Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 26, 2012

D. Farinacci
D. Lewis
D. Meyer
cisco Systems
C. White
Logical Elegance, LLC.
October 24, 2011

LISP Mobile Node
draft-meyer-lisp-mn-06

Abstract

This document describes how a lightweight version of LISP's ITR/ETR functionality can be used to provide seamless mobility to a mobile node. The LISP Mobile Node design described in this document uses standard LISP functionality to provide scalable mobility for LISP mobile nodes.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 26, 2012.

Copyright Notice

Table of Contents

1. Introduction

The Locator/ID Separation Protocol (LISP) [LISP] specifies a design and mechanism for replacing the addresses currently used in the Internet with two separate name spaces: Endpoint Identifiers (EIDs), used within sites, and Routing Locators (RLOCs), used by the transit networks that make up the Internet infrastructure. To achieve this separation, LISP defines protocol mechanisms for mapping from EIDs to RLOCs. The mapping infrastructure is comprised of LISP Map-Servers and Map-Resolvers [LISP-MS] and is tied together with LISP+ALT [LISP-ALT].

This document specifies the behavior of a new LISP network element: the LISP Mobile Node. The LISP Mobile Node implements a subset of the standard Ingress Tunnel Router and Egress Tunnel Router functionality [LISP]. Design goals for the LISP mobility design include:

- Allowing TCP connections to stay alive while roaming.
- Allowing the mobile node to communicate with other mobile nodes while either or both are roaming.
- Allowing the mobile node to multi-home (i.e., use multiple interfaces concurrently).
- Allowing the mobile node to be a server. That is, any mobile node or stationary node can find and connect to a mobile node as a server.
- Providing shortest path bidirectional data paths between a mobile node and any other stationary or mobile node.
- Not requiring fine-grained routes in the core network to support mobility.
- Not requiring a home-agent, foreign agent or other data plane network elements to support mobility. Note since the LISP mobile node design does not require these data plane elements, there is no triangle routing of data packets as is found in Mobile IP [RFC3344].
- Not requiring new IPv6 extension headers to avoid triangle routing [RFC3775].

The LISP Mobile Node design requires the use of the LISP Map-Server [LISP-ALT] and LISP Interworking [LISP-INTERWORK] technology to allow a LISP mobile node to roam and to be discovered in an efficient and

scalable manner. The use of Map-Server technology is discussed further in Section 5.

The protocol mechanisms described in this document apply those cases in which a node's IP address changes frequently. For example, when a mobile node roams, it is typically assigned a new IP address. Similarly, a broadband subscriber may have its address change frequently; as such, a broadband subscriber can use the LISP Mobile Node mechanisms defined in this specification.

The remainder of this document is organized as follows: Section 2 defines the terms used in this document. Section 3 provides a overview of salient features of the LISP Mobile Node design, and Section 4 describes design requirements for a LISP Mobile Node. Section 5 provides the detail of LISP Mobile Node data and control plane operation, and Section 6 discusses options for updating remote caches in the presence of unidirectional traffic flows. Section 7 specifies how the LISP Mobile Node protocol operates. Section 8 specifies multicast operation for LISP mobile nodes. Section 9 and Section 11 outline other considerations for the LISP-MN design and implementation. Finally, Section 12 outlines the security considerations for a LISP mobile node.

2. Definition of Terms

This section defines the terms used in this document.

Stationary Node (SN): A non-mobile node who's IP address changes infrequently. That is, its IP address does not change as frequently as a fast roaming mobile hand-set or a broadband connection and therefore the EID to RLOC mapping is relatively static.

Endpoint ID (EID): This is the traditional LISP EID [LISP], and is the address that a LISP mobile node uses as its address for transport connections. A LISP mobile node never changes its EID, which is typically a /32 or /128 prefix and is assigned to a loopback interface. Note that the mobile node can have multiple EIDs, and these EIDs can be from different address families.

Routing Locator (RLOC): This is the traditional LISP RLOC, and is in general a routable address that can be used to reach a mobile node. Note that there are cases in which an mobile node may receive an address that it thinks is an RLOC (perhaps via DHCP) which is either an EID or an RFC 1918 address [RFC1918]. This could happen if, for example, if the mobile node roams into a LISP domain or a domain behind a Network Address Translator (NAT)).

Ingress Tunnel Router (ITR): An ITR is a router that accepts an IP packet with a single IP header (more precisely, an IP packet that does not contain a LISP header). The router treats this "inner" IP destination address as an EID and performs an EID-to-RLOC mapping lookup. The router then prepends an "outer" IP header with one of its globally routable RLOCs in the source address field and the result of the mapping lookup in the destination address field. Note that this destination RLOC may be an intermediate, proxy device that has better knowledge of the EID-to-RLOC mapping closer to the destination EID. In general, an ITR receives IP packets from site end-systems on one side and sends LISP-encapsulated IP packets toward the Internet on the other side. A LISP mobile node, however, when acting as an ITR LISP encapsulates all packet that it originates.

Egress Tunnel Router (ETR): An ETR is a router that accepts an IP packet where the destination address in the "outer" IP header is one of its own RLOCs. The router strips the "outer" header and forwards the packet based on the next IP header found. In general, an ETR receives LISP-encapsulated IP packets from the Internet on one side and sends decapsulated IP packets to site end-systems on the other side. A LISP mobile node, when acting as an ETR, decapsulates packets that are then typically processed by the mobile node.

Proxy Ingress Tunnel Router (PITR): PITRs are used to provide interconnectivity between sites that use LISP EIDs and those that do not. They act as a gateway between the Legacy Internet and the LISP enabled Network. A given PITR advertises one or more highly aggregated EID prefixes into the public Internet and acts as the ITR for traffic received from the public Internet. Proxy Ingress Tunnel Routers are described in [LISP-INTERWORK].

Proxy Egress Tunnel Router (PETR): An infrastructure element used to decapsulate packets sent from mobile nodes to non-LISP sites. Proxy Egress Tunnel Routers are described in [LISP-INTERWORK].

LISP Mobile Node (LISP-MN): A LISP capable fast roaming mobile hand-set.

Map-cache: A data structure which contains an EID-prefix, its associated RLOCs, and the associated policy. Map-caches are typically found in ITRs and PITRs.

Negative Map-Reply: A Negative Map-Reply is a Map-Reply that contains a coarsely aggregated non-LISP prefix. Negative Map-Replies are typically generated by Map-Resolvers, and are used to inform an ITR (mobile or stationary) that a site is not a LISP

site. A LISP mobile node encapsulate packets to destinations covered by the negative Map-Reply are encapsulated to a PETR.

Roaming Event: A Roaming Event occurs when there is a change in a LISP mobile node's RLOC set.

3. Design Overview

The LISP-MN design described in this document uses the Map-Server/ Map-Resolver service interface in conjunction with a light-weight ITR/ETR implementation in the LISP-MN to provide scalable fast mobility. The LISP-MN control-plane uses a Map-Server as an anchor point, which provides control-plane scalability. In addition, the LISP-MN data-plane takes advantage of shortest path routing and therefore does not increase packet delivery latency.

4. Design Requirements

This section outlines the design requirements for a LISP-MN, and is divided into User Requirements (Section 4.1) and Network Requirements (Section 4.2).

4.1. User Requirements

This section describes the user-level functionality provided by a LISP-MN.

Transport Connection Survivability: The LISP-MN design must allow a LISP-MN to roam while keeping transport connections alive.

Simultaneous Roaming: The LISP-MN design must allow a LISP-MN to talk to another LISP-MN while both are roaming.

Multihoming: The LISP-MN design must allow for simultaneous use of multiple Internet connections by a LISP-MN. In addition, the design must allow for the LISP mobile node to specify ingress traffic engineering policies as documented in [LISP]. That is, the LISP-MN must be able to specify both active/active and active/ passive policies for ingress traffic.

Shortest Path Data Plane: The LISP-MN design must allow for shortest path bidirectional traffic between a LISP-MN and a stationary node, and between a LISP-MN and another LISP-MN (i.e., without triangle routing in the data path). This provides a low-latency data path between the LISP-MN and the nodes that it is communicating with.

4.2. Network Requirements

This section describes the network functionality that the LISP-MN design provides to a LISP-MN.

Routing System Scalability: The LISP-MN design must not require injection of fine-grained routes into the core network.

Mapping System Scalability: The LISP-MN design must not require additional state in the mapping system. In particular, any mapping state required to support LISP mobility must BE confined to the LISP-MN's Map-Server and the ITRs which are talking to the LISP-MN.

Component Reuse: The LISP-MN design must use existing LISP infrastructure components. These include map server, map resolver, and interworking infrastructure components.

Home Agent/Foreign Agent: The LISP-MN design must not require the use of home-agent or foreign-agent infrastructure components [RFC3344].

Readdressing: The LISP-MN design must not require TCP connections to be reset when the mobile node roams. In particular, since the IP address associated with a transport connection will not change as the mobile node roams, TCP connections will not reset.

5. LISP Mobile Node Operation

The LISP-MN design is built from three existing LISP components: A lightweight LISP implementation that runs in an LISP-MN, and the existing Map-Server [LISP-MS] and Interworking [LISP-INTERWORK] infrastructures. A LISP mobile node typically sends and receives LISP encapsulated packets (exceptions include management protocols such as DHCP).

The LISP-MN design makes a single mobile node look like a LISP site as described in in [LISP] by implementing ITR and ETR functionality. Note that one subtle difference between standard ITR behavior and LISP-MN is that the LISP-MN encapsulates all non-local, non-LISP site destined outgoing packets to a PETR.

When a LISP-MN roams onto a new network, it receives a new RLOC. Since the LISP-MN is the authoritative ETR for its EID-prefix, it must Map-Register it's updated RLOC set. New sessions can be established as soon as the registration process completes. Sessions that are encapsulating to RLOCs that did not change during the

roaming event are not affected by the roaming event (or subsequent mapping update). However, the LISP-MN must update the ITRs and PITRs that have cached a previous mapping. It does this using the techniques described in Section 6.

5.1. Addressing Architecture

A LISP-MN is typically provisioned with one or more EIDs that it uses for all transport connections. LISP-MN EIDs are provisioned from blocks reserved from mobile nodes much the way mobile phone numbers are provisioned today (such that they do not overlap with the EID space of any enterprise). These EIDs can be either IPv4 or IPv6 addresses. For example, one EID might be for a public network while another might be for a private network; in this case the "public" EID will be associated with RLOCs from the public Internet, while the "private" EID will be associated with private RLOCs. It is anticipated that these EIDs will change infrequently if at all, since the assignment of a LISP-MN's EID is envisioned to be a subscription time event. The key point here is that the relatively fixed EID allows the LISP-MN's transport connections to survive roaming events. In particular, while the LISP-MN's EIDs are fixed during roaming events, the LISP-MN's RLOC set will change. The RLOC set may be comprised of both IPv4 or IPv6 addresses.

A LISP-MN is also provisioned with the address of a Map-Server and a corresponding authentication key. Like the LISP-MN's EID, both the Map-Server address and authentication key change very infrequently (again, these are anticipated to be subscription time parameters). Since the LISP LISP-MN's Map-Server is configured to advertise an aggregated EID-prefix that covers the LISP-MN's EID, changes to the LISP-MN's mapping are not propagated further into the mapping system [LISP-ALT]. It is this property that provides for scalable fast mobility.

A LISP-MN is also be provisioned with the address of a Map-Resolver. A LISP-MN may also learn the address of a Map-Resolver though a dynamic protocol such as DHCP [RFC2131].

Finally, note that if, for some reason, a LISP-MN's EID is re-provisioned, the LISP-MN's Map-Server address may also have to change in order to keep LISP-MN's EID within the aggregate advertised by the Map-Server (this is discussed in greater detail in Section 5.2).

5.2. Control Plane Operation

A roaming event occurs when the LISP-MN receives a new RLOC. Because the new address is a new RLOC from the LISP-MN's perspective, it must update its EID-to-RLOC mapping with its Map-Server; it does this

using the Map-Register mechanism described in [LISP].

A LISP-MN may want the Map-Server to respond on its behalf for a variety of reasons, including minimizing control traffic on radio links and minimizing battery utilization. A LISP-MN may instruct its Map-Server to proxy respond to Map-Requests by setting the Proxy-Map-Reply bit in the Map-Register message [LISP]. In this case the Map-Server responds with a non-authoritative Map-Reply so that an ITR or PITR will know that the ETR didn't directly respond. A Map-Server will proxy reply only for "registered" EID-prefixes using the registered EID-prefix mask-length in proxy replies.

Because the LISP-MN's Map-Server is pre-configured to advertise an aggregate covering the LISP-MN's EID prefix, the database mapping change associated with the roaming event is confined to the Map-Server and those ITRs and PITRs that may have cached the previous mapping.

5.3. Data Plane Operation

A key feature of LISP-MN control-plane design is the use of the Map-Server as an anchor point; this allows control of the scope to which changes to the mapping system must be propagated during roaming events.

On the other hand, the LISP-MN data-plane design does not rely on additional LISP infrastructure for communication between LISP nodes (mobile or stationary). Data packets take the shortest path to and from the LISP-MN to other LISP nodes; as noted above, low latency shortest paths in the data-plane is an important goal for the LISP-MN design (and is important for delay-sensitive applications like gaming and voice-over-IP). Note that a LISP-MN will need additional interworking infrastructure when talking to non-LISP sites [LISP-INTERWORK]; this is consistent with the design of any host at a LISP site which talks to a host at a non-LISP site.

In general, the LISP-MN data-plane operates in the same manner as the standard LISP data-plane with one exception: packets generated by a LISP-MN which are not destined for the mapping system (i.e., those sent to destination UDP port 4342) or the local network are LISP encapsulated. Because data packets are always encapsulated to a RLOC, packets travel on the shortest path from LISP-MN to another LISP stationary or LISP-MN. When the LISP mobile node is sending packets to a stationary or LISP-MN in a non-LISP site, it sends LISP-encapsulated packets to a PETR which then decapsulates the packet and forwards it to its destination.

6. Updating Remote Caches

A LISP-MN has five mechanisms it can use to cause the mappings cached in remote ITRs and PITRs to be refreshed:

Map Versioning: If Map Versioning [LISP-VERSIONING] is used, an ETR can detect if an ITR is using the most recent database mapping. In particular, when mobile node's ETR decapsulates a packet and detects the Destination Map-Version Number is less than the current version for its mapping, in invokes the SMR procedure described in [LISP]. In general, SMRs are used to fix the out of sync mapping while Map-Versioning is used to detect they are out of sync. [LISP-VERSIONING] provides additional details of the Map Versioning process.

Data Driven SMRs: An ETR may elect to send SMRs to those sites it has been receiving encapsulated packets from. This will occur when an ITR is sending to an old RLOC (for which there is one-to-one mapping between EID-to-RLOC) and the ETR may not have had a chance to send an SMR the ITR.

Setting Small TTL on Map Replies: The ETR (or Map Server) may set a small Time to Live (TTL) on its mappings when responding to Map Requests. The TTL value should be chosen such that changes in mappings can be detected while minimizing control traffic. In this case the ITR is a SN and the ETR is the MN.

Piggybacking Mapping Data: If an ITR and ETR are co-located, an ITR may elect to send Map-Requests with piggybacked mapping data to those sites in its map cache or to which it has recently encapsulated data in order to inform the remote ITRs and PITRs of the change.

Temporary PITR Caching: The ETR can keep a cache of PITRs that have sent Map-Requests to it. The cache contains the RLOCs of the PITRs so later when the locator-set of a LISP-MN changes, SMR messages can be sent to all RLOCs in the PITR cache. This is an example of a control-plane driven SMR procedure.

7. Protocol Operation

There are five distinct connectivity cases considered by the LISP-MN design. The five mobility cases are:

LISP Mobile Node to a Stationary Node in a LISP Site.

LISP Mobile Node to a Non-LISP Site.

LISP Mobile Node to a LISP Mobile Node.

Non-LISP Site to a LISP Mobile Node.

LISP Site to a LISP Mobile Node.

The remainder of this section covers these cases in detail.

7.1. LISP Mobile Node to a Stationary Node in a LISP Site

After a roaming event, a LISP-MN must immediately register its new EID-to-RLOC mapping with its configured Map-Server(s). This allows LISP sites sending Map-Requests to the LISP-MN to receive the current mapping. In addition, remote ITRs and PITRs may have cached mappings that are no longer valid. These ITRs and PITRs must be informed that the mapping has changed. See Section 6 for a discussion of methods for updating remote caches.

7.1.1. Handling Unidirectional Traffic

A problem may arise when traffic is flowing unidirectionally between LISP sites. This can arise in communication flows between PITRs and LISP sites or when a site's ITRs and ETRs are not co-located. In these cases, data-plane techniques such as Map-Versioning and Data-Driven SMRs can't be used to update the remote caches.

For example, consider the unidirectional packet flow case depicted in Figure 1. In this case X is a non-LISP enabled SN (i.e., connected to the Internet) and Y is a LISP MN. Data traffic from X to Y will flow through a PITR. When Y changes its mapping (for example, during a mobility event), the PITR must update its mapping for Y. However, since data traffic from Y to X is unidirectional and does not flow though the PITR, it can not rely data traffic from Y to X to signal a mapping change at Y. In this case, the Y must use one or more of the techniques described in Section 6 to update the PITR's cache. Note that if Y has only one RLOC, then the PITR has to know when to send a Map-Request based on its existing state; thus it can only rely on the TTL on the existing mapping.

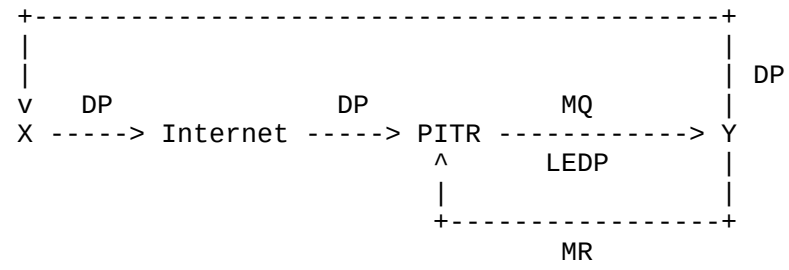

Figure 1: Unidirectional Packet Flow

7.2. LISP Mobile Node to a Non-LISP Stationary Node

LISP-MNs use the LISP Interworking infrastructure (specifically a PETR) to reach non-LISP sites. In general, the PETR will be co-located with the LISP-MN's Map-Server. This ensures that the LISP packets being decapsulated are from sources that have Map-Registered to the Map-Server. Note that when a LISP-MN roams it continues to uses its configured PETR and Map-Server which can have the effect of adding stretch to packets sent from a LISP-MN to a non-LISP destination.

7.3. LISP Mobile Node to LISP Mobile Node

LISP-MN to LISP-MN communication is an instance of LISP-to-LISP communication with three sub-cases:

- Both LISP-MNs are stationary (Section 7.1).
- Only one LISP-MN is roaming (Section 7.3.1).
- Both LISP-MNs are roaming. The case is analogous to the case described in Section 7.3.1.

7.3.1. One Mobile Node is Roaming

In this case, the roaming LISP-MN can find the stationary LISP-MN by sending Map-Request for its EID-prefix. After receiving a Map-Reply, the roaming LISP-MN can encapsulate data packets directly to the non-roaming LISP-MN node.

The roaming LISP-MN, on the other hand, must update its Map-Server

with the new mapping data as described in Section 7.1. It should also use the cache management techniques described in Section 6 to provide for timely updates of remote caches. Once the roaming LISP-MN has updated its Map-Server, the non-roaming LISP-MN can retrieve the new mapping data (if it hasn't already received an updated mapping via one of the mechanisms described in Section 6) and the stationary LISP-MN can encapsulate data directly to the roaming LISP-MN.

7.4. Non-LISP Site to a LISP Mobile Node

When a stationary ITR is talking to a non-LISP site, it may forward packets natively (unencapsulated) to the non-LISP site. This will occur when the ITR has received a negative Map Reply for a prefix covering the non-LISP site's address with the Natively-Forward action bit set [LISP]. As a result, packets may be natively forwarded to non-LISP sites by an ITR (the return path will through a PITR, however, since the packet flow will be non-LISP site to LISP site).

A LISP-MN behaves differently when talking to non-LISP sites. In particular, the LISP-MN always encapsulates packets to a PETR. The PETR then decapsulates the packet and forwards it natively to its destination. As in the stationary case, packets from the non-LISP site host return to the LISP-MN through a PITR. Since traffic forwarded through a PITR is unidirectional, a LISP-MN should use the cache management techniques described in Section 7.1.1.

7.5. LISP Site to LISP Mobile Node

When a LISP-MN roams onto a new network, it needs to update the caches in any ITRs that might have stale mappings. This is analogous to the case in that a stationary LISP site is renumbered; in that case ITRs that have cached the old mapping must be updated. This is done using the techniques described in Section 6.

When a LISP router in a stationary site is performing both ITR and ETR functions, a LISP-MN can update the stationary site's map-caches using techniques described in Section 6. However, when the LISP router in the stationary site is performing is only ITR functionality, these techniques can not be used because the ITR is not receiving data traffic from the LISP-MN. In this case, the LISP-MN should use the technique described in Section 7.1.1. In particular, a LISP-MN should set the TTL on the mappings in its Map-Replies to be in 1-2 minute range.

8. Multicast and Mobility

Since a LISP-MN performs both ITR and ETR functionality, it should also perform a lightweight version of multicast ITR/ETR functionality described in [MLISP]. When a LISP-MN originates a multicast packet, it will encapsulate the packet with a multicast header, where the source address in the outer header is one of it's RLOC addresses and the destination address in the outer header is the group address from the inner header. The interfaces in which the encapsulated packet is sent on is discussed below.

To not require PIM functionality in the LISP-MN as documented in [MLISP], the LISP-MN resorts to using encapsulated IGMP for joining groups and for determining which interfaces are used for packet origination. When a LISP-MN joins a group, it obtains the map-cache entry for the (S-EID,G) it is joining. It then builds a IGMP report encoding (S-EID,G) and then LISP encapsulates it with UDP port 4341. It selects an RLOC from the map-cache entry to send the encapsulated IGMP Report.

When other LISP-MNs are joining an (S-EID,G) entry where the S-EID is for a LISP-MN, the encapsulated IGMP Report will be received by the LISP-MN multicast source. The LISP-MN multicast source will remember the interfaces the encapsulated IGMP Report is received on and build an outgoing interface list for it's own (S-EID,G) entry. If the list is greater than one, then the LISP-MN is doing replication on the source-based tree for which it is the root.

When other LISP routers are joining (S-EID,G), they are instructed to send PIM encapsulated Join-Prune messages. However, to keep the LISP-MN as simple as possible, the LISP-MN will not be able to process encapsulated PIM Join-Prune messages. Because the map-cache entry will have a MN-bit indicating the entry is for a LISP-MN, the LISP router will send IGMP encapsulated IGMP Reports instead.

When the LISP-MN is sending a multicast packet, it can operate in two modes, multicast-origination-mode or unicast-origination-mode. When in multicast-origination-mode, the LISP-MN multicast-source can encapsulate a multicast packet in another multicast packet, as described above. When in unicast-origination-mode, the LISP-MN multicast source encapsulates the multicast packet into a unicast packet and sends a packet to each encapsulated IGMP Report sender.

These modes are provided depending on whether or not the mobile node's network it is currently connected can support IP multicast.

9. RLOC Considerations

This section documents cases where the expected operation of the LISP-MN design may require special treatment.

9.1. Mobile Node's RLOC is an EID

When a LISP-MN roams into a LISP site, the "RLOC" it is assigned may be an address taken from the site's EID-prefix. In this case, the LISP-MN will Map-Register a mapping from its statically assigned EID to the "RLOC" it received from the site. This scenario creates another level of indirection: the mapping from the LISP-MN's EID to a site assigned EID. The mapping from the LISP-MN's EID to the site assigned EID allow the LISP-MN to be reached by sending packets using the mapping for the EID; packets are delivered to site's EIDs use the same LISP infrastructure that all LISP hosts use to reach the site.

A packet egressing a LISP site destined for a LISP-MN that resides in a LISP site will have three headers: an inner header that is built by the host and is used by transport connections, a middle header that is built by the site's ITR and is used by the destination's ETR to find the current topological location of the LISP-MN, and an outer header (also built by the site's ITR) that is used to forward packets between the sites.

Consider a site A with EID-prefix 1.0.0.0/8 and RLOC A and a site B with EID-prefix 2.0.0.0/8 and RLOC B. Suppose that a host S in site A with EID 1.0.0.1 wants to talk to a LISP LISP-MN MN that has registered a mapping from EID 240.0.0.1 to "RLOC" 2.0.0.2 (where 2.0.0.2 allocated from site B's EID prefix, 2.0.0.0/8 in this case). This situation is depicted in Figure 2.

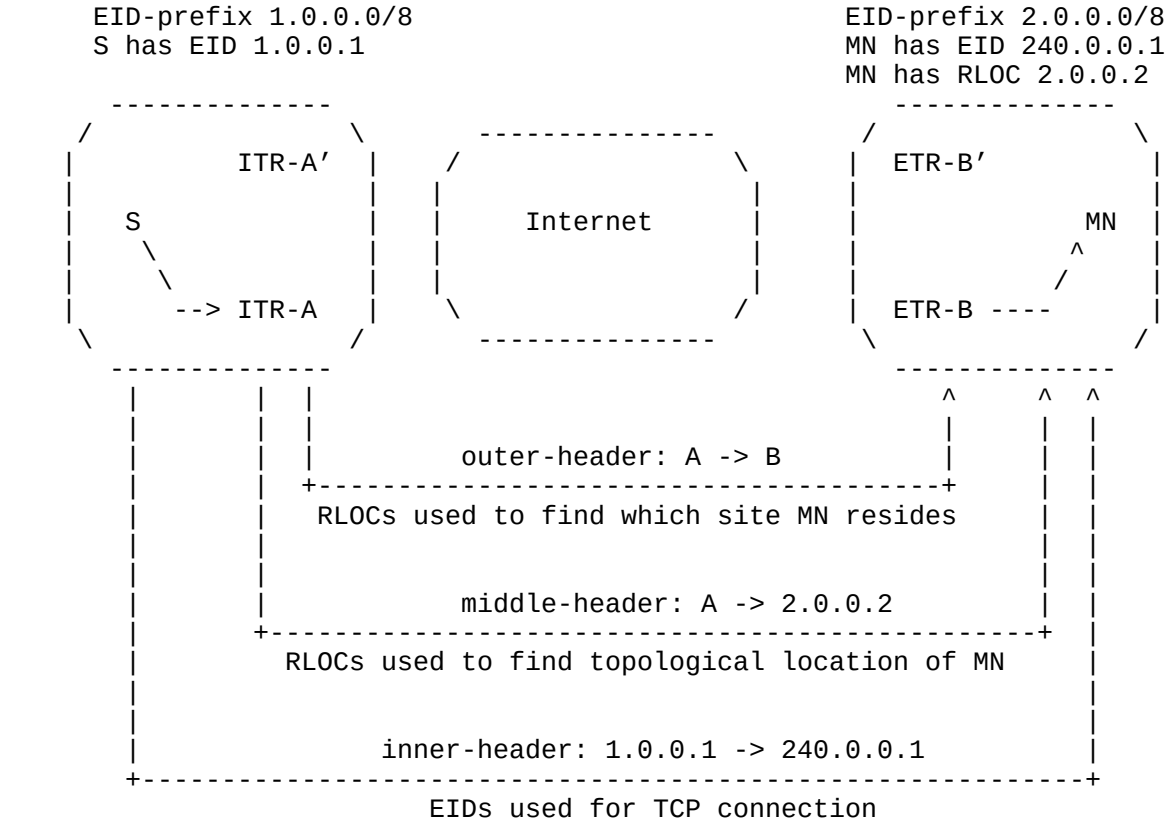

Figure 2: Mobile Node Roaming into a LISP Site

In this case, the inner header is used for transport connections, the middle header is used to find topological location of the LISP-MN (the LISP-MN Map-Registers the mapping 240.0.0.1 -> 2.0.0.2 when it roams into site B), and the outer header is used to move packets between sites (A and B in Figure 2).

In summary, when a LISP-MN roams into a LISP site and receives a new address (e.g., via DHCP) that is part of the site's EID space, the following sequence occurs:

1. The LISP-MN in the LISP site (call it Inside) registers its new RLOC (which is actually part of the sites EID prefix) to its map-server. Call its permanent EID E and the EID it DHCPs D. So it registers a mapping that looks like E->D.

2. The MN which is outside (call it Outside) sends a map request for inside's EID (E) and receives D (plus its policy). Outside realizes that D is an EID and sends a map request for D. This will return the site's RLOCs (by its ETR). Call this R.

3. Outside then double encapsulates the outbound packet with the inner destination being D and the outer destination being R.

4. The packet then finds its way to R, which strips the outer header and the packet is routed to D in the domain to Inside. Inside decapsulates the packet to serve the inner header to the application.

Note that both D and R could be returned to Inside in one query, so as not to incur the additional RTT.

10. Mobility Example

This section provides an example of how the LISP-MN is integrated into the base LISP Design [LISP].

10.1. Provisioning

The LISP-MN needs to be configured with the following information:

An EID, assigned to its loopback address

A key for map-registration

An IP address of a Map-Resolver (this could be learned dynamically)

An IP address of its Map-Server and Proxy ETR

10.2. Registration

After a LISP roams to a new network, it must immediately register its new mapping this new RLOC (and associated priority/weight data) with its Map-Server.

The LISP-MN may chose to set the 'proxy' bit in the map-register to indicate that it desires its Map-Server to answer map-requests on its behalf.

11. LISP Implementation in a Mobile Node

This section will describe a possible approach for developing a lightweight LISP-MN implementation. A LISP-MN will implement a LISP sub-layer inside of the IP layer of the protocol stack. The sub-layer resides between the IP layer and the link-layer.

For outgoing unicast packets, once the header that contains EIDs is built and right before an outgoing interface is chosen, a LISP header is prepended to the outgoing packet. The source address is set to the local RLOC address (obtained by DHCP perhaps) and the destination address is set to the RLOC associated with the destination EID from the IP layer. To obtain the RLOC for the EID, the LISP-MN maintains a map-cache for destination sites or destination LISP-MNs to which it is currently talking. The map-cache lookup is performed by doing a longest match lookup on the destination address the IP layer put in the first IP header. Once the new header is prepended, a route table lookup is performed to find the interface in which to send the packet or the default router interface is used to send the packet.

When the map-cache does not exist for a destination, the mobile node may queue or drop the packet while it sends a Map-Request to it’s configured Map-Resolver. Once a Map-Reply is returned, the map-cache entry stores the EID-to-RLOC state. If the RLOC state is empty in the Map-Reply, the Map-Reply is known as a Negative Map-Reply in which case the map-cache entry is created with a single RLOC, the RLOC of the configured Map-Server for the LISP-MN. The Map-Server that serves the LISP-MN also acts as a Proxy ETR (PETR) so packets can get delivered to hosts in non-LISP sites to which the LISP-MN is sending.

For incoming unicast packets, the LISP sub-layer simply decapsulates the packets and delivers to the IP layer. The loc-reach-bits can be processed by the LISP sub-layer. Specifically, the source EID from the packet is looked up in the map-cache and if the loc-reach-bits settings have changed, store the loc-reach-bits from the packet and note which RLOCs for the map-cache entry should not be used.

In terms of the LISP-MN detecting which RLOCs from each stored map-cache entry is reachable, it can use any of the Locator Reachability Algorithms from [LISP].

A background task that runs off a timer should be run so the LISP-MN can send periodic Map-Register messages to the Map-Server. The Map-Register message should also be triggered when the LISP-MN detects a change in IP address for a given interface. The LISP-MN should send Map-Registers to the same Map-Register out each of it’s operational links. This will provide for robustness on radio links with which

the mobile node is associated.

A LISP-MN receives a Map-Request when it has Map-Registered to a Map-Server with the Proxy-bit set to 0. This means that the LISP-MN wishes to send authoritative Map-Replies for Map-Requests that are targeted at the LISP-MN. If the Proxy-bit is set when the LISP-MN registers, then the Map-Server will send non-authoritative Map-Replies on behalf of the LISP-MN. In this case, the Map-Server never encapsulates Map-Requests to the LISP-MN. The LISP-MN can save resources by not receiving Map-Requests (note that the LISP-MN will receive SMRs which have the same format as Map-Requests).

To summarize, a LISP sub-layer should implement:

- Encapsulating and decapsulating data packets.
- Sending and receiving of Map-Request control messages.
- Receiving and optionally sending Map-Replies.
- Sending Map-Register messages periodically.

The key point about the LISP sub-layer is that no other components in the protocol stack need changing; just the insertion of this sub-layer between the IP layer and the interface layer-2 encapsulation/decapsulation layer.

12. Security Considerations

Security for the LISP-MN design builds upon the security fundamentals found in LISP [LISP] for data-plane security and the LISP Map Server [LISP-MS] registration security. Security issues unique to the LISP-MN design are considered below.

12.1. Proxy ETR Hijacking

The Proxy ETR (or PETR) that a LISP-MN uses as its destination for non-LISP traffic must use the security association used by the registration process outlined in Section 5.2 and explained in detail in the LISP-MS specification [LISP-MS]. These measures prevent third party injection of LISP encapsulated traffic into a Proxy ETR. Importantly, a PETR must not decapsulate packets from non-registered RLOCs.

12.2. LISP Mobile Node using an EID as its RLOC

For LISP packets to be sent to a LISP-MN which has an EID assigned to it as an RLOC as described in Section 9.1), the LISP site must allow for incoming and outgoing LISP data packets. Firewalls and stateless packet filtering mechanisms must be configured to allow UDP port 4341 and UDP port 4342 packets.

13. Acknowledgments

Albert Cabellos, Noel Chiappa, Pierre Francois, Michael Menth, Andrew Partan, Chris White and John Zwiebel provided insightful comments on the mobile node concept and on this document. A special thanks goes to Mary Nickum for her attention to detail and effort in editing early versions of this document.

14. IANA Considerations

This document creates no new requirements on IANA namespaces [RFC5226].

15. References

15.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

15.2. Informative References

[LISP] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-16.txt (work in progress), October 2011.

[LISP-ALT] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "LISP Alternative Topology (LISP-ALT)", draft-ietf-lisp-alt-09.txt (work in progress), September 2011.

[LISP-INTERWORK] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking LISP with IPv4 and IPv6", draft-ietf-lisp-interworking-02.txt (work in progress), June 2011.

[LISP-MS] Farinacci, D. and V. Fuller, "LISP Map Server", draft-ietf-lisp-ms-11.txt (work in progress), August 2011.

[LISP-VERSIONING] Iannone, L., Saucez, D., and O. Bonaventure, "LISP Mapping Versioning", draft-ietf-lisp-map-versioning-05.txt (work in progress), October 2011.

[MLISP] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "LISP for Multicast Environments", draft-ietf-lisp-multicast-11.txt (work in progress), October 2011.

Authors' Addresses

Dino Farinacci
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: dino@cisco.com

Darrel Lewis
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: darlewis@cisco.com

David Meyer
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: dmm@1-4-5.net

Chris White
Logical Elegance, LLC.
San Jose, CA 95134
USA

Email: chris@logicalelegance.com