


LISP Host Mobility Solution

The traditional IP addressing model associates both Location and Identity to a single IP address space, making mobility a very cumbersome process since identity and location are tightly bundled together. Because LISP creates two separate name spaces, separating IP addresses into Route Locators (RLOC) and End-point Identifiers (EID), and provides a dynamic mapping mechanism between these two address families, EIDs can be found at different RLOCs based on the EID-RLOC mappings. RLOCs remain associated with the topology and are reachable by traditional routing. However, EIDs can change locations dynamically and are reachable via different RLOCs, depending on where an EID attaches to the network. In a virtualized data center deployment, EIDs can be directly assigned to Virtual Machines that are hence free to migrate between data center sites preserving their IP addressing information.

The decoupling of Identity from the topology is the core principle on which the LISP Host Mobility solution is based. It allows the End-point Identifier space to be mobile without impacting the routing that interconnects the Locator IP space. When a move is detected, as illustrated in [Figure 3-1](#), the mappings between EIDs and RLOCs are updated by the new xTR. By updating the RLOC-to-EID mappings, traffic is redirected to the new locations without requiring the injection of host-routes or causing any churn in the underlying routing.

Figure 3-1 LISP Host Mobility



LISP Host Mobility detects moves by configuring xTRs to compare the source in the IP header of traffic received from a host against a range of prefixes that are allowed to roam. These prefixes are defined as Dynamic-EIDs in the LISP Host Mobility solution. When deployed at the first hop router (xTR), LISP Host Mobility devices also provide adaptable and comprehensive first hop router functionality to service the IP gateway needs of the roaming devices that relocate.

LISP Host Mobility Use Cases

The level of flexibility provided by the LISP Host Mobility functionality is key in supporting various mobility deployment scenarios/use cases ([Table 3-1](#)), where IP end-points must preserve their IP addresses to minimize bring-up time:


Table 3-1 *Use Case Requirements*

Use Case	Requirements
Geo-Clusters	Optimized routing to IP subnets extended with OTV or VPLS
Fast Bring-up of Disaster Recovery Facilities	Relocation of IP end-points across different subnets
Cloud Bursting	Relocation of IP end-points across organizations
Data Center Consolidation/Migration	Migration of workloads/applications between data center sites (often from brownfield to greenfield locations)

One main characteristic that differentiates the use cases listed in the table above is the fact that the workloads are moved in a “live” or “not live” (cold) fashion. Limiting the discussion to virtualized computing environments, the use of VMware vMotion is a typical example of a live workload migration, whereas VMware Site Recovery Manager (SRM) is an example of application used for cold/not live migrations.

LISP Host Mobility offers two different deployment models, which are usually associated to the different type of workload mobility scenarios:


- **LISP Host Mobility with an Extended Subnet:** in this model, represented in [Figure 3-2](#), the IP Subnets/VLANs are extended across the West and East data center sites leveraging OTV, VPLS, or any other LAN extension technology.

Figure 3-2 LISP Host Mobility with LAN Extension

In traditional routing, this poses the challenge of steering the traffic originated from remote clients to the data center site where the workload is located, given the fact that a specific IP subnet/VLAN is no longer associated to a single DC location. LISP Host mobility is hence used to provide seamless ingress path optimization by detecting the mobile EIDs dynamically, and updating the LISP Mapping system with its current EID-RLOC mapping.

This model is usually deployed when geo-clustering or live workload mobility is required between data center sites, so that the IP mobility functionality is provided by the LAN extension technology, whereas LISP takes care of inbound traffic path optimization.

- **LISP Host Mobility Across Subnets:** in this model, shown in [Figure 3-3](#), two different subnets exist between one of the two DC sites that are L2 connected (West or East) and a remote DC location (as for example a Disaster Recovery site or a Cloud Provider Data Center). In this scenario, LAN extension techniques are not deployed toward this remote location.

Figure 3-3 LISP Host Mobility Across Subnets

This model allows a workload to be migrated to a remote IP subnet while retaining its original IP address and can be generally used in cold migration scenarios (like Fast Bring-up of Disaster Recovery Facilities in a timely manner, Cloud Bursting or data center migration/consolidation). In these use cases, LISP provides both IP mobility and inbound traffic path optimization functionalities.

This document discusses both of these modes in detail, and provides configuration steps and deployment recommendations for each of the LISP network elements.

LISP Host Mobility Hardware and Software Prerequisites

Some of the hardware and software prerequisites for deploying a LISP Host Mobility solution are listed below.

- At time of writing of this document, LISP Host Mobility is only supported on Nexus 7000 platforms and ISR and ISRG2 platforms. This functionality will also be supported in the ASR 1000 router in 2HCY12 timeframe. This document only covers the configuration and support of LISP Host Mobility on Nexus 7000 platforms.
- For the Nexus 7000 platform, LISP Host Mobility is supported in general Cisco NX-OS Release 5.2(1) or later. LISP Host Mobility support requires the Transport Services license, which includes both LISP and OTV.
- From a switch modules perspective, LISP is currently supported only in N7K-M132XP-12 and in N7K-M132XP-12L linecards.



Caution

EPLD upgrade to the minimum version 186.008 is required for these linecards. For more information on how to upgrade the EPLDs refer to link:

http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_0/epld/release/notes/epld_rn.html


LISP is also supported on Nexus 7000 leveraging F1-series modules for site-facing interfaces, as long as one of the M1-32 series cards mentioned above is available in the chassis. This is because only M1-32 linecards can perform LISP encapsulation/decapsulation in HW, so it is required for L2 traffic received on F1 interfaces to be internally sent to these M1 cards for that purpose. This functionality available on Nexus 7000 platforms is known as "proxy routing".



Note

Only the F1 series line cards support proxy mode at the time of writing this document. Proxy mode is not supported between M-series cards, so the site facing interfaces can only be N7K-M132XP-12, N7K-M132XP-12L or F1-series linecards.

Figure 3-4 shows the specific HW deployment models for LISP Host Mobility.

Figure 3-4 Nexus 7000 LISP HW Support

Notice how traffic received on M-series linecards other than the M1-32 cards will not be processed by LISP. Therefore combining M1-32 cards with other M-series cards in a LISP enabled VDC will result in two different types of traffic processing depending on which interface receives the traffic. In deployments where other M-series cards (N7K-M148GT-11, N7K-M148GT-11L, N7K-M148GS-11, N7K-M148GS-11L or N7K-M108X2-12L) are part of the same VDC with the F1 and M1-32 cards, it is critical to ensure that traffic received on any F1 interfaces is internally sent (proxied) only to interfaces belonging to M1-32 cards. The "hardware proxy layer-3 forwarding use interface" command can be leveraged to list only these specific interfaces to be used for proxy-routing purpose.

**Note**

M2 linecards (not yet available at the time of writing of this paper) won't have support for LISP in HW, and the same is true for currently shipping F2 linecards.

LISP Host Mobility Operation

As previously mentioned, the LISP Host Mobility functionality allows any IP addressable device (host) to move (or "roam") from its subnet to a completely different subnet, or to an extension of its subnet in a different location (e.g. in a remote Data Center) - while keeping its original IP address. In LISP terminology, a device that moves is called a "roaming device," and its IP address is called its "dynamic-EID" (Endpoint ID).


The LISP xTR configured for LISP Host Mobility is usually positioned at the aggregation layer of the DC, so to allow the xTR function to be co-located on the device functioning also as default gateway. As clarified in the following sections of this paper, the LISP xTR devices dynamically determine when a workload moves into or away from one of its directly connected subnets. At time of writing of this document it is recommended to position the DC xTR as close to the hosts as possible, to facilitate this discovery of the mobile workload. This usually means enabling LISP on the first L3 hop device (default gateway).

Map-Server and Map-Resolver Deployment Considerations

As previously discussed, the Map-Server and Map-Resolver are key components in a LISP deployment. They provide capabilities to store and resolve the EID-to-RLOC mapping information for the LISP routers to route between LISP sites. The Map-Server is a network infrastructure component that learns EID-to-RLOC mapping entries from ETRs that are authoritative sources and that publish (register) their EID-to-RLOC mapping relationships with the Map-Server. A Map-Resolver is a network infrastructure component which accepts LISP encapsulated Map-Requests, typically from an ITR, and finds the appropriate EID-to-RLOC mapping by checking with a co-located Map-Server (typically), or by consulting the distributed mapping system.

This section details the Map-Server and Map-Resolver deployment consideration for an enterprise Data Center LISP deployment. It is recommended to deploy redundant standalone Map-Server and Map-Resolver systems, with the MS/MR functions co-located on the same device. When redundant standalone Map-Server/Map-Resolver are deployed, all xTRs must register to both Map-Servers so that each has a consistent view of the registered LISP EID namespace. For Map-Resolver functionality, using an Anycast IP address is desirable, as it will improve the mapping lookup performance by choosing the Map-Resolver that is topologically closest to the requesting ITR. The following topology in [Figure 3-5](#) details this deployment model.

Figure 3-5 Redundant MS/MR Deployment




Note

For large-scale LISP deployments, the Mapping database can be distributed and MS/MR functionality dispersed onto different nodes. The distribution of the mapping database can be achieved in many different ways, LISP-ALT and LISP-DDT being two examples. Large-scale LISP deployments using distributed mapping databases are not discussed in this document; please reference to lisp.cisco.com for more information on this matter.

It is worth noticing that a redundant standalone MS/MR model can be deployed by leveraging dedicated systems to perform these mapping functionalities (shown in [Figure 3-5](#)), or alternatively by deploying the Map-Server and Map-Resolver functionalities concurrently on the same network device already performing the xTR role ([Figure 3-6](#)).

Figure 3-6 Co-locating MS/MR and xTR Functionalities




The co-located model shown above is particularly advantageous as it reduces the overall number of managed devices required to roll out a LISP Host Mobility solution. Notice that the required configuration in both scenarios would remain identical, leveraging unique IP addresses to identify the Map-Servers and an Anycast IP address for the Map-Resolver.

Dynamic-EID Detection

When a dynamic-EID roams between data center sites, the local LISP Host Mobility xTRs need to detect its existence. As it will be discussed in detail in the next chapter, this information is then used to update the content of the mapping database, allowing the steering of the traffic to the new data center location.

[Figure 3-7](#) shows the dynamic detection of an EID migrated between DC sites.

Figure 3-7 Dynamic Discovery of a Migrated EID

The LISP xTR configured for LISP Host Mobility detects a new dynamic-EID move event if

1. It receives an IP data packet from a source (the newly arrived workload) that is not reachable from a routing perspective via the interface on which the packet was received.

**Note**

This mechanism is leveraged not only when moving a workload between different subnets (shown in Figure 3-7), but also in scenarios where an IP subnet is stretched between DC sites. A detailed discussion can be found in [Deploying LISP Host Mobility with an Extended Subnet](#), page 4-1 on how to trigger a check failure under these circumstances.

2. The source matches the dynamic-EID configuration applied to the interface.

At the time of writing of this paper, when leveraging NX-OS SW releases 5.2(5) and above, two kinds of packets can trigger a dynamic EID discovery event on the xTR:

1. IP packets sourced from the EID IP address. A discovery triggered by this kind of traffic is referred to as "data plane discovery".
2. ARP packets (which are L2 packets) containing the EID IP address in the payload. A discovery triggered by this kind of traffic is referred to as "control plane discovery". A typical scenario where this can be leveraged is during the boot-up procedure of a server (physical or virtual), since it would usually send out a Gratuitous ARP (GARP) frame to allow duplicate IP address detection, including in the payload its own IP address.

**Note**

The control plane discovery is only enabled (automatically) when deploying LISP Host Mobility Across Subnets. This is because in Extended Subnet Mode, the ARP messages would be sent also across the logical L2 extension between sites, potentially causing a false positive discovery event.

Once the LISP xTR has detected an EID, it is then responsible to register that information to the Mapping Database. This registration happens immediately upon detection, and then periodically every 60 seconds for each detected EID. More details on how this is happening will be discussed in the following two chapters of this paper.