

LISP  
Internet-Draft  
Intended status: Informational  
Expires: January 20, 2015

Y.H. Hertoghs  
M.B. Binderberger  
Cisco Systems  
July 21, 2014

**End Host Mobility Use Cases for LISP**  
**draft-hertoghs-lisp-mobility-use-cases-01**

Abstract

This memo proposes use cases for LISP in the area of end Host mobility. The applicability of end host mobility can be found in data centers, where Virtual Machines (VM's) can be moved freely from one physical server onto another physical server, independent of location, without having to change the IP/MAC-addresses inside those VMs, nor impacting traffic flows to and from those VMs. Wireless end hosts are another area of applicability. Although this draft will not address wireless end host mobility, most of the same principles apply.

Traditionally L2 extension technologies have been used to handle mobility events, but they could lead to suboptimal routing of traffic to and from the end host after the mobility event, as well as created big broadcast domains. This memo describes how LISP solves the traffic optimization issues caused by a mobility event of an end host like a Virtual Machine, as it decouples the identity of the end host from its location, such that traffic will always be forwarded to the correct location. More-over the LISP control plane can be leveraged to discover and distribute the reachability information of end hosts such that end to end broadcast domains, and their associated problems, are no longer needed.

Various sub-use cases will be looked at in this draft, depending on whether mobility is achieved at L2 (using MAC-addresses as EID) or at L3 (using IP addresses as EIDs), and whether subnets are L2 extended across LISP sites or not. This memo also describes how to handle mobility in the case where the default gateway of the end host is not capable of performing the LISP map-and-encap function, while the LISP xTR function is located one or more L3 hops away from the default gateway.

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 [RFC 2119](#) [[RFC2119](#)]

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).



Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <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 January 20, 2015.

## Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Terminology</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">LISP Use Cases for Mobility</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">End Host Mobility, extended subnet, IP as EID</a>	<a href="#">5</a>
<a href="#">3.2.</a>	<a href="#">End Host IP Mobility, non-extended subnet, IP as EID</a>	<a href="#">6</a>
<a href="#">3.3.</a>	<a href="#">End Host L2 Mobility, extended subnet/VLAN, MAC-address as EID</a>	<a href="#">8</a>
<a href="#">3.4.</a>	<a href="#">End Host Mobility, using a Combined L2/L3 LISP solution</a>	<a href="#">9</a>
<a href="#">3.5.</a>	<a href="#">End Host Mobility, using a Unified L2/L3 LISP solution</a>	<a href="#">10</a>
<a href="#">4.</a>	<a href="#">LISP Multi-hop Mobility</a>	<a href="#">10</a>
<a href="#">4.1.</a>	<a href="#">Overview</a>	<a href="#">11</a>
<a href="#">4.2.</a>	<a href="#">LISP Use Cases for Multihop Mobility</a>	<a href="#">12</a>
<a href="#">4.2.1.</a>	<a href="#">End Host IP Mobility, non-extended subnet</a>	<a href="#">12</a>
<a href="#">4.2.2.</a>	<a href="#">End Host IP Mobility, extended subnet</a>	<a href="#">13</a>
<a href="#">5.</a>	<a href="#">Protocol Extension Requirements</a>	<a href="#">14</a>
<a href="#">6.</a>	<a href="#">Acknowledgements</a>	<a href="#">14</a>
<a href="#">7.</a>	<a href="#">IANA Considerations</a>	<a href="#">14</a>
<a href="#">8.</a>	<a href="#">Security Considerations</a>	<a href="#">15</a>
<a href="#">9.</a>	<a href="#">References</a>	<a href="#">15</a>
<a href="#">9.1.</a>	<a href="#">Normative References</a>	<a href="#">15</a>

<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">15</a>
	Authors' Addresses . . . . .	<a href="#">16</a>

**[1.](#) Terminology**

LISP specific terminology such as Ingress-Tunnel-Router (iTR), Egress-Tunnel-Router (eTR), xTR, Proxy-iTR (PiTR), Proxy-eTR (PeTR), PxTR, Endstation IDentifier (EID), Routing Locator (RLOC), Mapping-Server (MS), Mapping-Resolver (MR), MS/MR can be found in [[RFC6830](#)] and [[RFC6832](#)].

## 2. Introduction

Moving a network node around while keeping it's IP address can be addressed in multiple ways. One way is to put the LISP xTR functionality on the mobile node, as described in [I-D.meyer-lisp-mn]. Another approach is to offload the mobility support to the site network. We divide the overall network into sites, with every site having one or more xTRs. Within the site mechanisms must be in place to detect a mobile host.

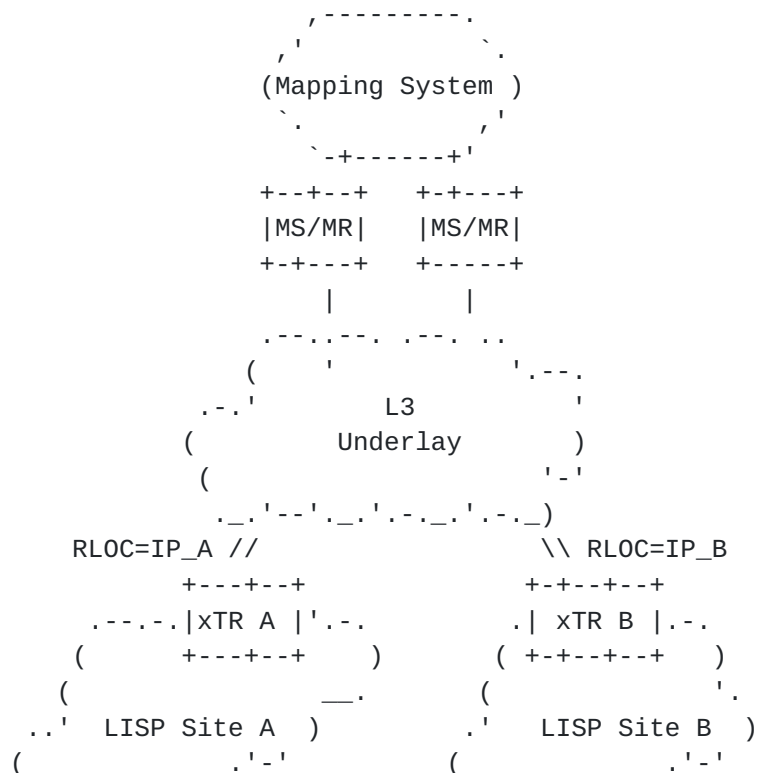
LISP [[RFC6830](#)] is a routing architecture that separates the location from the identity of hosts. A Host is part of a prefix known as an Endpoint Identity (EID), and an EID is attached to a LISP site. Traffic between EIDs at different locations is encapsulated by LISP xTR's, and the outer header's source and destination IP addresses are set to the source and destination Routing Locators (RLOCs) associated with the source and destination LISP Site. LISP eTR's register local EIDs and their associated RLOCs with the LISP Mapping System through LISP Map-Register Messages, while LISP iTRs request the destination RLOC for a given destination EID from the LISP Mapping System through LISP Map-Request Messages and associated LISP Map-Reply Messages.

As such it can be used to offer mobility support to end hosts e.g. to Virtual Machines (VMs) in data center networks, through the LISP xTR's registering the mobile end hosts IPv4 (/32), IPv6 (/128) and/or MAC-address (/48) as host-routes to the Mapping System, next to the existing least-specific LISP site EIDs where these hosts belong into. The LISP RLOC namespace functions as an underlay, while end host to end host communication is across the overlay in the EID namespace. Moves between sites are handled through a combination of LISP Map-Notify and LISP Solicit-Map-Request Messages. It should be noted that the EID movement is not limited to /32s or /128s. If e.g. a VM-cluster moves as a whole, and entire summary EID-prefix can move with it.



Figure 1 shows the reference architecture diagram we will use throughout this document. It shows two locations i.e. LISP sites A and B, with their specific xTR's XTR\_A and XTR\_B and Routing Locators (RLOCs) IP\_A and IP\_B. Note that within a typical Data Center architecture, these LISP sites can be as small as a set of interfaces where the end hosts are attached to on the (virtual) switch performing as a LISP xTR , or can be as large as a set of VLANs/EID subnets under one common Router performing as a LISP xTR. As such in these use cases, the name LISP site is a bit misleading as multiple 'LISP sites' can exist in one physical site, or even a rack in a data center, and is really dependent on the physical placement of the LISP xTR function. Typically a LISP xTR is placed as close as possible to the end hosts, and as such the scope of the LISP site is small. See [\[I-D.moreno-lisp-datacenter-deployment\]](#) for some deployment considerations.

Figure 1 also depicts a couple of end hosts X, Y and Z, each associated to a LISP Instance ID (IID) and a MAC-address based EID (/48) and/or an IP address based EID (/32 or /128). The IID used is dependent on the use cases, and can be used to identify a L2 instance or a L3 instance.







```

      '---'._.'. )\      '---'._.'. )\
      /      \      /      \
'-----'  '-----'  '-----'  '-----'
:  End  :  :  End  : ==> :  End  :  :  End  :
: Device :  : Device : ==> : Device :  : Device :
'-----'  '-----'  '-----'  '-----'
      EID=      EID=<IID1,MAC_W>      EID=
<IID2,MAC_X>      EID=<IID1,IP_W>      <IID2,MAC_Z>
<IID2,IP_X>      <IID2,IP_Z>

```

### 3. LISP Use Cases for Mobility

The following sections address the various ways in how LISP can be utilized to achieve end host mobility. The use cases differ in whether the end hosts subnet is extended towards the destination location or not, and whether IP-addresses (supporting IP flows) or MAC-addresses (supporting IP and non-IP traffic flows) are used to achieve mobility, while insuring no loss of sessions to and from the end host. The following use cases are addressed:

1. End host mobility, where the end host subnet is extended across the locations using some non-LISP L2 extension technique, where the EID is an IP address.
2. End host IP mobility, where the end host subnet is not extended across the locations, where the EID is an IP address.
3. End host L2 mobility, by using the MAC-Address as an EID. This effectively makes LISP a L2 extension technology, but without the disadvantage of traditional L2 extension techniques.
4. A Combined L2/L3 LISP based mobility solution, where Intra-subnet traffic is handled using the MAC-address as EID, while inter-subnet traffic is handled using the IP address as EID. This effectively combines use case 1 and 3.
5. A Unified L2/L3 LISP based mobility solution, where non-IP traffic is handled using the MAC-address as EID, while IP Intra- and Inter-subnet traffic are handled using IP addresses as EIDs. This is a combination of a generalized version of use case 2 and use case 3.

#### 3.1. End Host Mobility, extended subnet, IP as EID

This use case caters for both IP and non-IP traffic. The LISP protocol is only used to achieve optimized IP-based forwarding to hosts from outside the extended subnet i.e. LISP is only used for inter-subnet forwarding.



Two (or more) sites have the same subnet/prefixes configured, and the subnet is extended across them using a L2 extension. The control plane used for this L2 extension is responsible for non-IP traffic and intra-subnet forwarding (this could be e.g. flood-and-learn for 802.1d based bridging or VPLS), and the L2 extension will use MAC-address based forwarding. How the L2 extension works and how loop avoidance is achieved is out of scope of this memo.

For Inter-subnet based IP forwarding, xTR's at the site are configured with a set of prefixes containing EIDs of hosts which are mobile-eligible. This effectively causes discovered hosts to be registered as host-routes with the LISP Mapping System by the eTR function at the site. This has the advantage that traffic towards the mobile hosts will always be routed towards the correct site. This use case assumes that LISP xTRs are able to 'discover' hosts within the site which are part of the configured 'mobile-eligible' prefixes. This discovery can be triggered as a result of an ARP or IPv6 ND packet sourced by the Host, or by gleaning the source IP traffic of packets sourced from the Host, or through the use of host-to-network signalling e.g. VDP in IEEE802.1Qbp.

When a Host (e.g. Host X in instance 2 Figure 1) wants to send an IP packet to a moved Host (e.g. Host W in instance 1), which is in a different, but locally configured subnet, the local xTR (xTR\_A) will route those packets across the LISP overlay to the correct destination rather than locally route them. When the same Host sends a non-IP packet, or an IP packet destined within the same subnet (e.g. host Z in instance 2), it will be sent across the L2 extension.

Care needs to be taken that every site has a default gateway configured for the same prefix, and it uses the same (virtual) MAC-address in order to allow traffic from hosts to exit out of the local xTR rather than getting L2 switched back to another site. First Hop Redundancy Protocol (FHRP) originated packets (such as VRRP) have to be filtered between the two sites such that they cannot cross the L2 extension, and the FHRP protocol has to be configured identical in both sites, such that the virtual MAC-address of the default gateway is identical. In this way, the moved Host will always find the 'same' default gateway, irrespective of its location.

Hosts that are moving between sites will be discovered by the local xTR, and they will inform other xTR's which are connected to the extended subnet via LISP Map-Notify Messages. xTRs receiving traffic for a moved host will use LISP Solicit-Map-Request Messages to aid in clearing up the state of any eventual stale information at other xTRs.

For a discussion of the use case where the hosts are one or more L3 hops away from the site xTR, see [Section 4](#).

### **[3.2](#). End Host IP Mobility, non-extended subnet, IP as EID**

This use case caters for IP traffic only, for both intra-subnet and inter-subnet cases. The use case does not offer support for non-IP traffic flows. It is assumed that the LISP xTR in the site is the default gateway for those hosts that want mobility.

In this use case unique per-site IP EID prefixes are pre-configured in various LISP sites, and their location has been registered by the LISP eTR function at the site. A block of prefixes, part of the IP EID namespace associated with a site, can be configured across all sites as 'mobile-eligible'. If an xTR notices that one of these mobile-eligible prefixes match locally configured EID prefixes, the xTR will mark that site as 'home' of the prefix, when registering the prefix with the LISP Mapping System (standard LISP operation). The remaining prefixes are therefore 'away', when seen from that site perspective. All hosts discovered at an 'away' site which are part of one of those 'mobile-eligible' prefixes are registered with their /32 or /128 host route with the LISP Mapping System. In summary, prefixes are registered at home sites, as a result of provisioning, and host-route prefixes are registered at away sites, as a result of discovery.

When a host (e.g. Host W in Figure 1) moves from its LISP 'home site' A to LISP Site B, xTR\_B will notice the new Host, and will determine that it is part of a 'mobile-eligible' prefix. It will then register the new location of Host W with the LISP Mapping System, and inform xTR\_A of the fact that it has 'gone away' through LISP Map-Notify Messages. Sources sending traffic to the moved Host W might still send traffic to the old location site A. When receiving LISP encapsulated traffic, xTR\_A will notify the origin LISP xTR that it needs to update its mapping cache (this can be a LISP xTR, or a LISP PxTR in case the source is not part of a LISP site) via LISP Solicit-Map-Request Messages. That specific xTR will then consult the Mapping System to achieve the correct location of the end host, and update its local cache.

This use case assumes that LISP xTRs are able to 'discover' hosts within the site which are part of the configured 'mobile-eligible' prefixes. This discovery can be triggered as a result of an ARP or IPv6 ND packet sourced by the Host, or by gleaning the source IP traffic of packets sourced from the Host, or through the use of host-to-network signalling e.g. VDP in IEEE802.1Qbp.

The LISP xTR can be quite centrally located within the site i.e. a couple of L2 hops (or even L3 hops, see [Section 4](#)) away. In the case where the LISP xTR is a couple of L2 hops away, care needs to be taken making sure that ARP and IPv6 NDs tables of other hosts part of the same subnet as the mobile Host are synchronized after the move. There are three broad ways of achieving that:



1. When a LISP xTR gets notified of a Host that has 'moved away', it will issue an IPv4 GARP or an IPv6 Unsolicited Neighbour Announcement (NA) towards all hosts which are local. The GARP or NA will contain the MAC-address of the LISP xTR rather than the host's MAC-address. In this way it will attract all traffic that is sent to 'moved-away' hosts.
2. All traffic between hosts is always forced to be hairpinned through the local LISP xTR (i.e. all traffic from hosts underneath the xTR will flow 'through' the xTR, even intra-subnet traffic) and the LISP xTR can therefore catch and respond to all ARP and ND requests from hosts with a well-known per-subnet MAC-address that is shared between all xTRs that take care of 'mobile-eligible' prefixes. The hairpinning can be achieved by installing forwarding rules in the L2 switches underneath the LISP xTRs, achieving the hairpinning result, or by placing the LISP xTR function L2 adjacent to the mobile hosts (e.g on the Top Of Rack/End of Rack/Virtual Switch). How this is accomplished is out of scope of this draft. Every host's ARP/ND table will therefore have entries pointing to the same MAC-address i.e the local LISP xTR.
3. LISP xTRs register all active hosts with both their IP EID as well as their MAC EID. All traffic between hosts is always forced to be hairpinned through the local LISP xTR (see above), and the LISP xTR will do a LISP database lookup, either to respond on behalf of the Host with the real MAC-address of the Host, or by relaying the ARP/ND end to end. The LISP xTRs will also route all IP packets independent of their destination MAC-address, regardless of whether the destination is local or remote.

For a discussion of the use case where the hosts are one or more L3 hops away from the site xTR, see [Section 4](#).

### **3.3. End Host L2 Mobility, extended subnet/VLAN, MAC-address as EID**

This use case caters for both IP and non-IP traffic, by treating the IP traffic as L2 traffic. End hosts MAC-address /48 host routes are registered with the Mapping System rather than IP prefixes and IP host routes, effectively creating a LISP L2 extension solution.

[I-D.smith-lisp-layer2] documents how LISP can be used to register and forward based on MAC-addresses, while the underlay is IP based. LISP xTRs performing the L2 forwarding can be placed at any place in the hierarchical network topology of the site, and there is no need to hairpin all traffic upstream through the site's LISP xTR. However, the L2 nodes on a specific site have to clear their respective bridge table entry for all hosts that moved away. How this is done as a

result of a host moving is out of scope for this draft, although the most easy way is to colocate the LISP xTR function with the switch L2 adjacent to the 'mobile-eligible' hosts.



Loop avoidance in case of two LISP sites L2 interconnecting by some means unknown to LISP is needed, but is out of scope of this draft.

Although traditional bridging ('flood-and-learn') is used within the site, inter-site flooding is prevented, assuming that all active mobile hosts are registered with the Mapping System, and as such no flooding to unknown destinations needs to be performed as all hosts are known. Optionally the IP addresses can also be registered with the Mapping System, and this can aid in not having to broadcast ARP and ND packets across LISP sites, as the local LISP xTR can respond to the ARP/ND on behalf of the end-station. In case the ARP/ND packets do need to be relayed to the correct destination, registering the IP next to the MAC-address makes this easy to achieve and effectively turns the ARP/ND MAC level broadcast/multicast into an IP unicast in the underlay. MAC-level multicast and broadcasts can be encapsulated into multicasts in the RLOC namespace, or can be ingress replicated to the correct destination sites, using the techniques in [\[RFC6831\]](#) and [\[I-D.farinacci-lisp-mr-signaling\]](#). This significantly reduces the need for multicast in the underlay.

For a Network Virtualization Overlay (NVO3) specific based implementation and for a description of LISP Messages used when hosts move between sites, see [\[I-D.maino-nvo3-lisp-cp\]](#).

#### **3.4. End Host Mobility, using a Combined L2/L3 LISP solution**

This use case caters for both IP and non-IP traffic. This use case is effectively a combination of use case 1 ([Section 3.1](#)) and use case 3 ([Section 3.3](#)), where LISP provides the L2 extension functionality required.

Hosts IP addresses and MAC-addresses are independently registered with the LISP Mapping System upon discovery. The placement of the xTR functions for both namespaces needs to be co-located on the same device. This functionality is often referred to a 'Integrated Routing and Bridging'. The combined L2/L3 LISP xTR can be placed at any place in the hierarchical network topology of the site, and there is no need to hairpin all traffic upstream through the site's LISP xTR, although making the LISP xTR function colocate within the switch L2 adjacent to the mobile hosts can have its advantages, see [Section 3.3](#). All LISP xTR's which offer mobility support for the same prefix, need to be configured with the same virtual MAC-address, such that hosts will think they talk to the same default gateway independent of which site they are located at.

Intra-subnet traffic (that includes both IP and non-IP traffic) is handled within the MAC-address EID namespace as per [Section 3.3](#), where as Inter-subnet traffic is handled within the IP-address EID



namespace as per [Section 3.1](#).

### **3.5. End Host Mobility, using a Unified L2/L3 LISP solution**

This use case caters for both IP and non-IP traffic. It differs with the use case in [Section 3.4](#) in that it handles all IP traffic i.e. Intra-subnet and Inter-subnet within the IP EID namespace, while it handles all non-IP traffic within the MAC-address EID namespace. In other words it is a combination of the use case in [Section 3.2](#), and the use case in [Section 3.3](#) for non-IP traffic. Again the L2 and L3 xTR functions need to be colocated on the same physical node.

Since the Unified L2/L3 is also responsible for intra-subnet IP forwarding, all traffic from within the subnet/VLAN needs to be hairpinned across the Unified L2/L3 LISP xTR. The simplest way to achieve this is to make it L2 adjacent to the mobile hosts. Other methods of achieving this hairpinning are out of scope of this draft. As a result the notion of a subnet equaling a broadcast domain goes away. The subnet is purely used as a common address pool across participating LISP sites. The Unified L2/L3 LISP xTR will route all IP packets, independent of their destination MAC-address and whether these packets are part of Intrasubnet or Inter-subnet flows.

An important distinction of this use case is the fact that, for IP traffic, also the MAC-address will be registered, next to the IP address with the Mapping System. This allows the LISP xTR to optimise ARP and IPv6 ND handling for intra-subnet traffic. For non-IP traffic, the MAC-address of the host is also registered as a separate entry with the LISP Mapping System.

The Unified L2/L3 LISP xTRs are configured with a uniform default gateway MAC-address and IP address across all LISP sites. Mobile hosts will thus think they talk to the same default gateway independent of which site they are located at.

For a Network Virtualization Overlay (NVO3) specific based implementation of the Unified L2/L3 LISP solution and for a description of LISP Messages used when hosts move between sites, see [[I-D.hertoghs-nvo3-lisp-controlplane-unified](#)].

## **4. LISP Multi-hop Mobility**

There are cases when the detection of a mobile host needs to be separated from the LISP encapsulation/decapsulation. The detection typically happens on the layer 3 hop that is connected to the mobile host, i.e. it happens close to the mobile host. The LISP encapsulation/decapsulation may be performed by other equipment further "up" in the diagrams Figure 1 and Figure 2, i.e. closer to the entry/exit of the data center. We name the detection node a First

Hop (FH) node from here on. The FH does not need to support LISP on the data plane but must implement certain LISP control plane

In theory one could combine detection and xTR functionality on the first hop nodes FH-1 and FH-2; these are the scenarios discussed in the earlier sections. Sometimes though a requirement exists for firewalls, IDS, load-balancers, WAN optimizers and other functionality in the "per site network" boxes in the figure above.

Some of these services require access to the EID-space packet and may not have the ability to look into the LISP data packet. This is why detection and actual LISP encapsulation are separated for the following discussion.

## **4.2. LISP Use Cases for Multihop Mobility**

The use cases are similar to the scenarios discussed in [Section 3](#). Our focus will be on L3 though as the services implemented in the "per site network" are typically IP based:

- o End host IP mobility , where the end host subnet is not extended across the locations, where the EID is an IP address.
- o End host IP mobility, where the end host subnet is extended across the locations using some non-LISP L2 extension technique, where the EID is an IP address.

The difference between multihop and singlehop host mobility (as in [Section 3](#)) is mainly in the additional routing setup required in each site to match the host detection and LISP signaling. This will be discussed in the following sections.

### **4.2.1. End Host IP Mobility, non-extended subnet**

The discussion of [Section 3.2](#) applies for this case as well. The aspects covering ARP are handled by the First Hop routers while LISP registration and associated signaling is handled by the xTR. As a minimum the site xTR must register moved-in hosts to the mapping system, based on the detection on the FH router. This will attract traffic for the hosts that moved-in from the LISP network.

What needs to be added is the routing inside the sites. Assume host H1 has moved away from site-1 into site-2. For site-1 in this example two fundamental options exist: xTR-1 will be informed by the LISP Mapping System that host H1 has been detected elsewhere. This information can be used to inject a host route for H1 from xTR-1 into site-1 IGP. The FH-1 would inject the network prefix P1 into site-1's IGP. A detection of active hosts with an address within P1 would not be necessary as long as a detection of the return of absent hosts is guaranteed. xTR-1 would register the network prefix P1. For foreign hosts from network prefix P2 the FH must detect them and inject a host route into site-1 IGP. The xTR of site-1 would register these foreign hosts. In other words the IGP would carry the network prefix P1 and hosts routes for local, foreign hosts and for the absent hosts belonging to network P1. The routing for undetected host from prefix P1 would point to the FH router.





As an alternative all FH's could detect the active hosts on their LAN, both for hosts that are home at the site as well as hosts that moved into the site's network. The FH routers would inject a host route into the site IGP for every detected host. The xTR-1 injects more specific subnets of network P1 into the IGP to attract the traffic for P1. As an example an xTR would inject 10.1.0.0/17 and 10.1.128.0/17 to attract traffic for the prefix 10.1.0.0/16. The xTR-n would still register network prefix P1 and the foreign hosts with the mapping system. The IGP would carry the subnets of prefix P and hosts routes for all local hosts. The default routing for prefix P would point to the xTR. Practically a mix of these two options is possible to optimize the number of routes, the number of registrations and/or the number of mapping requests.

Both sites can choose their internal routing scheme independently. This is possible as the registration scheme is the same: register the home network and the foreign hosts. In all cases a default route pointing to the xTR is completing the routing, to reach all other EID prefixes.

#### **4.2.2. End Host IP Mobility, extended subnet**

The discussion of [Section 3.1](#) applies here as well, covering aspects of ARP handled by the First Hop routers and LISP registration handled by the xTR.

For extended subnets it doesn't make sense to talk about the home site or a host returning home. All registrations with the LISP mapping system must be on a per-host basis, based on the locally detected hosts. Additional registration of the network prefix P would be necessary when the setup requires the xTRs of every site to know about all remote hosts. Such registrations of P would have the downside to attract traffic from the LISP network that finally may be dropped when no receiving host is found.

For the routing we can identify again two fundamental cases. The first case is that the FH, when detecting a mobile host, is not only informing the site's xTR but it also redistributing a host route into the site IGP. The xTR would redistribute subnets of network prefix P into the site IGP to attract all addresses of P that are not detected in the site. A standard default route pointing to the xTR would complete the site routing. In summary the site IGP would carry host routes for all locally detected hosts and the default routing for prefix P would be towards the xTR.

The other case is that the xTR knows about all the remotely registered hosts and redistributes them as host routes into the site IGP. The FH router would inject a route for network P into the site

IGP. A default route pointing to the xTR would complete the routing for the non-mobile EID prefixes. The IGP would carry the host routes of the remotely detected hosts and default routing for P would point to the FH.

Again sites can chose their internal routing independently as the common way to register all locally detected hosts guarantees interoperability of the site routings.

The extended network scenario allows to handle the case of a "silent host", i.e. a host that behaves passive until it receives a request. This means a silent host can be detected only when traffic is sent to it. To support silent host detection at least one site covering the extended subnet must advertise the subnet's prefix P to the LISP mapping system and the site's routing must forward unknown host within prefix P to the FH. Advertising prefix P from the FH into the IGP and the xTR redistributing remotely registered hosts into the IGP fulfills this requirement.

## **5. Protocol Extension Requirements**

This section is not an exhaustive discussion nor description of the LISP protocol extensions used for the use cases. It only briefly mentions the requirements that allow for the design of the use cases in the previous sections.

- o Multiple registrations for the same prefix and parent registrations: registering a host and overriding the previous registration and then informing every site that needs to be updated is the core of how the mapping system synchronizes the sites. Imagine the extended subnet case with 2 sites and a host is activated at site-1. When site-2 uses a routing that requires knowledge of all remote hosts then site-2 needs to be updated. The proposed protocol extension requires site-2 to register the network P and the map server will notify every registerer of P when a host registration falls into the registered network P.
- o Reliable Notifications: notifications are a key aspect of how LISP mapping services updates the sites. A mechanism is proposed to acknowledge received notifications and retry sending them when the ACK is missing.
- o Separation of messages between xTR and FH against map server: an expected setup is to run both map resolver and map server on xTR routers. Messages from FH to xTR must be distinguishable from map registrations to avoid confusing the map server.

## **6. Acknowledgements**

The authors want to thank Dino Farinacci, Patrice Bellagamba, Johnson Leong, Victor Moreno and Fabio Maino for the early review, insightful comments and suggestions. The multihop discussion is using work from Chris Cassar and Isidor Kouvelas.

## **7. IANA Considerations**

This memo includes no request to IANA.

## **8. Security Considerations**

See [[I-D.ietf-lisp-sec](#)] for a list of security considerations for LISP.

## **9. References**

### **9.1. Normative References**

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### **9.2. Informative References**

[I-D.farinacci-lisp-mr-signaling]  
Farinacci, D. and M. Napierala, "LISP Control-Plane Multicast Signaling", Internet-Draft [draft-farinacci-lisp-mr-signaling-04](#), March 2014.

[I-D.hertoghs-nvo3-lisp-controlplane-unified]  
Hertoghs, Y., Maino, F., Moreno, V., Smith, M., Farinacci, D. and L. Iannone, "A Unified LISP Mapping Database for L2 and L3 Network Virtualization Overlays", Internet-Draft [draft-hertoghs-nvo3-lisp-controlplane-unified-01](#), February 2014.

[I-D.ietf-lisp-sec]  
Maino, F., Ermagan, V., Cabellos-Aparicio, A. and D. Saucez, "LISP-Security (LISP-SEC)", Internet-Draft [draft-ietf-lisp-sec-06](#), April 2014.

[I-D.maino-nvo3-lisp-cp]  
Maino, F., Ermagan, V., Hertoghs, Y., Farinacci, D. and M. Smith, "LISP Control Plane for Network Virtualization Overlays", Internet-Draft [draft-maino-nvo3-lisp-cp-03](#), October 2013.

[I-D.meyer-lisp-mn]  
Farinacci, D., Lewis, D., Meyer, D. and C. White, "LISP Mobile Node", Internet-Draft [draft-meyer-lisp-mn-10](#), January 2014.

[I-D.moreno-lisp-datacenter-deployment]  
Moreno, V., Maino, F., Lewis, D., Smith, M. and S. Sinha, "LISP Deployment Considerations in Data Center Networks", Internet-Draft [draft-moreno-lisp-datacenter-deployment-00](#), February 2014.

[I-D.smith-lisp-layer2]

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

- [RFC6830] Farinacci, D., Fuller, V., Meyer, D. and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), January 2013.
- [RFC6831] Farinacci, D., Meyer, D., Zwiebel, J. and S. Venaas, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", [RFC 6831](#), January 2013.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D. and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", [RFC 6832](#), January 2013.

#### Authors' Addresses

Yves Hertoghs  
Cisco Systems  
De Kleetlaan 6a  
Diegem, 1831  
BE

Email: [yhertogh@cisco.com](mailto:yhertogh@cisco.com)

Marc Binderberger  
Cisco Systems  
510 McCarthy Blvd.  
Milpitas, California 95035  
USA

Email: [mbinderb@cisco.com](mailto:mbinderb@cisco.com)

