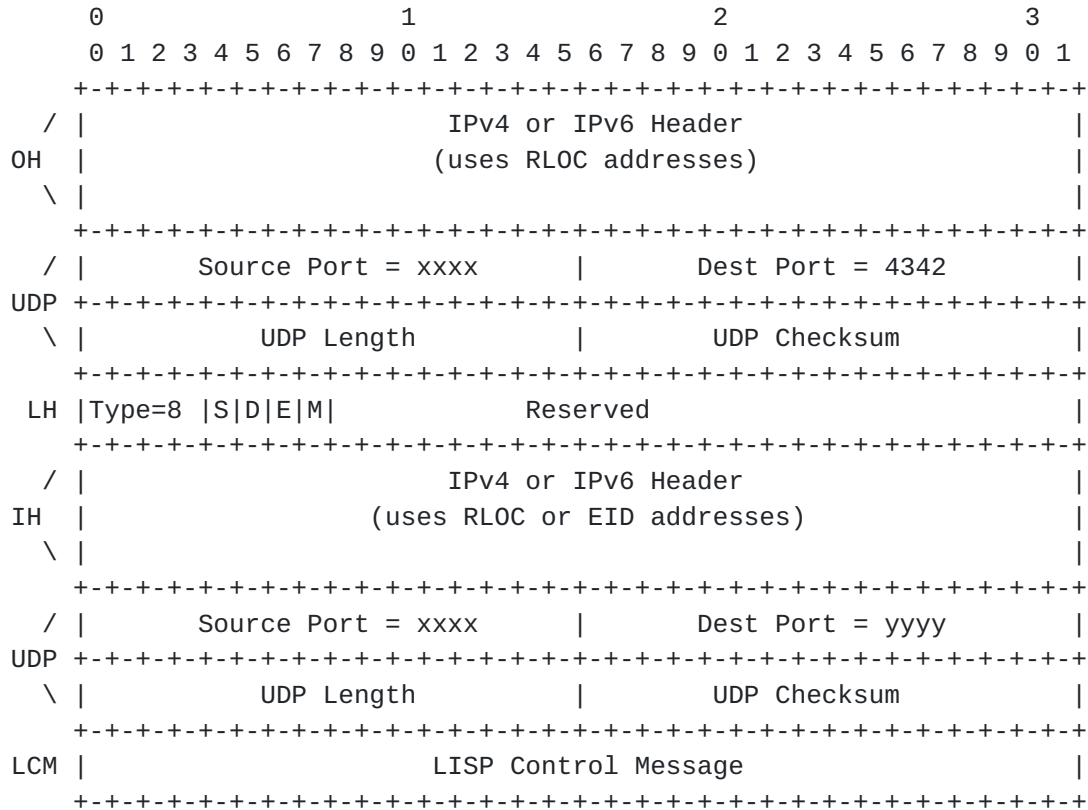
Type: 4/5 (Map-Notify/Map-Notify-Ack)

The Map-Notify message has the same contents as a Map-Register message. See the Map-Register section for field descriptions.

The Map-Notify-Ack message has the same contents as a Map-Notify message. It is used to acknowledge the receipt of a Map-Notify and for the sender to stop retransmitting a Map-Notify with the same nonce.

4.8. Encapsulated Control Message Format

An Encapsulated Control Message (ECM) is used to encapsulate control packets sent between xTRs and 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.

LH: 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.

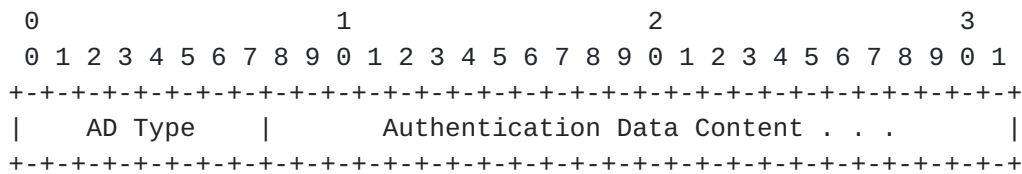
Type: 8 (Encapsulated Control Message (ECM))

S: This is the Security bit. When set to 1, the procedures from [\[I-D.ietf-lisp-sec\]](#) are followed.

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 [[I-D.ietf-lisp-ddt](#)].

E: This is the to-ETR bit. When set to 1, the Map-Server's intention is to forward the ECM to an authoritative ETR.

M: This is the to-MS bit. When set to 1, a Map-Request is being sent to a co-located Map-Resolver and Map-Server where the message can be processed directly by the Map-Server versus the Map-Resolver using the LISP-DDT procedures in [[I-D.ietf-lisp-ddt](#)].



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 this section. At this time, only Map-Request messages are allowed to be encapsulated. In the future, PIM Join/Prune messages [[RFC6831](#)] might be allowed. Encapsulating other types of LISP control messages is for further study. 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.

5. Interactions with Other LISP Components

5.1. 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. In particular, the ITR need not connect to the LISP+ALT infrastructure or implement the BGP and GRE protocols that it uses.

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 for an EID-Prefix that matches the requested EID but that does not have an actively registered, more-specific ID-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 5.2](#) 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 5.4](#) 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.

5.2. 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 shared secret or other relevant authentication information. A Map-Server's configuration must 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. As developers and operators gain experience with the mapping system, additional, stronger security measures may be added to the registration process.

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 5.3](#) 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.meyer-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 [[I-D.ietf-lisp-rfc6830bis](#)] 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.

5.3. 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 (received over the ALT if LISP+ALT is in use), 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 may 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 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.

5.4. 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.

5.4.1. Anycast Map-Resolver 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.

Note that Map-Server associations with ETRs should not use anycast addresses, as registrations need to be established between an ETR and a specific set of Map-Servers, each identified by a specific registration association.

6. Security Considerations

The 2-way LISP header nonce exchange documented in [\[I-D.ietf-lisp-rfc6830bis\]](#) can be used to avoid ITR spoofing attacks.

To publish an authoritative EID-to-RLOC mapping with a Map-Server, an ETR includes authentication data that is a hash of the message using a pair-wise shared key. An implementation must support use of HMAC-SHA-1-96 [\[RFC2104\]](#) and should support use of HMAC-SHA-256-128 [\[RFC6234\]](#) (SHA-256 truncated to 128 bits).

As noted in [Section 5.2](#), a Map-Server should verify that all EID-Prefixes registered by an ETR match the configuration stored on the Map-Server.

The currently defined authentication mechanism for Map-Register messages does not provide protection against "replay" attacks by a "man-in-the-middle". Additional work is needed in this area.

[I-D.ietf-lisp-sec] defines a proposed mechanism for providing origin authentication, integrity, anti-replay protection, and prevention of man-in-the-middle and "overclaiming" attacks on the Map-Request/Map-Reply exchange. Work is ongoing on this and other proposals for resolving these open security issues.

While beyond the scope of securing an individual Map-Server or Map-Resolver, it should be noted that a BGP-based LISP+ALT network (if ALT is used as the mapping database infrastructure) can take advantage of standards work on adding security to BGP.

A complete LISP threat analysis has been published in [[RFC7835](#)]. Please refer to it for more security related details.

7. References

7.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), DOI 10.17487/RFC2104, February 1997, <<http://www.rfc-editor.org/info/rfc2104>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc4086>>.
- [RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", [BCP 107](#), [RFC 4107](#), DOI 10.17487/RFC4107, June 2005, <<http://www.rfc-editor.org/info/rfc4107>>.
- [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, <<http://www.rfc-editor.org/info/rfc4868>>.

- [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, <<http://www.rfc-editor.org/info/rfc6234>>.
- [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, <<http://www.rfc-editor.org/info/rfc6831>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc6837>>.
- [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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc8060>>.

7.2. 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.ermagan-lisp-nat-traversal] Ermagan, V., Farinacci, D., Lewis, D., Skriver, J., Maino, F., and C. White, "NAT traversal for LISP", [draft-ermagan-lisp-nat-traversal-12](#) (work in progress), March 2017.

[I-D.ietf-lisp-ddt]

Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "LISP Delegated Database Tree", [draft-ietf-lisp-ddt-09](#) (work in progress), January 2017.

[I-D.ietf-lisp-introduction]

Cabellos-Aparicio, A. and D. Saucez, "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", [draft-ietf-lisp-introduction-13](#) (work in progress), April 2015.

[I-D.ietf-lisp-rfc6830bis]

Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", [draft-ietf-lisp-rfc6830bis-00](#) (work in progress), December 2016.

[I-D.ietf-lisp-sec]

Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D. Saucez, "LISP-Security (LISP-SEC)", [draft-ietf-lisp-sec-12](#) (work in progress), November 2016.

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

Moreno, V. and D. Farinacci, "Signal-Free LISP Multicast", [draft-ietf-lisp-signal-free-multicast-02](#) (work in progress), October 2016.

[I-D.lewis-lisp-gpe]

Lewis, D., Agarwal, P., Kreeger, L., Maino, F., Quinn, P., Smith, M., and N. Yadav, "LISP Generic Protocol Extension", [draft-lewis-lisp-gpe-02](#) (work in progress), July 2014.

[I-D.meyer-lisp-mn]

Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", [draft-meyer-lisp-mn-16](#) (work in progress), December 2016.

[I-D.portoles-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-portoles-lisp-eid-mobility-01](#) (work in progress), October 2016.

[I-D.quinn-vxlan-gpe]

Quinn, P., Manur, R., Kreeger, L., Lewis, D., Maino, F., Smith, M., Agarwal, P., Yong, L., Xu, X., Elzur, U., Garg, P., and D. Melman, "Generic Protocol Extension for VXLAN", [draft-quinn-vxlan-gpe-04](#) (work in progress), February 2015.

[LISP-CONS]

Brim, S., Chiappa, N., Farinacci, D., Fuller, V., Lewis, D., and D. Meyer, "LISP-CONS: A Content distribution Overlay Network Service for LISP", Work in Progress, April 2008.

Appendix A. Acknowledgments

The 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 on caching with LISP-CONS, some of which may be used by Map-Resolvers.

Appendix B. Document Change Log

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

B.1. 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.2. 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.3. 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 MSSs, 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 [[I-D.ietf-lisp-ddt](#)], 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

Vince Fuller
Cisco Systems

E-Mail: vaf@vaf.net

Dino Farinacci
Cisco Systems

E-Mail: farinacci@gmail.com

Albert Cabellos
UPC/BarcelonaTech
Campus Nord, C. Jordi Girona 1-3
Barcelona, Catalunya
Spain

E-Mail: acabello@ac.upc.edu

