Internet Engineering Task Force (IETF)
Request for Comments: 8060
Category: Experimental
ISSN: 2070-1721

D. Farinacci
lispers.net
D. Meyer
Brocade
J. Snijders
NTT

February 2017

LISP Canonical Address Format (LCAF)

Abstract

This document defines a canonical address format encoding used in Locator/ID Separation Protocol (LISP) control messages and in the encoding of lookup keys for the LISP Mapping Database System.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. 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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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 http://www.rfc-editor.org/info/rfc8060.

Table of Contents

1. Introduction

The LISP architecture and protocol [RFC6830] introduces two new numbering spaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). To provide flexibility for current and future applications, these values can be encoded in LISP control messages using a general syntax that includes Address Family Identifier (AFI), length, and value fields.

Currently defined AFIs include IPv4 and IPv6 addresses, which are formatted according to code-points assigned in the "Address Family Numbers" registry [AFN] as follows:

IPv4-Encoded Address:

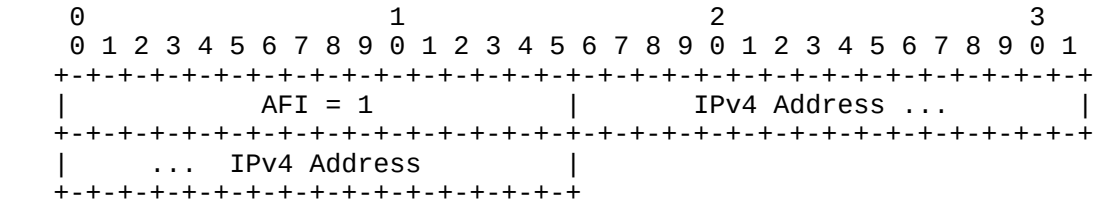

IPv6-Encoded Address:

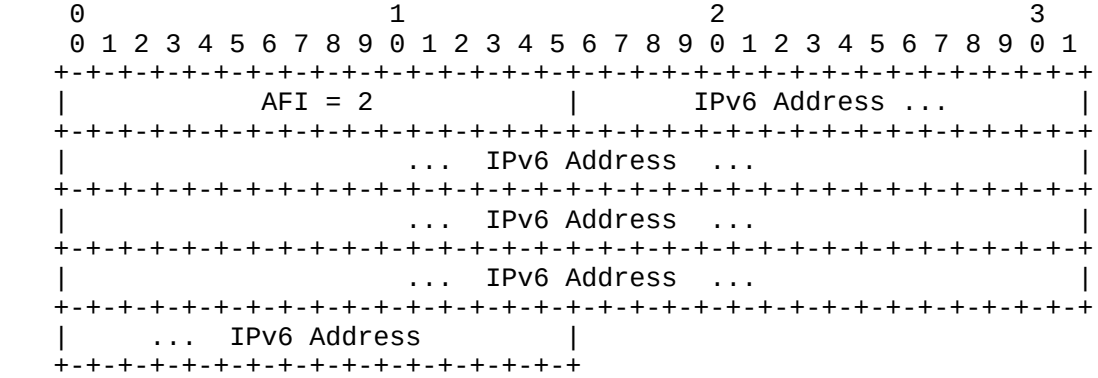

This document describes the currently defined AFIs that LISP uses along with their encodings and introduces the LISP Canonical Address Format (LCAF) that can be used to define the LISP-specific encodings for arbitrary AFI values.

Specific detailed uses for the LCAF Types defined in this document can be found in the use-case documents that implement them. The same LCAF Type may be used by more than one use-case document. As an Experimental specification, this work is, by definition, incomplete.

The LCAF Types defined in this document are to support experimentation and are intended for cautious use in self-contained environments in support of the corresponding use-case documents. This document provides assignment for an initial set of approved LCAF Types (registered with IANA) and additional unapproved LCAF Types [RFC6830]. The unapproved LCAF encodings are defined to support further study and experimentation.

2. Terminology

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.2. Definition of Terms

Address Family Identifier (AFI): a term used to describe an address encoding in a packet. Address families are defined for IPv4 and IPv6. See [AFN] and [RFC3232] for details. The reserved AFI value of 0 is used in this specification to indicate an unspecified encoded address where the length of the address is 0 bytes following the 16-bit AFI value of 0.

Unspecified Address Format:

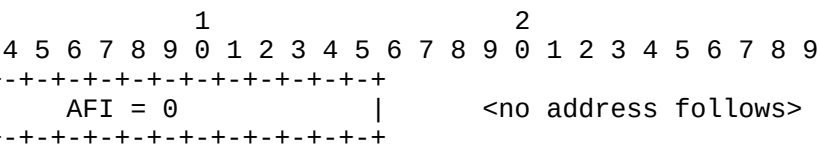

Endpoint ID (EID): a 32-bit (for IPv4) or 128-bit (for IPv6) value used in the source and destination address fields of the first (most inner) LISP header of a packet. The host obtains a destination EID the same way it obtains a destination address today, for example, through a DNS lookup or SIP exchange. The source EID is obtained via existing mechanisms used to set a host's "local" IP address. An EID is allocated to a host from an EID-prefix block associated with the site where the host is located. An EID can be used by a host to refer to other hosts.

Routing Locator (RLOC): the IPv4 or IPv6 address of an Egress Tunnel Router (ETR). It is the output of an EID-to-RLOC mapping lookup. An EID maps to one or more RLOCs. Typically, RLOCs are numbered from topologically aggregatable blocks that are assigned to a site at each point to which it attaches to the global Internet; where

the topology is defined by the connectivity of provider networks, RLOCs can be thought of as Provider-Assigned (PA) addresses. Multiple RLOCs can be assigned to the same ETR device or to multiple ETR devices at a site.

3. LISP Canonical Address Format Encodings

IANA has assigned AFI value 16387 (0x4003) to the LISP Canonical Address Format (LCAF). This specification defines the encoding format of the LISP Canonical Address (LCA). This section defines all Types for which an initial allocation in the LISP-LCAF registry is requested. See Section 7 for the complete list of such Types.

The AFI definitions in [AFN] only allocate code-points for the AFI value itself. The length of the address or entity that follows is not defined and is implied based on conventional experience. When LISP uses LCAF definitions from this document, the AFI-based address lengths are specified in this document. When new LCAF definitions are defined in other use-case documents, the AFI-based address lengths for any new AFI-encoded addresses are specified in those documents.

The first 6 bytes of a LISP Canonical Address are followed by a variable number of fields of variable length:

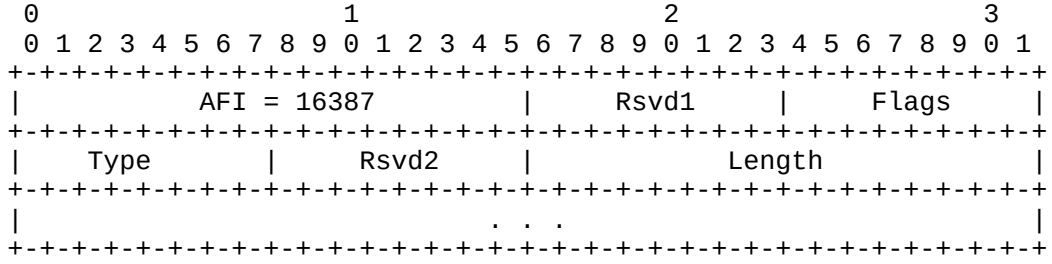

Rsvd1/Rsvd2: these 8-bit fields are reserved for future use and MUST be transmitted as 0 and ignored on receipt.

Flags: this 8-bit field is for future definition and use. For now, set to zero on transmission and ignored on receipt.

Type: this 8-bit field is specific to the LISP Canonical Address Format encodings. Both approved and unapproved values are listed below. Unapproved values are indicated; see Section 5 for more details.

Type 0: Null Body

Type 1: AFI List

Type 2: Instance ID

Type 3: AS Number

Type 4: Application Data (unapproved; see Section 5)

Type 5: Geo-Coordinates

Type 6: Opaque Key (unapproved; see Section 5)

Type 7: NAT-Traversal

Type 8: Nonce Locator (unapproved; see Section 5)

Type 9: Multicast Info

Type 10: Explicit Locator Path

Type 11: Security Key

Type 12: Source/Dest Key

Type 13: Replication List Entry

Type 14: JSON Data Model (unapproved; see Section 5)

Type 15: Key/Value Address Pair (unapproved; see Section 5)

Type 16: Encapsulation Format (unapproved; see Section 5)

Length: this 16-bit field is in units of bytes and covers all of the LISP Canonical Address payload, starting and including the byte after the Length field. When including the AFI, an LCAF-encoded address will have a minimum length of 8 bytes when the Length field is 0. The 8 bytes include the AFI, Flags, Type, Rsvd1, Rsvd2, and Length fields. When the AFI is not next to an encoded address in a control message, the encoded address will have a minimum length of 6 bytes when the Length field is 0. The 6 bytes include the Flags, Type, Rsvd1, Rsvd2, and Length fields.

[RFC6830] states RLOC-records based on an IP address are sorted when encoded in control messages, so the locator-set has consistent order across all xTRs for a given EID. The sort order is based on sort-key {afi, RLOC-address}. When an RLOC based on an IP address is LCAF encoded, the sort-key is {afi, LCAF-Type}. Therefore, when a locator-set has a mix of AFI records and LCAF records, they are ordered from smallest to largest AFI value.

4. LISP Canonical Address Applications

The following sections define the LCAF for the currently approved initial set of Type values.

4.1. Segmentation Using LISP

When multiple organizations inside of a LISP site are using private addresses [RFC1918] as EID prefixes, their address spaces must remain segregated due to possible address duplication. An Instance ID in the address encoding can aid in making the entire AFI-based address unique.

Another use for the Instance ID LISP Canonical Address Format is when creating multiple segmented VPNs inside of a LISP site where keeping EID-prefix-based subnets is desirable.

Instance ID LISP Canonical Address Format:

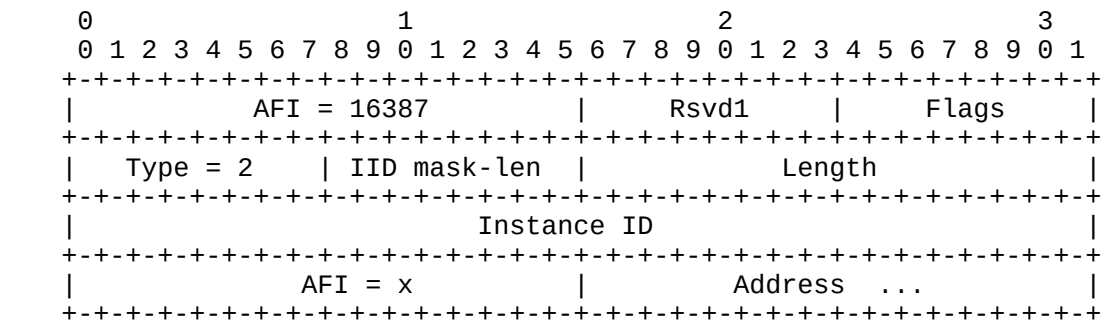

IID mask-len: if the AFI is set to 0, then this format is not encoding an extended EID prefix, but rather an Instance ID range where the 'IID mask-len' indicates the number of high-order bits used in the Instance ID field for the range. The low-order bits of the Instance ID field must be 0.

Length: length in bytes starting and including the byte after this Length field.

Instance ID: the low-order 24 bits that can go into a LISP data header when the I bit is set. See [RFC6830] for details. The reason for the length difference is so that the maximum number of instances supported per mapping system is 2^32, while conserving space in the LISP data header. This comes at the expense of limiting the maximum number of instances per xTR to 2^24. If an xTR is configured with multiple Instance IDs where the value in the high-order 8 bits is the same, then the low-order 24 bits MUST be unique.

AFI = x: x can be any AFI value from [AFN].

This LISP Canonical Address Type can be used to encode either EID or RLOC addresses.

Usage: When used as a lookup key, the EID is regarded as an extended-EID in the mapping system. This encoding is used in EID-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP Delegated Database Tree (LISP-DDT) [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages.

4.2. Carrying AS Numbers in the Mapping Database

When an Autonomous System (AS) number is stored in the LISP Mapping Database System for either policy or documentation reasons, it can be encoded in a LISP Canonical Address.

AS Number LISP Canonical Address Format:

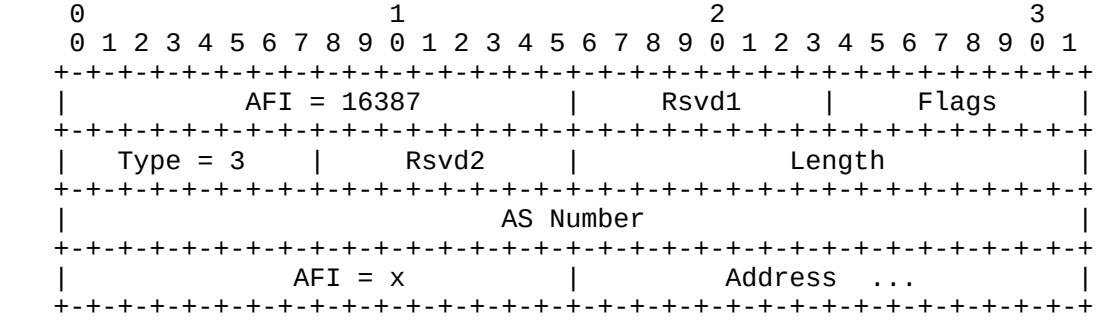

Length: length in bytes starting and including the byte after this Length field.

AS Number: the 32-bit AS number of the autonomous system that has been assigned to either the EID or RLOC that follows.

AFI = x: x can be any AFI value from [AFN].

The AS Number LCAF Type can be used to encode either EID or RLOC addresses. The former is used to describe the LISP-ALT AS number the EID prefix for the site is being carried for. The latter is used to describe the AS that is carrying RLOC based prefixes in the underlying routing system.

Usage: This encoding can be used in EID-records or RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages.

4.3. Assigning Geo-Coordinates to Locator Addresses

If an ETR desires to send a Map-Reply describing the Geo-Coordinates for each locator in its locator-set, it can use the Geo-Coordinates LCAF Type to convey physical location information.

Coordinates are specified using the WGS 84 (World Geodetic System 1984) reference coordinate system [WGS-84].

Geo-Coordinates LISP Canonical Address Format:

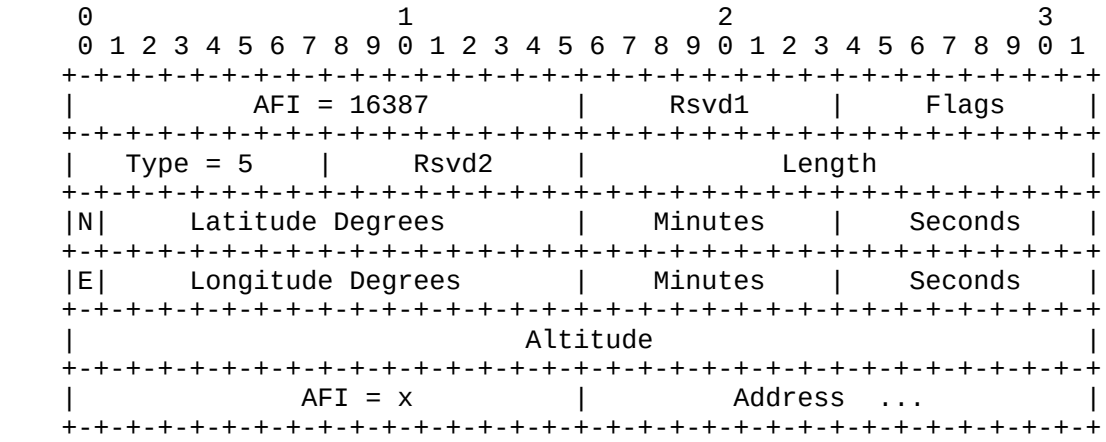

Length: length in bytes starting and including the byte after this Length field.

N: When set to 1 means north; otherwise, south.

Latitude Degrees: Valid values range from 0 to 90 degrees above or below the equator (northern or southern hemisphere, respectively).

Latitude Minutes: Valid values range from 0 to 59.

Latitude Seconds: Valid values range from 0 to 59.

E: When set to 1 means east; otherwise, west.

Longitude Degrees: Valid values are from 0 to 180 degrees right or left of the Prime Meridian.

Longitude Minutes: Valid values range from 0 to 59.

Longitude Seconds: Valid values range from 0 to 59.

Altitude: Height relative to sea level in meters. This is a two's complement signed integer meaning that the altitude could be below sea level. A value of 0x7fffffff indicates no Altitude value is encoded.

AFI = x: x can be any AFI value from [AFN].

The Geo-Coordinates LCAF Type can be used to encode either EID or RLOC addresses. When used for EID encodings, you can determine the physical location of an EID along with the topological location by observing the locator-set.

Usage: This encoding can be used in EID-records or RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages.

The use of the Geo-Coordinates LCAF encoding raises privacy issues as location information is privacy sensitive, and possibly unexpectedly privacy-sensitive information may be conveyed, e.g., if the location information corresponds to a router located in a person's home. Therefore, this encoding should not be used unless needed for operation of a LISP deployment. Before electing to utilize this encoding, care should be taken to ensure the appropriate policies are being used by the EID for controlling the conveyed information.

4.4. NAT Traversal Scenarios

When a LISP system is conveying global-address and mapped-port information when traversing through a NAT device, the NAT-Traversal LCAF Type is used. See [NAT-LISP] for details.

NAT-Traversal Canonical Address Format:

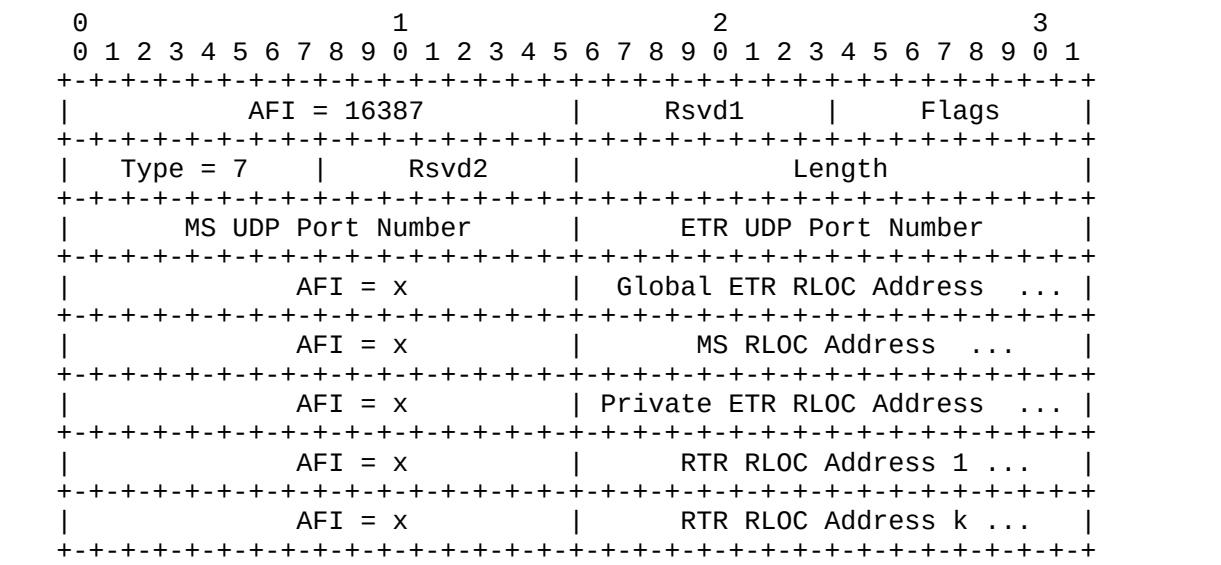

Length: length in bytes starting and including the byte after this Length field.

MS UDP Port Number: this is the UDP port number of the Map-Server and is set to 4342.

ETR UDP Port Number: this is the port number returned to a LISP system that was copied from the source port from a packet that has flowed through a NAT device.

AFI = x: x can be any AFI value from [AFN].

Global ETR RLOC Address: this is an address known to be globally unique built by NAT-traversal functionality in a LISP router.

MS RLOC Address: this is the address of the Map-Server used in the destination RLOC of a packet that has flowed through a NAT device.

Private ETR RLOC Address: this is an address known to be a private address inserted in this LCAF by a LISP router that resides on the private side of a NAT device.

RTR RLOC Address: this is an encapsulation address used by an Ingress Tunnel Router (ITR) or Proxy Ingress Tunnel Router (PITR) that resides behind a NAT device. This address is known to have state in a NAT device so packets can flow from it to the LISP ETR

behind the NAT. There can be one or more NAT Re-encapsulating Tunnel Router (RTR) [NAT-LISP] addresses supplied in these set of fields. The number of RTRs encoded is determined by parsing each field. When there are no RTRs supplied, the RTR fields can be omitted and reflected by the LCAF length field or an AFI of 0 can be used to indicate zero RTRs encoded.

Usage: This encoding can be used in Info-Request and Info-Reply messages. The mapping system does not store this information. The information is used by an xTR and Map-Server to convey private and public address information when traversing NAT and firewall devices.

Care should be taken to protect privacy against the adverse use of a Global or Private ETR RLOC Address by ensuring policy controls are used during EID registrations that use this LCAF Type in RLOC-records. Refer to the use-case documents for additional information.

4.5. Multicast Group Membership Information

Multicast group information can be published in the mapping database. So a lookup on a group address EID can return a replication list of RLOC group addresses or RLOC unicast addresses. The intent of this type of unicast replication is to deliver packets to multiple ETRs at receiver LISP multicast sites. The locator-set encoding for this EID-record Type can be a list of ETRs when they each register with "Merge Semantics". The encoding can be a typical AFI-encoded locator address. When an RTR list is being registered (with multiple levels according to [LISP-RE]), the Replication List Entry LCAF Type is used for locator encoding.

This LCAF encoding can be used to send broadcast packets to all members of a subnet when an EID is away from its home subnet location.

Multicast Info Canonical Address Format:

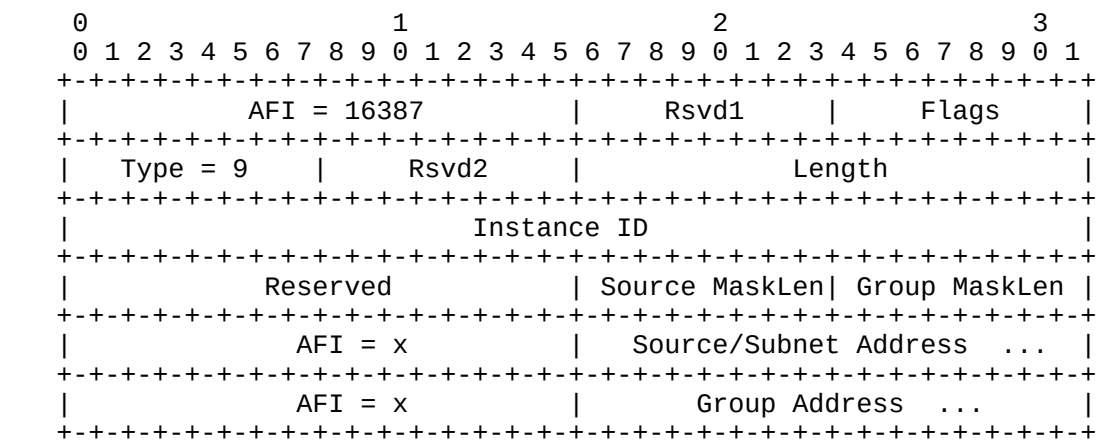

Length: length in bytes starting and including the byte after this Length field.

Reserved: must be set to zero and ignored on receipt.

Instance ID: the low-order 24 bits that can go into a LISP data header when the I bit is set. See [RFC6830] for details. The use of the Instance ID in this LCAF Type is to associate a multicast forwarding entry for a given VPN. The Instance ID describes the VPN and is registered to the mapping database system as a 3-tuple of (Instance ID, S-prefix, G-prefix).

Source MaskLen: the mask length of the source prefix that follows. The length is the number of high-order mask bits set.

Group MaskLen: the mask length of the group prefix that follows. The length is the number of high-order mask bits set.

AFI = x: x can be any AFI value from [AFN]. When a specific address family has a multicast address semantic, this field must be either a group address or a broadcast address.

Source/Subnet Address: the source address or prefix for encoding an (S,G) multicast entry.

Group Address: the group address or group prefix for encoding (S,G) or (*,G) multicast entries.

Usage: This encoding can be used in EID-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages.

4.6. Traffic Engineering Using Re-encapsulating Tunnels

For a given EID lookup into the mapping database, this LCAF can be returned to provide a list of locators in an explicit re-encapsulation path. See [LISP-TE] for details.

Explicit Locator Path (ELP) Canonical Address Format:

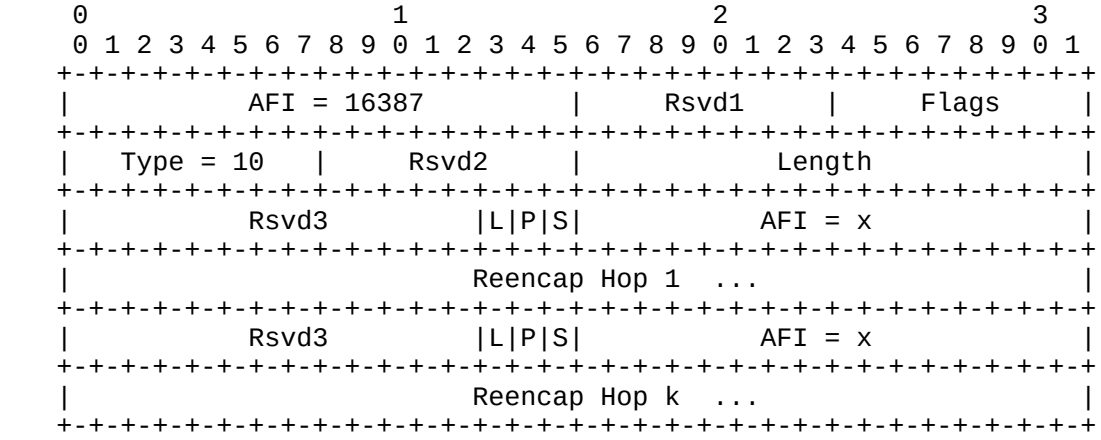

Length: length in bytes starting and including the byte after this Length field.

Rsvd3: this field is reserved for future use and MUST be transmitted as 0 and ignored on receipt.

Lookup bit (L): this is the Lookup bit used to indicate to the user of the ELP not to use this address for encapsulation but to look it up in the mapping database system to obtain an encapsulating RLOC address.

RLOC Probe bit (P): this is the RLOC Probe bit that means the Reencap Hop allows RLOC-probe messages to be sent to it. When the R bit is set to 0, RLOC-probes must not be sent. When a Reencap Hop is an anycast address then multiple physical Reencap Hops are using the same RLOC address. In this case, RLOC-probes are not needed because when the closest RLOC address is not reachable, another RLOC address can be reachable.

Strict bit (S): this is the Strict bit, which means the associated Reencap Hop is required to be used. If this bit is 0, the re-encapsulator can skip this Reencap Hop and go to the next one in the list.

AFI = x: x can be any AFI value from [AFN]. When a specific AFI has its own encoding of a multicast address, this field must be either a group address or a broadcast address.

Usage: This encoding can be used in RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. This encoding does not need to be understood by the mapping system for mapping database lookups, since this LCAF Type is not a lookup key.

4.7. Storing Security Data in the Mapping Database

When a locator in a locator-set has a security key associated with it, this LCAF will be used to encode key material. See [LISP-DDT] for details.

Security Key Canonical Address Format:

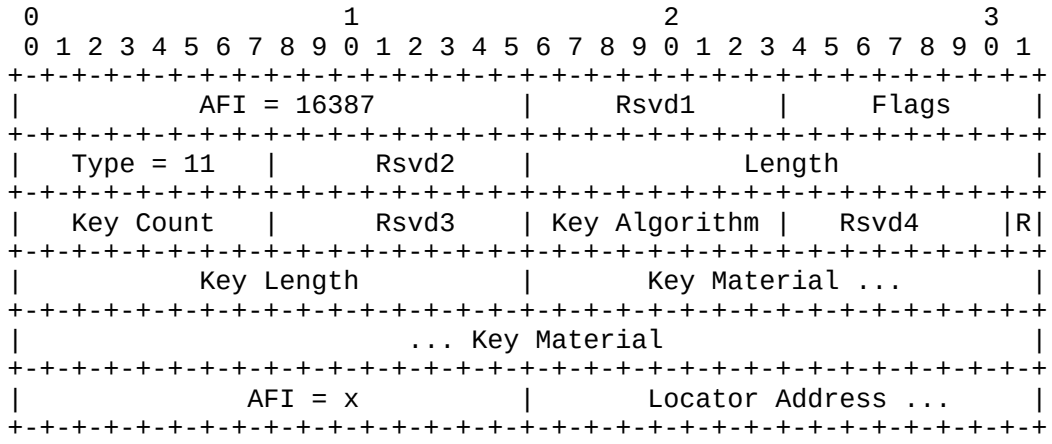

Length: length in bytes starting and including the byte after this Length field.

Key Count: the Key Count field declares the number of Key sections included in this LCAF. A Key section is made up of Key Length and Key Material fields.

Rsvd3: this field is reserved for future use and MUST be transmitted as 0 and ignored on receipt.

Key Algorithm: the Key Algorithm field identifies the key's cryptographic algorithm and specifies the format of the Public Key field. Refer to the [LISP-DDT] and [RFC8061] use cases for definitions of this field.

Rsvd4: this field is reserved for future use and MUST be transmitted as 0 and ignored on receipt.

R bit: this is the Revoke bit and, if set, it specifies that this key is being revoked.

Key Length: this field determines the length in bytes of the Key Material field.

Key Material: the Key Material field stores the key material. The format of the key material stored depends on the Key Algorithm field.

AFI = x: x can be any AFI value from [AFN]. This is the locator address that owns the encoded security key.

Usage: This encoding can be used in EID-records or RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages.

4.8. Source/Destination 2-Tuple Lookups

When both a source and destination address of a flow need consideration for different locator-sets, this 2-tuple key is used in EID fields in LISP control messages. When the Source/Dest key is registered to the mapping database, it can be encoded as a source-prefix and destination-prefix. When the Source/Dest is used as a key for a mapping database lookup, the source and destination come from a data packet.

Source/Dest Key Canonical Address Format:

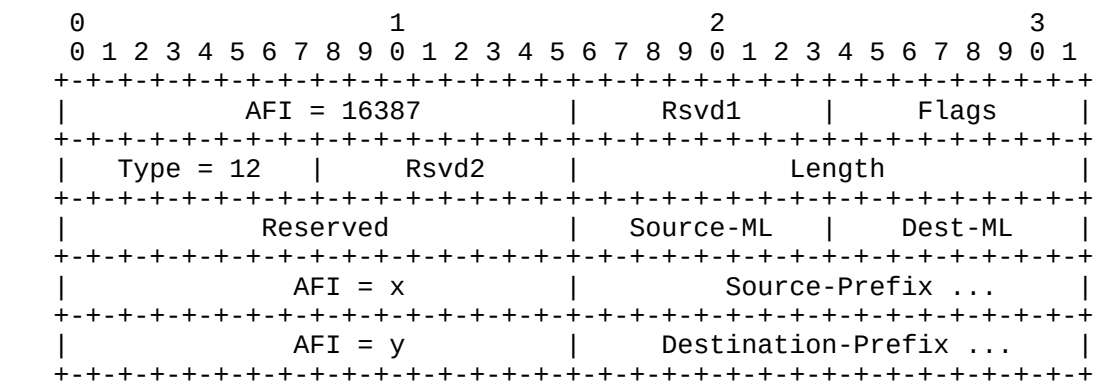

Length: length in bytes starting and including the byte after this Length field.

Reserved: must be set to zero and ignored on receipt.

Source-ML: the mask length of the source prefix that follows. The length is the number of high-order mask bits set.

Dest-ML: the mask length of the destination prefix that follows. The length is the number of high-order mask bits set.

AFI = x: x can be any AFI value from [AFN].

AFI = y: y can be any AFI value from [AFN]. When a specific address family has a multicast address semantic, this field must be either a group address or a broadcast address.

Usage: This encoding can be used in EID-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages. Refer to [LISP-TE] for usage details of this LCAF Type.

4.9. Replication List Entries for Multicast Forwarding

The Replication List Entry LCAF Type is an encoding for a locator being used for unicast replication according to the specification in [LISP-RE]. This locator encoding is pointed to by a Multicast Info LCAF Type and is registered by Re-encapsulating Tunnel Routers (RTRs) that are participating in an overlay distribution tree. Each RTR will register its locator address and its configured level in the distribution tree.

Replication List Entry Canonical Address Format:

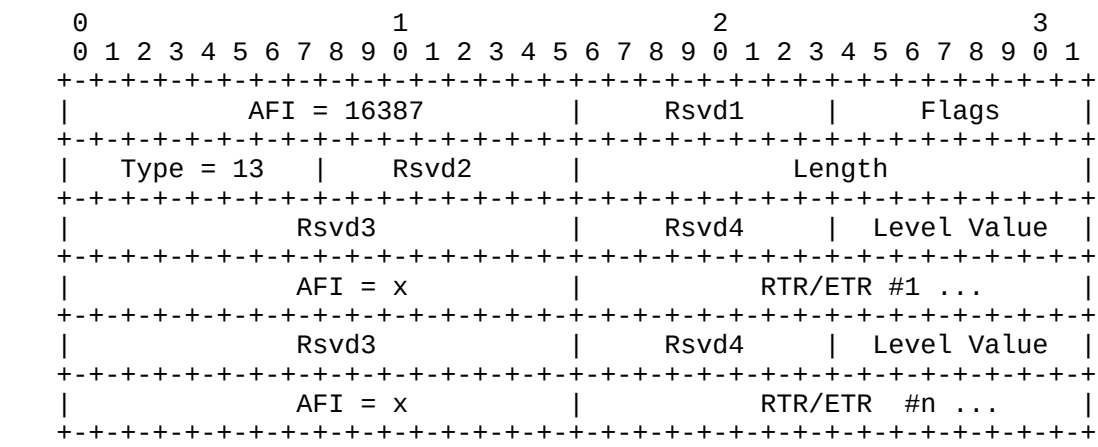

Length: length in bytes starting and including the byte after this Length field.

Rsvd3/Rsvd4: must be set to zero and ignored on receipt.

Level Value: this value is associated with the level within the overlay distribution tree hierarchy where the RTR resides. The level numbers are ordered from lowest value being close to the ITR (meaning that ITRs replicate to level-0 RTRs) and higher levels are further downstream on the distribution tree closer to ETRs of multicast receiver sites.

AFI = x: x can be any AFI value from [AFN]. A specific AFI has its own encoding of either a unicast or multicast locator address. For efficiency reasons, all RTR/ETR entries for the same level should be combined by a Map-Server to avoid searching through the entire multilevel list of locator entries in a Map-Reply message.

Usage: This encoding can be used in RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages.

4.10. Applications for AFI List LCAF Type

4.10.1. Binding IPv4 and IPv6 Addresses

When header translation between IPv4 and IPv6 is desirable, a LISP Canonical Address can use the AFI List LCAF Type to carry a variable number of AFIs in one LCAF AFI.

Address Binding LISP Canonical Address Format:

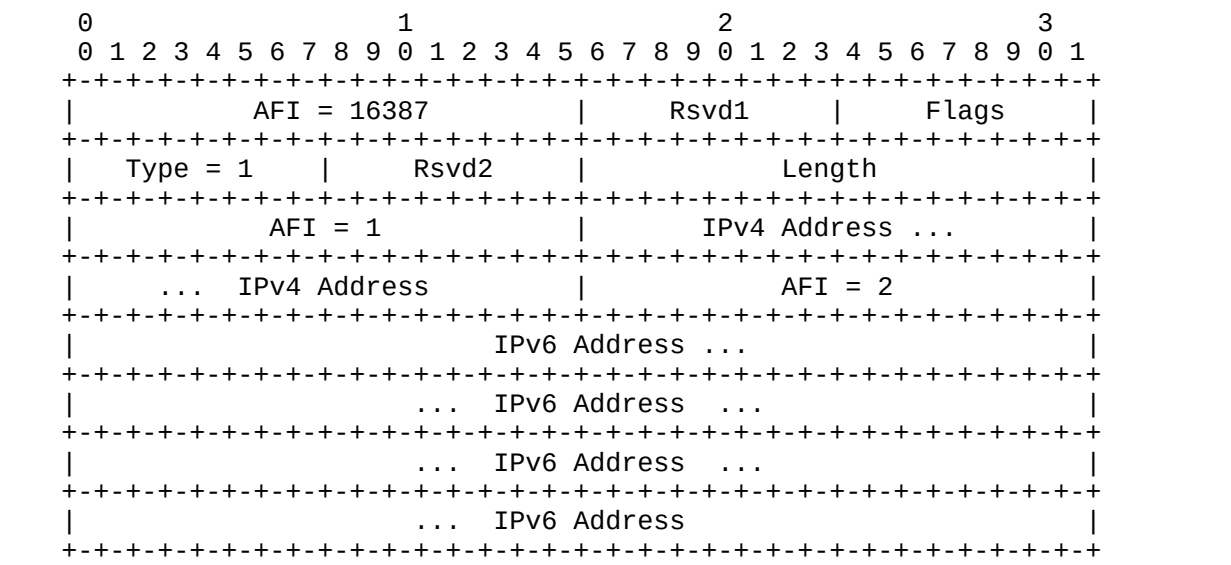

Length: length in bytes starting and including the byte after this Length field.

This type of address format can be included in a Map-Request when the address is being used as an EID, but the LISP Mapping Database System lookup destination can use only the IPv4 address. This is so a Mapping Database Service Transport System, such as LISP-ALT [RFC6836], can use the Map-Request destination address to route the control message to the desired LISP site.

Usage: This encoding can be used in EID-records or RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. See the other subsections in this section for specific use cases.

4.10.2. Layer 2 VPNs

When Media Access Control (MAC) addresses are stored in the LISP Mapping Database System, the AFI List LCAF Type can be used to carry AFI 6.

MAC Address LISP Canonical Address Format:

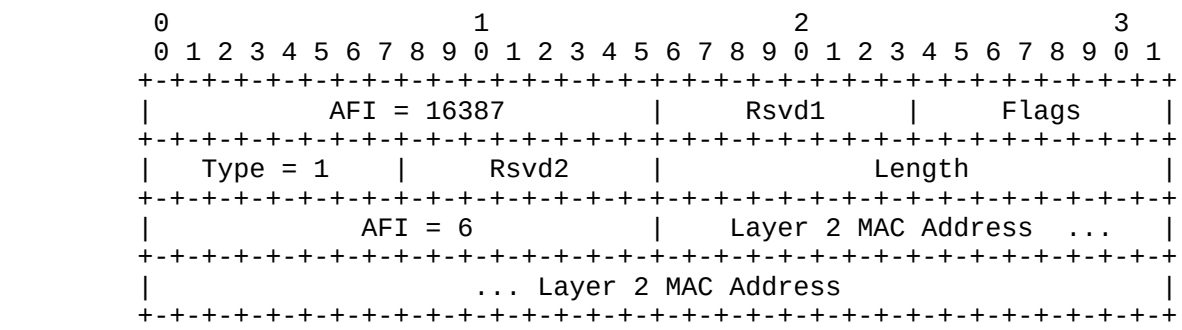

Length: length in bytes starting and including the byte after this Length field.

This address format can be used to connect Layer 2 domains together using LISP over an IPv4 or IPv6 core network to create a Layer 2 VPN. In this use case, a MAC address is being used as an EID, and the locator-set that this EID maps to can be an IPv4 or IPv6 RLOC, or even another MAC address being used as an RLOC. See [EID-MOBILITY] for how Layer 2 VPNs operate when doing EID mobility.

Care should be taken to protect privacy against the adverse use of a Layer 2 MAC address by ensuring policy controls are used during EID registrations that use AFI=6 encodings in RLOC-records. Refer to the use-case documents for additional information.

4.10.3. ASCII Names in the Mapping Database

If DNS names [RFC1035] or URIs [RFC3986] are stored in the LISP Mapping Database System, the AFI List LCAF Type can be used to carry an ASCII string.

ASCII LISP Canonical Address Format:

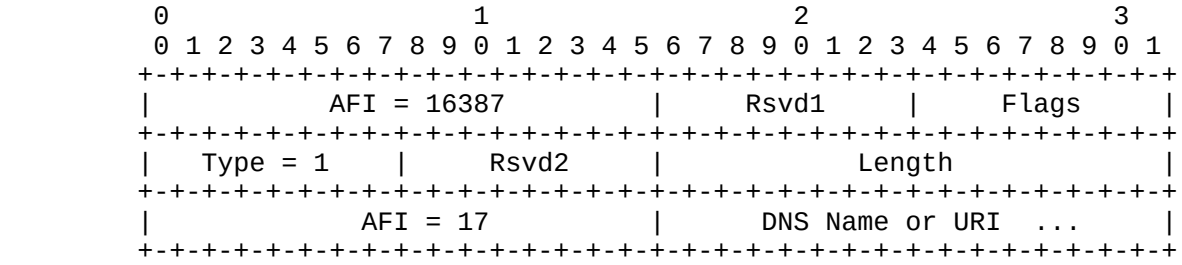

Length: length in bytes starting and including the byte after this Length field.

An example for using DNS names is when an ETR registers a mapping with an EID-record encoded as (AFI=1, 10.0.0.0/8) with an RLOC-record (AFI=17, "router.abc.com").

4.10.4. Using Recursive LISP Canonical Address Encodings

When any combination of above is desirable, the AFI List LCAF Type value can be used to carry within the LCAF AFI another LCAF AFI (for example, Application-Specific Data in Section 5.1).

Recursive LISP Canonical Address Format:

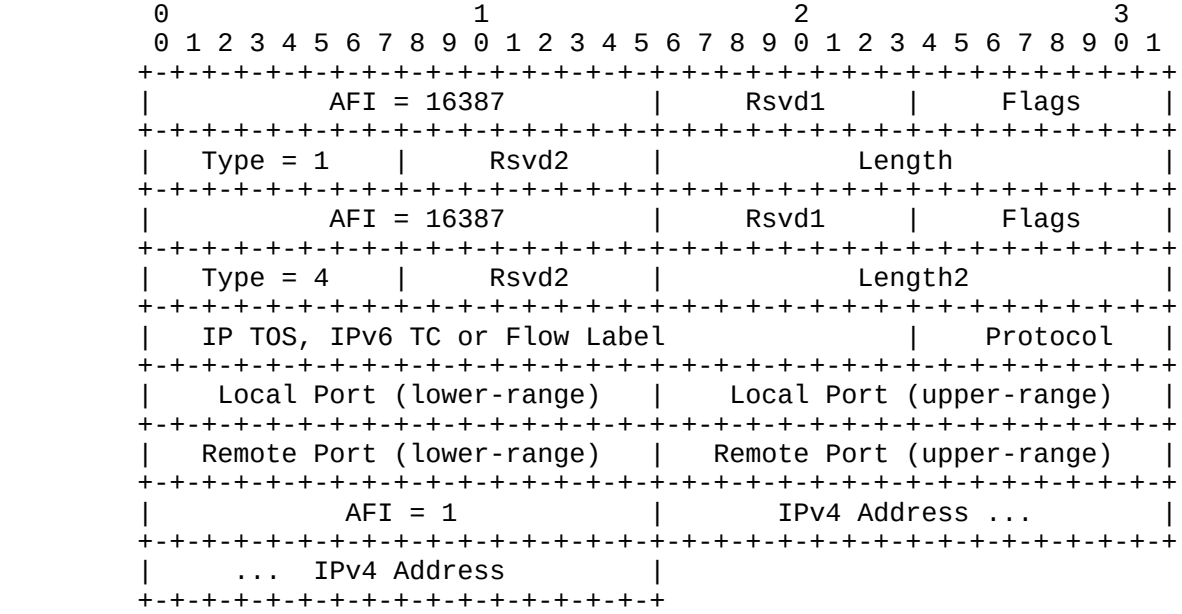

Length: length in bytes starting and including the byte after this Length field.

Length2: length in bytes starting and including the byte after this Length2 field.

This format could be used by a Mapping Database Service Transport System, such as LISP-ALT [RFC6836], where the AFI=1 IPv4 address is used as an EID and placed in the Map-Request destination address by the sending LISP system. The ALT system can deliver the Map-Request to the LISP destination site independent of the Application Data LCAF

Type AFI payload values. When this AFI is processed by the destination LISP site, it can return different locator-sets based on the type of application or level of service that is being requested.

4.10.5. Compatibility Mode Use Case

A LISP system should use the AFI List LCAF Type format when sending to LISP systems that do not support a particular LCAF Type used to encode locators. This allows the receiving system to be able to parse a locator address for encapsulation purposes. The list of AFIs in an AFI List LCAF Type has no semantic ordering and a receiver should parse each AFI element no matter what the ordering.

Compatibility Mode Address Format:

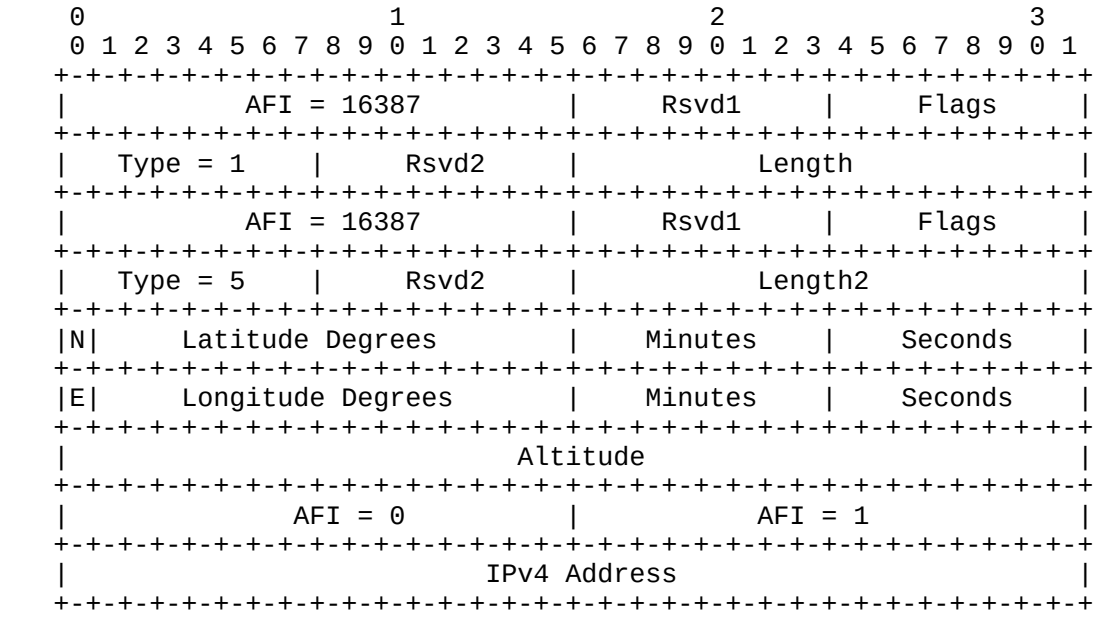

Length: length in bytes starting and including the byte after this Length field.

Length2: length in bytes starting and including the byte after this Length2 field.

If a system does not recognized the Geo-Coordinates LCAF Type that is accompanying a locator address, an encoder can include the Geo-Coordinates LCAF Type embedded in an AFI List LCAF Type where the AFI

in the Geo-Coordinates LCAF Type is set to 0 and the AFI encoded next in the list is encoded with a valid AFI value to identify the locator address.

A LISP system is required to support the AFI List LCAF Type to use this procedure. It would skip over 10 bytes of the Geo-Coordinates LCAF Type to get to the locator address encoding (an IPv4 locator address). A LISP system that does support the Geo-Coordinates LCAF Type can support parsing the locator address within the Geo-Coordinates LCAF Type encoding or in the locator encoding that follows in the AFI List LCAF Type.

5. Experimental LISP Canonical Address Applications

The following sections describe experimental LCAF encodings. These LCAF Types are not approved (i.e., not registered with IANA). The inclusion of these encodings in this document is in support of further study and experimentation to determine whether these encodings are functional, if there is a demand for these use cases, and to better understand deployment considerations. As noted previously, these LCAF Types are restricted to cautious use in self-contained environments in support of the corresponding use-case documents.

5.1. Convey Application-Specific Data

When a locator-set needs to be conveyed based on the type of application or the Per-Hop Behavior (PHB) of a packet, the Application Data LCAF Type can be used.

Application Data LISP Canonical Address Format:

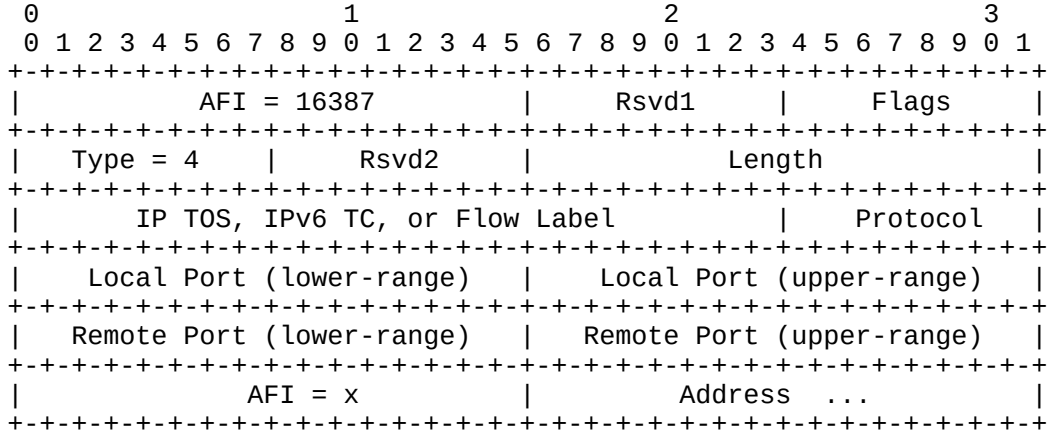

Length: length in bytes starting and including the byte after this Length field.

IP TOS, IPv6 TC, or Flow Label: this field stores the 8-bit IPv4 TOS field used in an IPv4 header, the 8-bit IPv6 Traffic Class or Flow Label used in an IPv6 header.

Local Port/Remote Port Ranges: these fields are from the TCP, UDP, or Stream Control Transmission Protocol (SCTP) transport header. A range can be specified by using a lower value and an upper value. When a single port is encoded, the lower and upper value fields are the same.

AFI = x: x can be any AFI value from [AFN].

The Application Data LCAF Type is used for an EID encoding when an ITR wants a locator-set for a specific application. When used for an RLOC encoding, the ETR is supplying a locator-set for each specific application is has been configured to advertise.

Usage: This encoding can be used in EID-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. When LISP-DDT [LISP-DDT] is used as the mapping system mechanism, extended EIDs are used in Map-Referral messages. This LCAF Type is used as a lookup key to the mapping system that can return a longest-match or exact-match entry.

5.2. Generic Database Mapping Lookups

When the LISP Mapping Database System holds information accessed by a generic formatted key (where the key is not the usual IPv4 or IPv6 address), an opaque key may be desirable.

Opaque Key LISP Canonical Address Format:

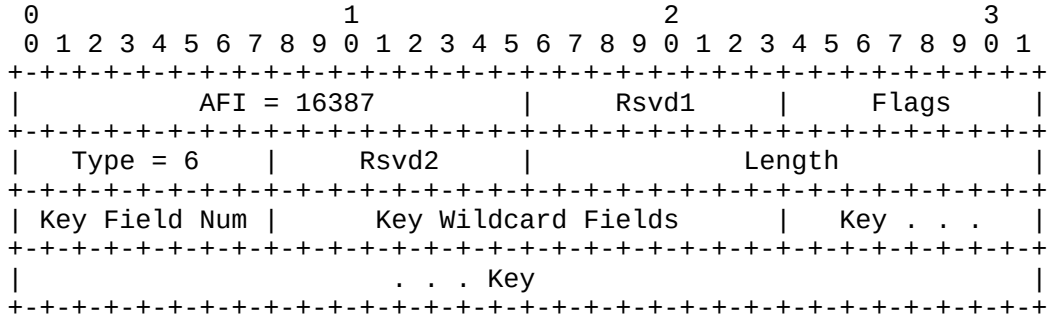

Length: length in bytes starting and including the byte after this Length field.

Key Field Num: the value of this field is the number of "Key" sub-fields minus 1, the Key field can be broken up into. So, if this field has a value of 0, there is one sub-field in the "Key". The width of the sub-fields are fixed length. So, for a key size of 8 bytes, with a Key Field Num of 3, four sub-fields of 2 bytes each in length are allowed. Allowing for a reasonable number of 16 sub-field separators, valid values range from 0 to 15.

Key Wildcard Fields: describes which fields in the key are not used as part of the key lookup. This wildcard encoding is a bitfield. Each bit is a don’t-care bit for a corresponding field in the key. Bit 0 (the low-order bit) in this bitfield corresponds the first field, the low-order field in the key, bit 1 the second field, and so on. When a bit is set in the bitfield, it is a don’t-care bit and should not be considered as part of the database lookup. When the entire 16 bits are set to 0, then all bits of the key are used for the database lookup.

Key: the variable length key used to do a LISP Mapping Database System lookup. The length of the key is the value n (as shown above).

Usage: This is an experimental Type where the usage has not yet been defined.

5.3. PETR Admission Control Functionality

When a public Proxy Egress Tunnel Router (PETR) device wants to verify who is encapsulating to it, it can check for a specific nonce value in the LISP-encapsulated packet. To convey the nonce to admitted ITRs or PITRs, this LCAF is used in a Map-Register or Map-Reply locator-record.

Nonce Locator Canonical Address Format:

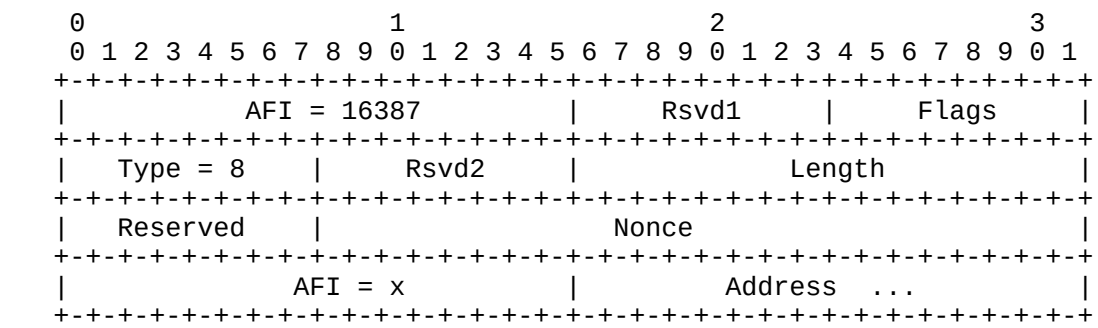

Length: length in bytes starting and including the byte after this Length field.

Reserved: must be set to zero and ignored on receipt.

Nonce: a nonce value returned by an ETR in a Map-Reply locator-record to be used by an ITR or PITR when encapsulating to the locator address encoded in the AFI field of this LCAF Type. This nonce value is inserted in the nonce field in the LISP header encapsulation.

AFI = x: x can be any AFI value from [AFN].

Usage: This is an experimental Type where the usage has not yet been defined.

5.4. Data Model Encoding

This Type allows a JSON data model to be encoded as either an EID or an RLOC.

JSON Data Model Type Address Format:

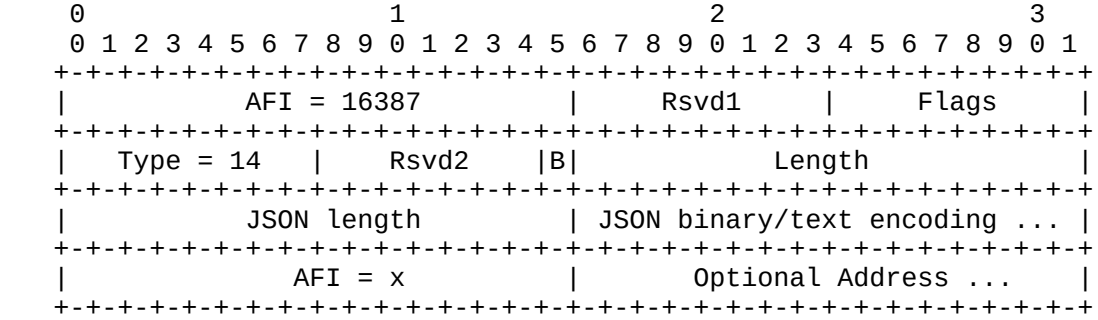

Length: length in bytes starting and including the byte after this Length field.

B bit: indicates that the JSON field is binary encoded according to [JSON-BINARY] when the bit is set to 1. Otherwise, the encoding is based on text encoding according to [RFC7159].

JSON length: length in octets of the following JSON binary/text encoding field.

JSON binary/text encoding: a variable-length field that contains either binary or text encodings.

AFI = x: x can be any AFI value from [AFN]. A specific AFI has its own encoding of either a unicast or multicast locator address. All RTR/ETR entries for the same level should be combined by a Map-Server to avoid searching through the entire multilevel list of locator entries in a Map-Reply message.

Usage: This is an experimental Type where the usage has not yet been defined. An example mapping is an EID-record encoded as a distinguished-name "cpe-router" and an RLOC-record encoded as a JSON string "{ "router-address" : "1.1.1.1", "router-mask" : "8" }".

5.5. Encoding Key/Value Address Pairs

The Key/Value pair is, for example, useful for attaching attributes to other elements of LISP packets, such as EIDs or RLOCs. When attaching attributes to EIDs or RLOCs, it’s necessary to distinguish between the element that should be used as EID or RLOC and, hence, as the key for lookups and additional attributes. This is especially the case when the difference cannot be determined from the Types of the elements, such as when two IP addresses are being used.

Key/Value Address Pair Address Format:

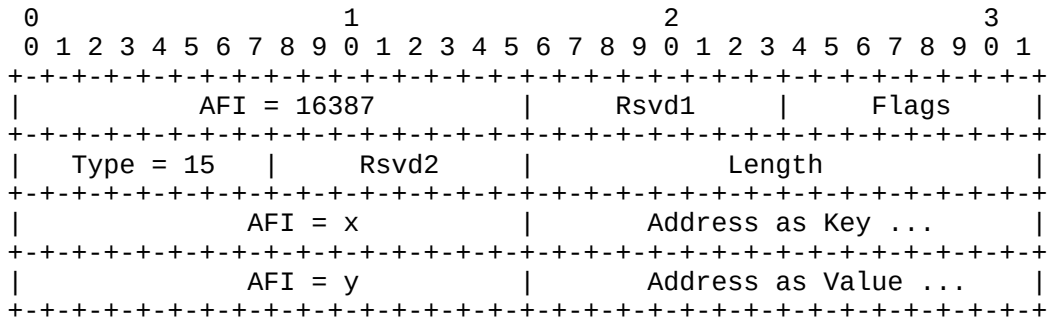

Length: length in bytes starting and including the byte after this Length field.

AFI = x: x is the "Address as Key" AFI that can have any value from [AFN]. A specific AFI has its own encoding of either a unicast or a multicast locator address. All RTR/ETR entries for the same level should be combined by a Map-Server to avoid searching through the entire multilevel list of locator entries in a Map-Reply message.

Address as Key: AFI-encoded address that will be attached with the attributes encoded in "Address as Value", which follows this field.

AFI = y: y is the "Address of Value" AFI that can have any value from [AFN]. A specific AFI has its own encoding of either a unicast or a multicast locator address. All RTR/ETR entries for the same level should be combined by a Map-Server to avoid searching through the entire multilevel list of locator entries in a Map-Reply message.

Address as Value: AFI-encoded address that will be the attribute address that goes along with "Address as Key" which precedes this field.

Usage: This is an experimental Type where the usage has not yet been defined.

5.6. Multiple Data-Planes

Overlays are becoming popular in many parts of the network, which has created an explosion of data-plane encapsulation headers. Since the LISP mapping system can hold many types of address formats, it can represent the encapsulation format supported by an RLOC as well. When an encapsulator receives a Map-Reply with an Encapsulation Format LCAF Type encoded in an RLOC-record, it can select an encapsulation format, that it can support, from any of the encapsulation protocols that have the bit set to 1 in this LCAF Type.

Encapsulation Format Address Format:

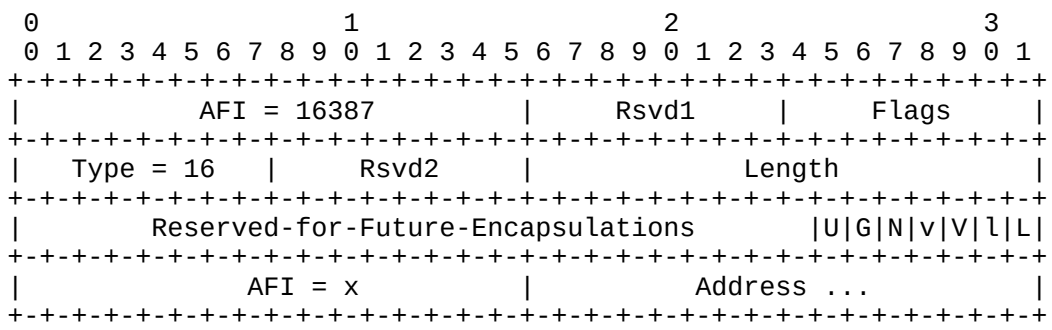

Length: length in bytes starting and including the byte after this Length field.

Reserved-for-Future-Encapsulations: must be set to zero and ignored on receipt. This field will get bits allocated to future encapsulations, as they are created.

U: The RLOCs listed in the AFI-encoded addresses in the next longword can accept Generic UDP Encapsulation (GUE) using destination UDP port 6080 [GUE].

G: The RLOCs listed in the AFI-encoded addresses in the next longword can accept Geneve encapsulation using destination UDP port 6081 [GENEVE].

N: The RLOCs listed in the AFI-encoded addresses in the next longword can accept NV-GRE (Network Virtualization - Generic Routing Encapsulation) using IPv4/IPv6 protocol number 47 [RFC7637].

v: The RLOCs listed in the AFI-encoded addresses in the next longword can accept VXLAN-GPE (Generic Protocol Extension) encapsulation using destination UDP port 4790 [GPE-VXLAN].

V: The RLOCs listed in the AFI-encoded addresses in the next longword can accept Virtual eXtensible Local Area Network (VXLAN) encapsulation using destination UDP port 4789 [RFC7348].

l: The RLOCs listed in the AFI-encoded addresses in the next longword can accept Layer 2 LISP encapsulation using destination UDP port 8472 [LISP-L2].

L: The RLOCs listed in the AFI-encoded addresses in the next longword can accept Layer 3 LISP encapsulation using destination UDP port 4341 [RFC6830].

Usage: This encoding can be used in RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages.

6. Security Considerations

This document is classified as Experimental. The LCAF encodings defined in this document are intended to be used with their corresponding use cases and in self-contained environments. Users should carefully consider how the [LISP-SEC] threat model applies to their particular use case.

The use of the Geo-Coordinates LCAF Type may raise physical privacy issues. Care should be taken when configuring the mapping system to use specific policy parameters so geolocation information is not returned gratuitously. It is recommended that any documents that specify the use of the Geo-Coordinates LCAF Type should consider the applicability of RFC 6280 (BCP 160) [RFC6280] for location-based privacy protection.

Additional privacy concerns have arisen since publication of BCP 160, and future work on LISP should examine potential threats beyond BCP 160 and address improving privacy and security for LISP deployments.

7. IANA Considerations

This document defines a canonical address format encoding used in LISP control messages and in the encoding of lookup keys for the LISP Mapping Database System. Such an address format is based on a fixed AFI (16387) and a LISP LCAF Type field.

The LISP LCAF Type field is an 8-bit field specific to the LISP Canonical Address Format encodings. IANA has created a new registry (as outlined in [RFC5226]) titled "LISP Canonical Address Format (LCAF) Types". Initial values for the "LISP Canonical Address Format (LCAF) Types" registry are given below. Future assignments are to be made using the Specification Required policy [RFC5226]. Assignments consist of a LISP LCAF Type Name and its associated value:

Value	LISP LCAF Type Name	Reference
0	Null Body	Section 3
1	AFI List	Section 3
2	Instance ID	Section 3
3	AS Number	Section 3
5	Geo-Coordinates	Section 3
7	NAT-Traversal	Section 3
9	Multicast Info	Section 3
10	Explicit Locator Path	Section 3
11	Security Key	Section 3
12	Source/Dest Key	Section 3
13	Replication List Entry	Section 3

Table 1: Initial Values in the
"LISP Canonical Address Format (LCAF) Types" Registry

8. References

8.1. 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>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3232] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, DOI 10.17487/RFC3232, January 2002, <http://www.rfc-editor.org/info/rfc3232>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <http://www.rfc-editor.org/info/rfc3986>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, DOI 10.17487/RFC6280, July 2011, <http://www.rfc-editor.org/info/rfc6280>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <http://www.rfc-editor.org/info/rfc6830>.

[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>.

[RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <http://www.rfc-editor.org/info/rfc7159>.

[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>.

[RFC7637] Garg, P., Ed. and Y. Wang, Ed., "NVGRE: Network Virtualization Using Generic Routing Encapsulation", RFC 7637, DOI 10.17487/RFC7637, September 2015, <http://www.rfc-editor.org/info/rfc7637>.

8.2. Informative References

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

[EID-MOBILITY] Portoles-Comeras, M., Ashtaputre, V., Moreno, V., Maino, F., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", Work in Progress, draft-portoles-lisp-eid-mobility-01, October 2016.

[GENEVE] Gross, J., Ganga, I., and T. Sridhar, "Geneve: Generic Network Virtualization Encapsulation", Work in Progress, draft-ietf-nvo3-geneve-03, September 2016.

[GPE-VXLAN] Maino, F., Kreeger, L., and U. Elzur, "Generic Protocol Extension for VXLAN", Work in Progress, draft-ietf-nvo3-vxlan-gpe-03, October 2016.

[GUE] Herbert, T., Yong, L., and O. Zia, "Generic UDP Encapsulation", Work in Progress, draft-ietf-nvo3-gue-05, October 2016.

[JSON-BINARY] "Universal Binary JSON Specification", <http://ubjson.org>.

[LISP-DDT] Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "LISP Delegated Database Tree", Work in Progress, draft-ietf-lisp-ddt-09, January 2017.

[LISP-L2] Smith, M., Dutt, D., Farinacci, D., and F. Maino, "Layer 2 (L2) LISP Encapsulation Format", Work in Progress, draft-smith-lisp-layer2-03, September 2013.

[LISP-RE] Coras, F., Cabellos-Aparicio, A., Domingo-Pascual, J., Maino, F., and D. Farinacci, "LISP Replication Engineering", Work in Progress, draft-coras-lisp-re-08, November 2015.

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

[LISP-TE] Farinacci, D., Kowal, M., and P. Lahiri, "LISP Traffic Engineering Use-Cases", Work in Progress, draft-farinacci-lisp-te-11, September 2016.

[NAT-LISP] Ermagan, V., Farinacci, D., Lewis, D., Skriver, J., Maino, F., and C. White, "NAT traversal for LISP", Work in Progress, draft-ermagan-lisp-nat-traversal-11, August 2016.

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

[WGS-84] National Imagery and Mapping Agency, "Department of Defense World Geodetic System 1984", NIMA TR8350.2, January 2000, <http://earth-info.nga.mil/GandG/publications/tr8350.2/wgs84fin.pdf>.

Acknowledgments

The authors would like to thank Vince Fuller, Gregg Schudel, Jesper Skriver, Luigi Iannone, Isidor Kouvelas, and Sander Steffann for their technical and editorial commentary.

The authors would like to thank Victor Moreno for discussions that led to the definition of the Multicast Info LCAF Type.

The authors would like to thank Parantap Lahiri and Michael Kowal for discussions that led to the definition of the Explicit Locator Path (ELP) LCAF Type.

The authors would like to thank Fabio Maino and Vina Ermagan for discussions that led to the definition of the Security Key LCAF Type.

The authors would like to thank Albert Cabellos-Aparicio and Florin Coras for discussions that led to the definition of the Replication List Entry LCAF Type.

Thanks goes to Michiel Blokzijl and Alberto Rodriguez-Natal for suggesting new LCAF Types.

Thanks also goes to Terry Manderson for assistance obtaining a LISP AFI value from IANA.

And finally, the authors thank Stephen Farrell (Security Area Director) and Deborah Brungard (Routing Area Director) for their suggested text to get the document through IESG review.

Authors' Addresses

Dino Farinacci
lispers.net
San Jose, CA
United States of America

Email: farinacci@gmail.com

Dave Meyer
Brocade
San Jose, CA
United States of America

Email: dmm@1-4-5.net

Job Snijders
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
The Netherlands

Email: job@ntt.net