

Internet Engineering Task Force (IETF)
Request for Comments: 9301
Obsoletes: 6830, 6833
Category: Standards Track
ISSN: 2070-1721

D. Farinacci
lispers.net
F. Maino
Cisco Systems
V. Fuller
vaf.net Internet Consulting
A. Cabellos, Ed.
Universitat Politecnica de Catalunya
October 2022

Locator/ID Separation Protocol (LISP) Control Plane

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 provide a simplified "front end" for one or more Endpoint IDs (EIDs) 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; this behavior 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, the implementation and operational complexity of the overall cost and effort of deploying LISP is reduced.

This document obsoletes RFCs 6830 and 6833.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9301>.

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
 - 1.1. Scope of Applicability
2. Requirements Notation
3. Definitions of Terms

4. Basic Overview
 5. LISP IPv4 and IPv6 Control Plane Packet Formats
 - 5.1. LISP Control Packet Type Allocations
 - 5.2. Map-Request Message Format
 - 5.3. EID-to-RLOC UDP Map-Request Message
 - 5.4. Map-Reply Message Format
 - 5.5. EID-to-RLOC UDP Map-Reply Message
 - 5.6. Map-Register Message Format
 - 5.7. Map-Notify and Map-Notify-Ack Message Formats
 - 5.8. Encapsulated Control Message Format
 6. Changing the Contents of EID-to-RLOC Mappings
 - 6.1. Solicit-Map-Request (SMR)
 7. Routing Locator Reachability
 - 7.1. RLOC-Probing Algorithm
 8. Interactions with Other LISP Components
 - 8.1. ITR EID-to-RLOC Mapping Resolution
 - 8.2. EID-Prefix Configuration and ETR Registration
 - 8.3. Map-Server Processing
 - 8.4. Map-Resolver Processing
 - 8.4.1. Anycast Operation
 9. Security Considerations
 10. Privacy Considerations
 11. Changes Related to RFCs 6830 and 6833
 12. IANA Considerations
 - 12.1. LISP UDP Port Numbers
 - 12.2. LISP Packet Type Codes
 - 12.3. LISP Map-Reply EID-Record Action Codes
 - 12.4. LISP Address Type Codes
 - 12.5. LISP Algorithm ID Numbers
 - 12.6. LISP Bit Flags
 13. References
 - 13.1. Normative References
 - 13.2. Informative References
- Acknowledgments
- Authors' Addresses

1. Introduction

The Locator/ID Separation Protocol [RFC9300] (see also [RFC9299]) specifies an architecture and mechanism for dynamic tunneling by logically separating the addresses currently used by IP in two separate namespaces: 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 and Mapping Service can be used by many different encapsulation-based or translation-based data planes, including but not limited to those defined in LISP [RFC9300], the LISP Generic Protocol Extension (LISP-GPE) [RFC9305], Virtual eXtensible Local Area Networks (VXLANs) [RFC7348], VXLAN-GPE [NVO3-VXLAN-GPE], GRE [RFC2890], the GPRS Tunneling Protocol (GTP) [GTP-3GPP], Identifier-Locator Addressing (ILA) [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) servers [RFC1035]; 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 this document doesn't assume any particular database mapping infrastructure to illustrate certain aspects of Map-Server and Map-Resolver operations. 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 Section 1.1 of [RFC9300].

This document obsoletes [RFC6830] and [RFC6833].

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 uses for LISP have been found and are being implemented. As such, the design and development of LISP have 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 Security (LISP-SEC) MUST be implemented [RFC9303]. 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 (ECMs) (Section 5.8).
2. Implementations SHOULD use 'HMAC-SHA256-128+HKDF-SHA256' as the Algorithm ID (Section 12.5) in the Map-Register message (Section 5.6) and MUST NOT use 'None' or 'HMAC-SHA-1-96-None' as the 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. Definitions 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 control plane message that queries 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 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 an RLOC-Probe Reply.

Encapsulated Map-Request: A LISP Map-Request carried within an ECM. This Map-Request 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 (reformatted 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 [RFC9300].

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 would 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, 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.

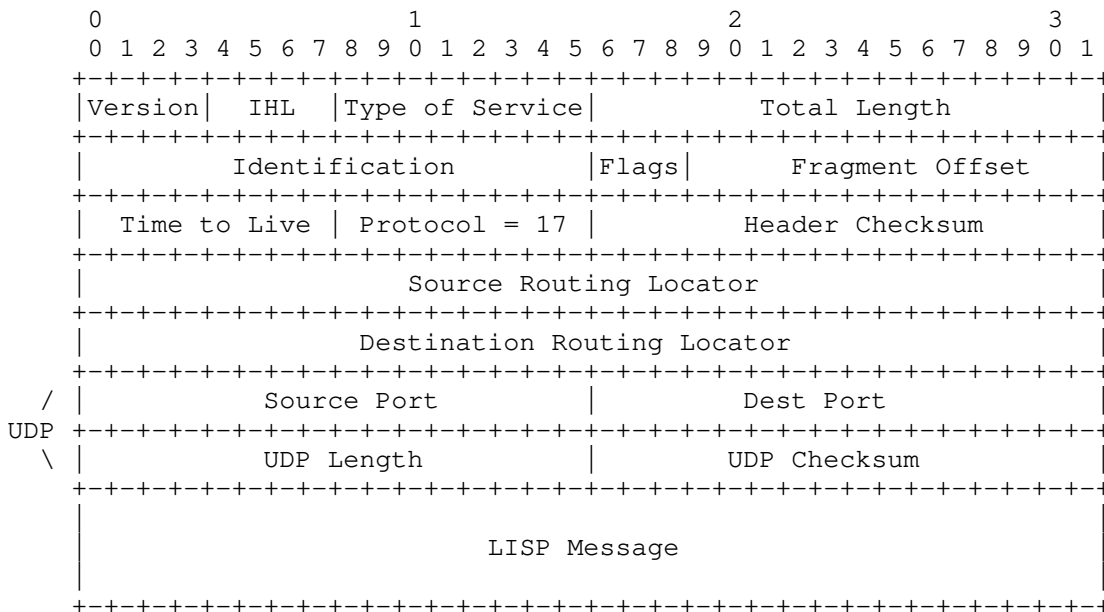
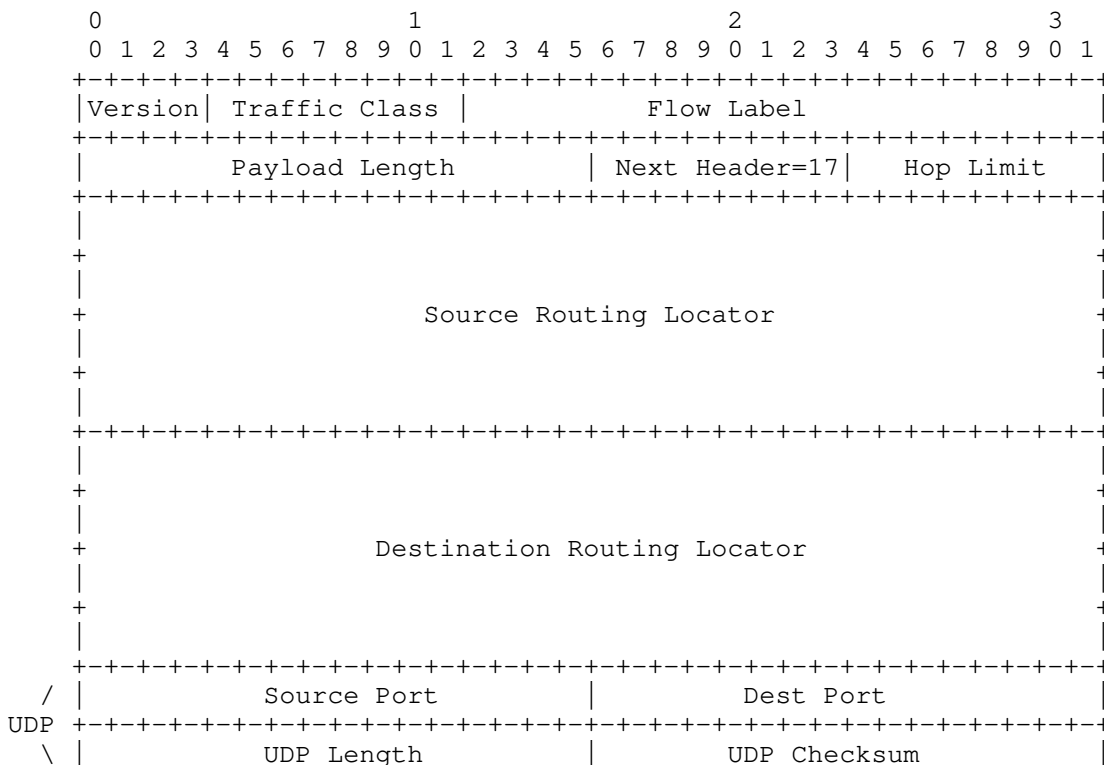


Figure 1: IPv4 UDP LISP Control Message

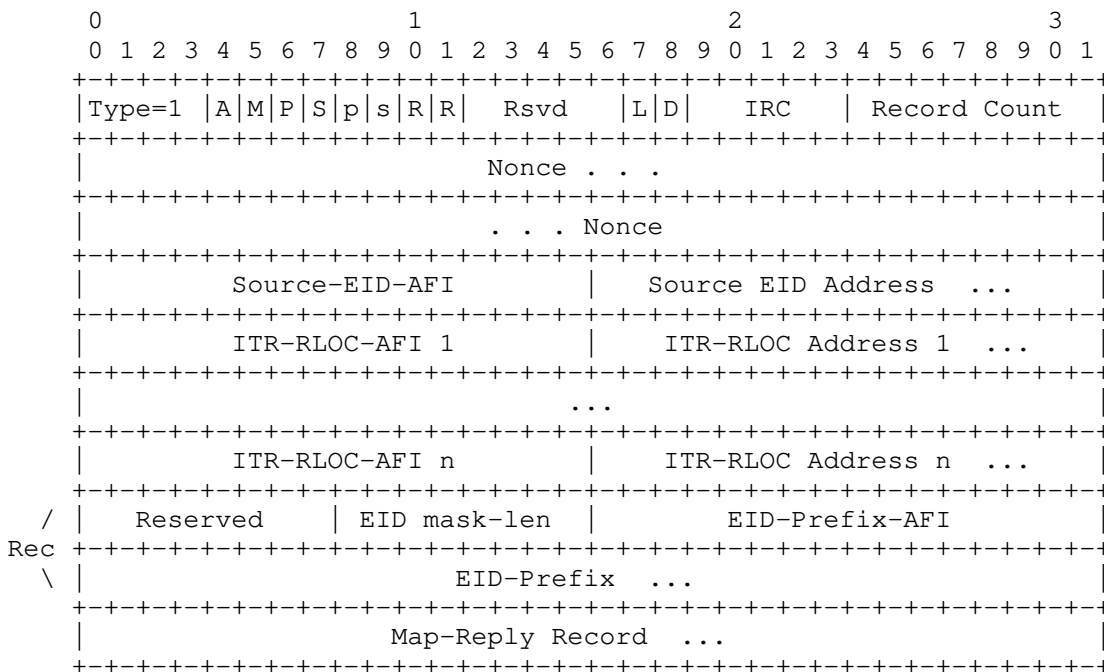


Protocol designers experimenting with new message formats are recommended to use the LISP Shared Extension Message Type described in [RFC9304].

All LISP control plane messages use Address Family Identifiers (AFIs) [AFN] or LISP Canonical Address Format (LCAF) entries [RFC8060] to encode either fixed-length 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.

5.2. Map-Request Message Format



Packet field descriptions:

Type: 1 (Map-Request)

- A: This is an authoritative bit. It is set to 1 when an ITR wants the destination site to return the Map-Reply rather than the mapping database system returning a Map-Reply and is set to 0 otherwise.
- 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 MUST be treated as a Locator reachability probe. The receiver MUST 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 Algorithm" (Section 7.1) for more details. This RLOC-Probe Map-Request MUST NOT be sent to the Mapping System. If a Map-Resolver or Map-Server receives a Map-Request with the probe-bit set, it MUST drop the message.
- S: This is the Solicit-Map-Request (SMR) bit. See "Solicit-Map-Request (SMR)" (Section 6.1) for details.
- p: This is the Proxy Ingress Tunnel Router (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; for example, see [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 of 0 is used, and this field is of zero length.

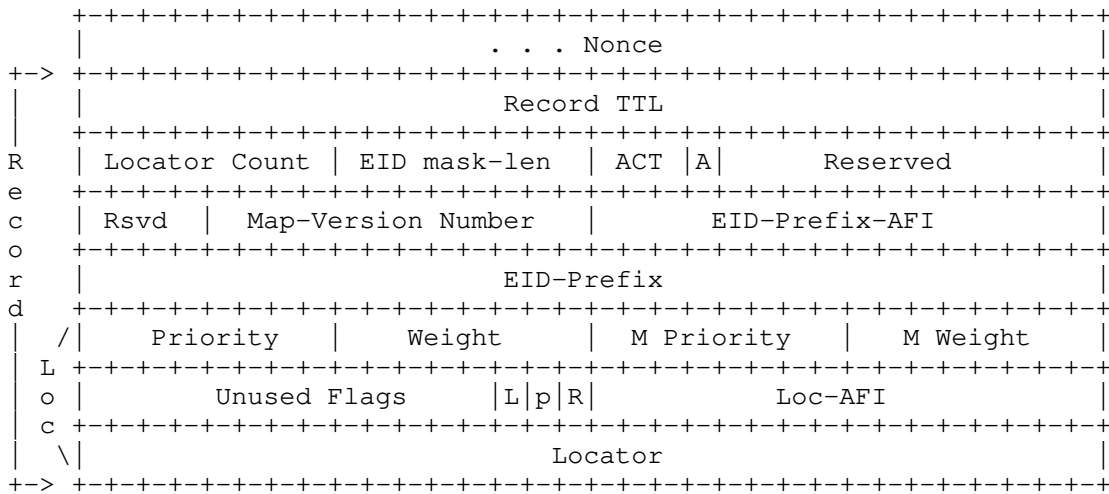
ITR-RLOC-AFI: This is the address family of the 'ITR-RLOC Address' field that follows this field.

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 [AFN] 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 [AFN], the address length varies, and for the LCAF AFI, the format is defined




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 Algorithm" (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, it 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 Section 10.1 ("Echo-Nonce Algorithm") of [RFC9300] 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. Details can be found in [RFC9303].



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-specific prefixes 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 an 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 so 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/Auth-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, see [RFC9302] for details.

EID-Prefix-AFI: This is the address family of the EID-Prefix according to [AFN] 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 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 second and third Locators will each get 25% of the traffic, and the fourth 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 Section 12 ("Routing Locator Hashing") of [RFC9300] 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 Routing 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-Replies MUST be rate limited. It is RECOMMENDED that a Map-Reply for the same destination RLOC be sent to no more than one packet every 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-Probing Algorithm" (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 [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, and 2001:db8:1:2::/64, filling out the /48 with more-specific prefixes that exist in the Mapping System.

Note that not all overlapping EID-Prefixes need to be returned but only the more-specific entries (note in the second example above that 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 TTL value so they can all time out at the same time. When a more-specific EID-Prefix is received later, its TTL 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 Map-Reply 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 Routing Locator Address is considered numerically "less than" an IPv6 Routing 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 Routing 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 chosen local IP addresses; this allows Unicast Reverse Path

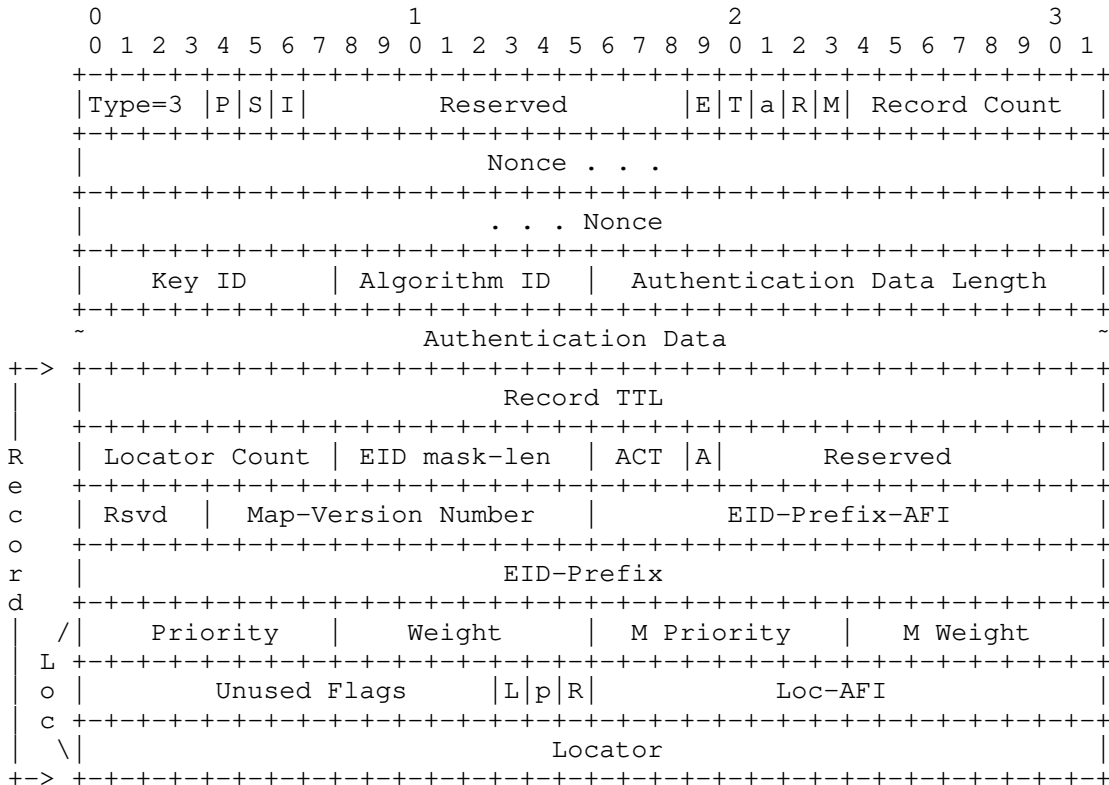
Forwarding (uRPF) checks to succeed in the upstream service provider. The destination port of a Map-Reply message is copied from the source port of the Map-Request or Data-Probe, and the source port of the Map-Reply message is set to the well-known UDP port 4342.

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 [RFC9303] are supported.

I: This is the ID-present bit. This bit is set to 1 to indicate that a 128-bit 'xTR-ID' field and a 64-bit 'Site-ID' field 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 that discovers a dynamic EID. This EID-notify-

based Map-Register is sent by the First-Hop Router to a same site xTR that propagates the Map-Register to the Mapping System. The site xTR keeps state to later Map-Notify the First-Hop Router after the EID has moved away. See [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 acknowledgment 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 that a replay attack could have occurred.

Key ID: This is 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 rotating between multiple pre-shared secrets in a nondisruptive 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 'Algorithm ID' field 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 'Algorithm ID' field 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: per-msg-key=KDF(nonce+PSK[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: per-msg-key = PSK[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 the 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 an xTR. The same xTR-ID value MUST NOT be used in two different xTRs in the scope of the Site-ID.

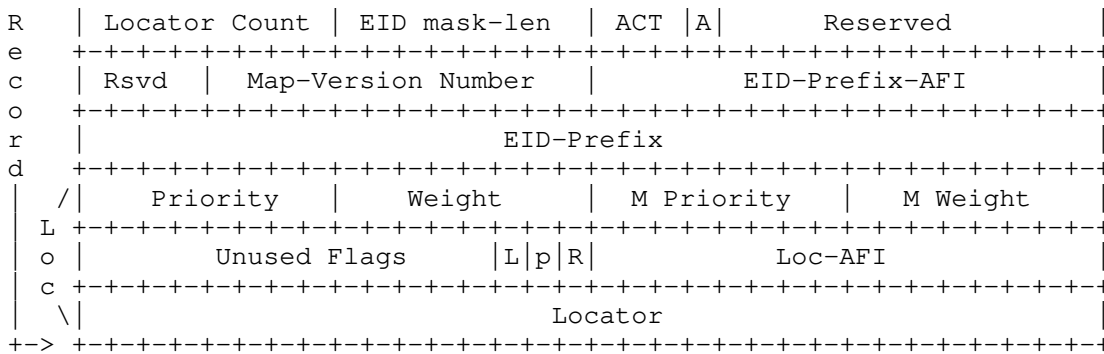
Site-ID: 'Site-ID' is a 64-bit field at the end of the Map-Register message, following the xTR-ID. The Site-ID is used to uniquely identify to which site the xTR that sent the message belongs. This document does not specify a strict meaning for the 'Site-ID' field. Informally, it provides an indication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

5.7. Map-Notify and Map-Notify-Ack Message Formats

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 "Map-Register Message Format" (Section 5.6) for field descriptions and "Map-Reply Message Format" (Section 5.4) 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 in an unsolicited manner. This topic is out of scope for this document. See [LISP-PUBSUB] for details.

After sending a Map-Register, if a Map-Notify is not received after 1 second, the transmitter MUST retransmit 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 Section 3.1 ("Congestion Control Guidelines") of [RFC8085] and Section 3.3 ("Reliability Guidelines") 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.

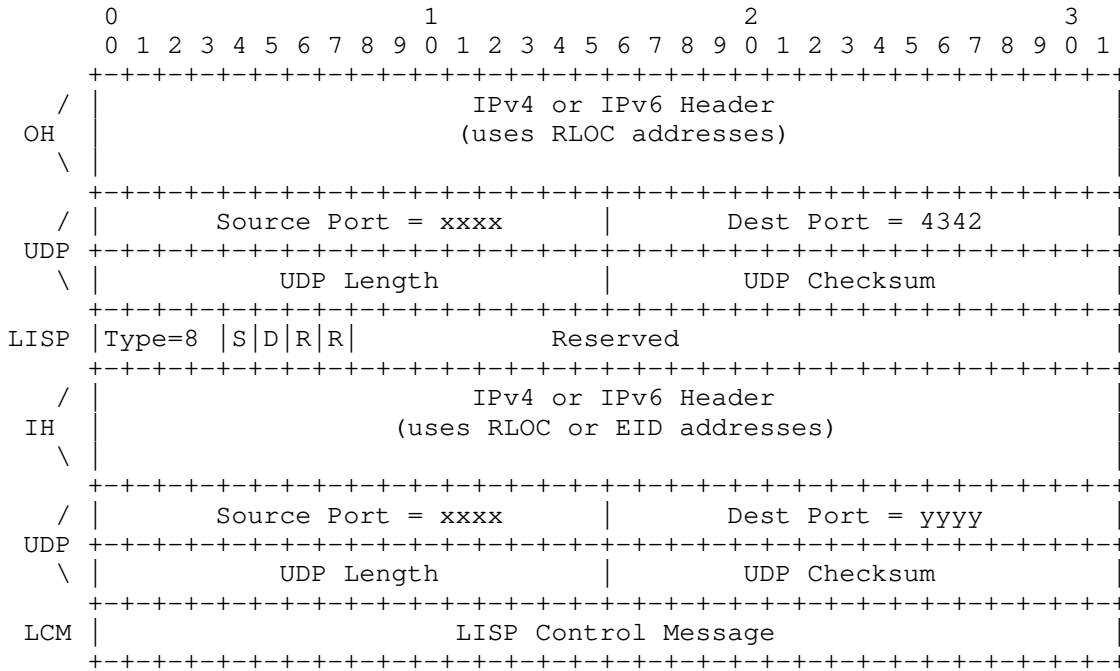
The Map-Notify-Ack message has the same contents as a Map-Notify message. It is used to acknowledge the receipt of an unsolicited Map-Notify and, once the Authentication Data is validated, allows the sender to stop retransmitting a Map-Notify with the same nonce and (validated) Authentication Data. The 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 a 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: This is the outer IPv4 or IPv6 header, which uses RLOC addresses in the source and destination header address fields.

UDP: This is 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 [RFC9303] 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 [RFC9303].



D: This is the DDT-bit. When set to 1, the sender is requesting a Map-Referral message to be returned. Details regarding 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: This is 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: This is 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 mechanisms to update mappings [RFC9300]. 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 [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, and when an ETR is also acting as an 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 it receives 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 limit a Map-Request for the same destination RLOC to no more than one packet every 3 seconds. It is RECOMMENDED that the SMR responder rate limit a Map-Request for the same EID-Prefix to no more than once every 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 [RFC9300].

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 [OPSEC-ICMP-FILTER].

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 Sections 5.3 and 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 Routing Locator Address from one of its own Routing 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:

- * An immediate Negative Map-Reply (with action code "Natively-Forward" and a 15-minute 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.
- * A Negative Map-Reply (with action code "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 a discussion of aggregate EID-Prefixes and details regarding Map-Server EID-Prefix matching.
- * 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 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 [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 requester 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/Auth-Failure action is returned. If either of these actions results as a temporary state in policy or authentication, then a Send-Map-Request action with a 1-minute TTL MAY be returned to allow the requester 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 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 that 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]. Here, 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 these security considerations, the Mapping System is considered as one trusted element.
2. The ETRs have a preconfigured trust relationship with the Mapping System, including some form of shared secret. The Mapping System is aware of which EIDs an ETR can advertise. How those keys and mappings are established is out of scope for this document.
3. LISP-SEC [RFC9303] MUST be implemented. Network operators should carefully weigh 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 into 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 [RFC9303]. 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 against 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, 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). See Section 3 of [RFC9303] for definitions of "EID-AD", "OTK-AD", and "PKT-AD".

Mechanisms to encrypt, support privacy, and prevent eavesdropping and packet tampering for messages exchanged between xTRs, between xTRs and the Mapping System, and between nodes that make up the Mapping System SHOULD be deployed. Examples of this are DTLS [RFC9147] 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 [RFC9300], LISP focuses on use cases where entities communicate over the public Internet while keeping separate addressing and topology. Here, 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. The RLOCs represent 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 [ECDSA-AUTH] in the Mapping System to make sure that only authorized users can access this information (data minimization). Use of ephemeral EIDs [EID-ANONYMITY] to achieve anonymity is another mechanism to lessen persistency and identity tracking.

11. Changes Related to RFCs 6830 and 6833

For implementation considerations, the following major changes have been made to this document since [RFC6830] and [RFC6833] were published:

- * The 16-bit 'Key ID' field of the Map-Register and Map-Notify messages as defined in [RFC6830] has been split into an 8-bit 'Key ID' field and an 8-bit 'Algorithm ID' field. Note that this change also applies to the Map-Notify-Ack message defined by this document. See Sections 5.6 and 5.7.

- * This document defines a Map-Notify-Ack message 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.
- * This document has incorporated 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.
- * Bits L and D have been added to the Map-Request message. See Section 5.3 for details.
- * Bits S, I, E, T, a, R, and M have been added to the Map-Register message. See Section 5.6 for details.
- * The nonce and the Authentication Data in the Map-Register message each behave differently; see Section 5.6 for details.
- * This document adds two new action values that are in an EID-Record that appears in Map-Reply, Map-Register, Map-Notify, and Map-Notify-Ack messages. These new action values are Drop/Policy-Denied and Drop/Auth-Failure. See Section 5.4 for details.

12. IANA Considerations

This section provides guidance to IANA regarding registration of values related to this LISP control plane specification, in accordance with BCP 26 [RFC8126].

- * LISP IANA registry allocations should not be made for purposes unrelated to LISP routing or transport protocols.
- * The following policies are used here with the meanings defined in BCP 26 [RFC8126]: "Specification Required", "IETF Review", "Experimental Use", and "First Come First Served".

There are three namespaces (listed in the sub-sections below) in LISP that have been registered (see [RFC9299]).

12.1. LISP UDP Port Numbers

IANA allocated UDP port number 4342 for the LISP control plane. IANA has updated the description for UDP port 4342 to reflect the following:

Service Name	Port Number	Transport Protocol	Description	Reference
lisp-control	4342	udp	LISP Control Packets	RFC 9301

Table 2

12.2. LISP Packet Type Codes

IANA is now authoritative for LISP Packet Type definitions, so they have replaced the registry references to [RFC6830] with references to this document.

Based on deployment experience related to [RFC6830], the Map-Notify-Ack message (message type 5) is defined in this document. IANA has registered it in the "LISP Packet Types" registry.

Message	Code	Reference
---------	------	-----------

LISP Map-Notify-Ack	5	RFC 9301
---------------------	---	----------

Table 3

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 has replaced the reference pointing to [RFC6830] to point to this document. This specification changes the Action name of value 3 from "Drop" to "Drop/No-Reason". It also adds the following new ACT values.

Value	Action	Description	Reference
4	Drop/Policy-Denied	A packet matching this Map-Cache entry is dropped because the target EID is policy-denied by the xTR or the Mapping System.	RFC 9301
5	Drop/Auth-Failure	A packet matching this Map-Cache entry is dropped because the Map-Request for the target EID fails an authentication check by the xTR or the Mapping System.	RFC 9301

Table 4: LISP Map-Reply Action Values

In addition, LISP has a number of flag fields and reserved fields, such as the flags of the LISP header fields [RFC9300]. 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] has an 8-bit Type 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 "LISP Canonical Address Format (LCAF) Types" registry is used for LCAF types. The registry for LCAF types uses 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]. Per this document, the registry has been closed.

12.5. LISP Algorithm ID Numbers

In [RFC6830], a request for a "LISP Key ID Numbers" registry was submitted. Per this document, IANA has renamed the registry to "LISP Algorithm ID Numbers" and listed this document as the registry reference.

The following Algorithm ID values are defined by this specification, as used in any packet type that references an '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]	None
HMAC-SHA256-128+HKDF-SHA256	3	[RFC4868]	[RFC4868]

Table 5

Number values are in the range of 0 to 255. Values are assigned 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 Map-Request messages, Map-Reply messages, Map-Register messages, and Encapsulated Control Messages. Bit allocations are also requested for EID-Records and RLOC-Records. The registry created should be named "LISP Control Plane Header Bits". A subregistry needs to be created per each message and EID-Record. The name of each subregistry is indicated below, along with its format and allocation of bits defined in this document. Any additional bit allocations require a specification, in accordance with policies defined in [RFC8126].

Subregistry: 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

Table 6: LISP Map-Request Header Bits

Subregistry: Map-Reply 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	2	3	4	5	6	7	8	9
Type=2										P E S										Reserved										Record Count									

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
---	-------------	---	--------------

Table 7: LISP Map-Reply Header Bits

Subregistry: Map-Register Header Bits (Section 5.6):

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=3	P S I	Reserved	E T a R M Record Count

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 Bit

Table 8: LISP Map-Register Header Bits

Subregistry: 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

Table 9: LISP Encapsulated Control Message (ECM) Header Bits

Subregistry: 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

Table 10: LISP EID-Record Header Bits

Subregistry: RLOC-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

Unused Flags	L p R	Loc-AFI	
Spec Name	IANA Name	Bit Position	Description
L	RLOC-Record-L	13	Local RLOC Bit
p	RLOC-Record-p	14	RLOC-Probe Reply Bit
R	RLOC-Record-R	15	RLOC Reachable Bit

Table 11: LISP RLOC-Record Header Bits

13. References

13.1. 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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9302] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 9302, DOI 10.17487/RFC9302, October 2022, <<https://www.rfc-editor.org/info/rfc9302>>.

[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>>.

[RFC9304] Boucadair, M. and C. Jacquenet, "Locator/ID Separation Protocol (LISP): Shared Extension Message and IANA Registry for Packet Type Allocations", RFC 9304, DOI 10.17487/RFC9304, October 2022, <<https://www.rfc-editor.org/info/rfc9304>>.

13.2. Informative References

[AFN] IANA, "Address Family Numbers", <<http://www.iana.org/assignments/address-family-numbers/>>.

[ECDSA-AUTH] Farinacci, D. and E. Nordmark, "LISP Control-Plane ECDSA Authentication and Authorization", Work in Progress, Internet-Draft, draft-ietf-lisp-ecdsa-auth-09, 11 September 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-ecdsa-auth-09>>.

[EID-ANONYMITY] Farinacci, D., Pillay-Esnault, P., and W. Haddad, "LISP EID Anonymity", Work in Progress, Internet-Draft, draft-ietf-lisp-eid-anonymity-13, 11 September 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-eid-anonymity-13>>.

[EID-MOBILITY] Portoles, 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-10, 10 July 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-eid-mobility-10>>.

[GTP-3GPP] 3GPP, "General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)", TS.29.281, June 2022, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1699>>.

[INTAREA-ILA] Herbert, T. and P. Lapukhov, "Identifier-locator addressing for IPv6", Work in Progress, Internet-Draft, draft-herbert-intarea-ila-01, 5 March 2018, <<https://datatracker.ietf.org/doc/html/draft-herbert-intarea-ila-01>>.

[LISP-MN] Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", Work in Progress, Internet-Draft, draft-ietf-lisp-mn-12, 24 July 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-mn-12>>.

[LISP-PUBSUB] Rodriguez-Natal, A., Ermagan, V., Cabellos-Aparicio, A., Barkai, S., and M. Boucadair, "Publish/Subscribe Functionality for LISP", Work in Progress, Internet-Draft, draft-ietf-lisp-pubsub-09, 28 June 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-pubsub-09>>.

[NVO3-VXLAN-GPE] Maino, F., Ed., Kreeger, L., Ed., and U. Elzur, Ed., "Generic Protocol Extension for VXLAN (VXLAN-GPE)", Work in Progress, Internet-Draft, draft-ietf-nvo3-vxlan-gpe-12, 22 September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-nvo3-vxlan-gpe-12>>.

[OPSEC-ICMP-FILTER]

Gont, F., Gont, G., and C. Pignataro, "Recommendations for filtering ICMP messages", Work in Progress, Internet-Draft, draft-ietf-opsec-icmp-filtering-04, 3 July 2013, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsec-icmp-filtering-04>>.

- [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>>.
- [RFC2890] Dommety, 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>>.
- [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>>.
- [RFC9147] Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", RFC 9147, DOI 10.17487/RFC9147, April 2022, <<https://www.rfc-editor.org/info/rfc9147>>.
- [RFC9299] Cabellos, A. and D. Saucez, Ed., "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", RFC 9299, DOI 10.17487/RFC9299, October 2022, <<https://www.rfc-editor.org/info/rfc9299>>.
- [RFC9305] Maino, F., Ed., Lemon, J., Agarwal, P., Lewis, D., and M. Smith, "Locator/ID Separation Protocol (LISP) Generic Protocol Extension", RFC 9305, DOI 10.17487/RFC9305, October 2022, <<https://www.rfc-editor.org/info/rfc9305>>.

Acknowledgments

The original authors would like to thank Greg Schudel, Darrel Lewis, John Zwiebel, Andrew Partan, Dave Meyer, Isidor Kouvelas, Jesper Skriver, 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 the 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 Kühlewind, Francis Dupont, Benjamin Kaduk, Eric Rescorla, Alvaro Retana, Alexey Melnikov, Alissa Cooper, Suresh Krishnan, Alberto Rodriguez-Natal, Vina Ermagan, 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.

Authors' Addresses

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

Fabio Maino
Cisco Systems
San Jose, CA
United States of America
Email: fmaino@cisco.com

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

Albert Cabellos (editor)
Universitat Politecnica de Catalunya
c/ Jordi Girona s/n
08034 Barcelona
Spain
Email: acabello@ac.upc.edu