


Deploying LISP Host Mobility Across Subnets

Figure 5-1 shows the Enterprise datacenter deployment topology with a West DC having the subnet 10.17.1.0/24 and an East DC having the subnet 10.18.1.0/24. The mobile subnet is 10.17.0.0/24 associated to VLAN 1301, which is locally defined in the West Data Center.

Figure 5-1 LISP Host Mobility across Subnets Topology



In the rest of this chapter, the specific site where the mobile subnet is natively defined may be called the “home site”, whereas the subnet itself can be referred to as the “home subnet”. The IP address space available in the East Data Center is instead completely independent from the one defined in the West DC. In the example above, the workloads that are moved away from the West DC site will land on a different 10.18.1.0/24 subnet locally defined in the East DC. The remote site is deployed with a Cisco IOS device hosting the 10.10.10.0/24 network. This section describes steps to configure these data center sites and remote IOS device sites as LISP sites with their corresponding EID spaces. This section also describes the required configuration to enable EIDs to move between data centers and highlights how client-server and inter-DC traffic flows can be established.


LISP Host Mobility Across Subnets: Sample Config

With reference to [Figure 5-1](#), the following are the basic configuration steps required on the various devices in the network to enable LISP Host Mobility.

Nexus 7000 N7K1A and N7K1B West DC-xTRs Configuration

The required configuration for the xTRs deployed in the West DC is shown in [Figure 5-2](#).

Figure 5-2 LISP Host Mobility across Subnets Configuration for xTRs in West DC



As it is easy to notice, the LISP configuration is pretty much identical across the two devices part of the same DC site and very similar to the configuration used for LISP Host Mobility with Extended Subnet. Below is the explanation of the various portion of the configuration.

- As first step, it is required to enable the LISP functionality on the Nexus devices and specify that they are going to perform the roles of LISP ETR (for decapsulating LISP traffic received from the L3 domain of the network) and ITR (for encapsulating LISP traffic destined to remote locations).

```

feature lisp
ip lisp itr-etr

```

- A global database mapping is then configured, including an aggregate prefix that ideally identifies all the IP subnets deployed in this specific Data Center site. Notice that this aggregate prefix may include both “mobile subnets” and “static subnets”. The former ones represent the IP subnets where the mobile workloads will be connected. An additional piece of configuration is required to specifically identify the mobile subnets, as discussed in a following section.

```

ip lisp database-mapping 10.17.0.0/16 12.1.1.1 priority 1 weight 50
ip lisp database-mapping 10.17.0.0/16 12.1.1.2 priority 1 weight 50

```

The mapping above associates the aggregate prefix 10.17.0.0/16 to two IP addresses, which are the RLOCs identifying each of the local DC xTR devices. The recommendation is to define a loopback interface on each device as RLOC, so that communication to that IP address will remain successful as long as a valid L3 path connects the xTR to the L3 domain of the network.

Notice also how a priority and a weight can be associated to each mapping statement: these values can be tuned to influence the inbound traffic, preferring for example the path through a specific xTR. In the configuration above the values are identical to ensure that inbound traffic can be load-balanced across both DC xTRs.

**Note**

The definition of the global database-mapping statements are particularly important to enable communication between a remote client and mobile workloads connected to the mobile subnet in the home site, as it will be explained in detail in the “[East-West Traffic Flows Considerations](#)” section on [page 4-19](#).

- The next step consists of defining the IP addresses of the Map-Servers and Map-Resolvers.

```
ip lisp itr map-resolver 1.1.1.100
ip lisp etr map-server 1.1.1.1 key abcd
ip lisp etr map-server 2.2.2.1 key abcd
```

As already mentioned, in a typical Enterprise deployment, two devices will perform the roles of MS/MR and work in a complete stateless fashion. As a consequence, on the xTRs we need to specify the IP addresses of the two Map-Servers (so that each xTR can register with both MS the EID prefixes) and the Anycast IP address of the Map-Resolver (so that the Map-Requests will be received by the MR that is closer from a routing perspective).

- A dynamic mapping is then required to identify the IP subnets to which the mobile workloads belong.

```
lisp dynamic-eid MOBILE_SUBNET
  database-mapping 10.17.1.0/24 12.1.1.1 priority 1 weight 25
  database-mapping 10.17.1.0/24 12.1.1.2 priority 1 weight 25
  map-notify-group 239.1.1.1
```

In this example, the mobile subnet is a /24 prefix, which is associated to the same two RLOCs previously used for the global mapping. Priorities and weights are kept the same also in this case, to achieve inbound load balancing for traffic destined to the mobile subnet. A multicast address (named “map-notify-group”) must also be associated to the dynamic-eid mapping. Its use will be clarified in the following sections of the document.

Some additional considerations around the length of the network prefix specified in the dynamic-eid mapping:

- If multiple mobile subnets are configured, it is possible to define a different “lisp dynamic-eid” construct for each subnet or to use a coarser prefix. The multicast address of the map-notify-group can be the same across multiple constructs.
- The mask associated to the dynamic-eid prefix should always be more specific than the one used in the global mapping statements.
- Differently than in LISP with Extended Subnet mode, there is no specific relationship between the length of the mask associated to the dynamic-eid prefix and the network mask of the interface where LISP mobility is enabled. It is anyway common practice to match the two values, as shown in the example here under discussion.

- Finally, the LISP configuration must be applied under the L3 interface connecting to the mobile subnet. Since the DC xTR is positioned at the aggregation layer, the L3 interface is usually a VLAN Interface (SVI). Notice how this is the only piece of configuration that is different between the two xTRs belonging to the same site (because of IP addressing and HSRP commands).

N7K1A

```
interface vlan 1301
  ip address 10.17.1.3/24
  lisp mobility MOBILE_SUBNET
  ip proxy-arp
  hsrp 1
    mac-address 0000.aaaa.bbbb
    preempt delay reload 300
    priority 110
    ip 10.17.1.1
```

N7K1B

```
interface vlan 1301
  ip address 10.17.1.2/24
  lisp mobility MOBILE_SUBNET
  ip proxy-arp
  hsrp 1
    mac-address 0000.aaaa.bbbb
    ip 10.17.1.1
```

The “lisp mobility” command is used to attach the dynamic-eid construct to this interface, whereas “ip proxy-arp” is required to permit workloads to communicate successfully after the migration to the remote data center site (this mechanism will be discussed in detail in the following sections). Notice the static configuration of the vMAC as part of the HSRP configuration. The need for this is also clarified in the following sections.


**Note**

No specific multicast configuration (like for example enabling PIM) is required under the SVI. Simply enabling LISP Host Mobility on the interface ensures that map-notify-group multicast frames can be sent and received successfully.

Nexus 7000 N7K2A and N7K2B East DC-xTRs Configuration

The required configuration for the xTRs deployed in the East DC is shown in [Figure 5-3](#).

Figure 5-3 LISP Host Mobility across Subnets Configuration for xTRs in East DC



The explanation of the various part of the configuration has been already done in the previous section. The few things to notice when comparing it with the West DC xTRs are the following:

- The global mapping must be different from the one configured on the West DC xTRs. Again, the assumption here is that the IP subnets deployed in the East site can be aggregated by a unique IP prefix (10.18.0.0/16 in this case). Also, the RLOC associated to the global prefixes are now identifying the xTR devices in the East DC site.
- The prefix in the dynamic-eid mapping identifying the mobile subnet **must be identical** to the one defined on the West xTRs, since it identifies the IP subnet where the mobile workloads are connected. This is the reason why 10.17.1.0/24 is specified also on the xTRs in the East DC. However, the RLOCs associated to the mobile subnet are now identifying the xTRs in the East site.



Note The same considerations would apply if we wanted to enable migration of workloads connected to the East DC as home site (for example belonging to the IP subnet 10.18.1.0/24).

- Despite the fact that mobile workloads from the 10.17.1.0/24 subnet can be moved to the East DC site and retain their original IP addresses, there is no need to define an SVI belonging to 10.17.1.0/24 in the East DC. This is a consequence of the true IP Mobility functionality across L3 network domains enabled by LISP.
- The map-notify-group associated to the dynamic-eid mapping must be different to the one configured for the xTRs in the West site. This is because, differently from what discussed for LISP Host Mobility with Extended Subnet Mode, the multicast communication on this group must be limited between the LISP xTR devices connected to the same local DC site.

■ LISP Host Mobility Across Subnets: Sample Config

- Notice how the HSRP group 2 configured on the LISP xTRs in the East DC is different from the one (HSRP group 1) used in the West DC site. In this case, it is recommended to statically configure the **same** HSRP vMAC (000.aaaa.bbbb in the example above), to ensure that a migrated workload can consistently send traffic to the default gateway. However, since the deployment of LISP Host Mobility Across Subnets is currently positioned for “cold migration” scenarios, it is likely that the ARP cache of the migrated workload is initially empty. In this case, just having proxy ARP enabled on the SVI 1301 would ensure that the workload can successfully resolve the MAC address of the localized gateway, making the static vMAC configuration optional.




Note To cover all the scenarios, it is recommended to have a consistent vMAC across different DC sites. The easiest way to achieve that is by configuring the same HSRP group everywhere, as discussed for the LISP Mobility with Extended Subnets scenario.

Remote Site Cisco IOS-xTR Configuration

The configuration of the branch xTR is shown in [Figure 5-4](#).

Figure 5-4 Remote xTR IOS Configuration



Compared to the DC xTR, the configuration for the remote xTR is very simple, since there are usually no Host Mobility requirements for the EIDs belonging to remote locations. The explanation of the different commands is almost self-explanatory. Notice how IOS requires that the LISP configuration is added under a “router lisp” construct, in a similar fashion on how a routing protocol is usually enabled.

- Define the EID space where the clients that will communicate to the DC workloads are deployed.

```
database-mapping 10.10.10.0/24 12.0.0.1 priority 1 weight 100
```

The RLOC address associated to the EID prefix may be a loopback address (as recommended for the DC xTR devices) or, in scenarios where the remote xTR connects to separate SP connections, the IP address of the physical links toward the providers may be used as RLOCs. This last option is usually recommended when it is desirable to tune the priority and weight parameters associated to each RLOC, to influence the inbound traffic policies.

- Configure the Map-Servers and the Map-Resolver Anycast address.

```
ipv4 itr map-resolver 1.1.1.100  
ipv4 etr map-server 1.1.1.1 key abcd  
ipv4 etr map-server 2.2.2.1 key abcd
```




Note When deploying redundant xTR devices at the remote locations, multiple RLOCs are usually associated to the same EID subnet, similarly to what previously shown for the DC xTRs.

- Enable the ITR and ETR functionalities on the device.

```
ipv4 itr  
ipv4 etr
```

NX-OS Map-Server and Map-Resolver Configuration

Considerations around the recommended MS/MR deployment options in LISP Enterprise deployments have been already discussed in the previous “[Map-Server and Map-Resolver Deployment Considerations](#)” section on page 3-6. Figure 5-5 shows the configuration required when deploying the MS/MR on NX-OS platforms also configured as LISP DC xTR devices.

Figure 5-5 NX-OS and IOS MS/MR Configurations

Notice how the configuration on the two MS/MR devices is basically identical, with the only exception of the IP address used as Map-Server identifier. The different parts of the NX-OS configuration are explained below (the equivalent IOS configuration is also shown, in scenarios where dedicated standalone MS/MR are deployed).

Step 1 Enable the MS and MR functionalities on the device.
NX-OS

```

feature lisp
ip lisp map-resolver
ip lisp map-server

```

IOS

```

router lisp
  ipv4 map-server
  ipv4 map-resolver

```

Step 2 Define the Loopback interfaces used as IP addresses for the Map-Resolver and Map-Server functions.
NX-OS

```

interface loopback0
description Map-Resolver IP Address
ip address 1.1.1.100/32
!
interface loopback1
description Map-Server IP Address

```

```
ip address 1.1.1.1/32
```

IOS

```
interface loopback0
  description Map-Resolver IP Address
  ip address 1.1.1.100 255.255.255.255
!
interface loopback1
  description Map-Server IP Address
  ip address 1.1.1.1 255.255.255.255
```

Both Map-Resolvers in [Figure 5-5](#) are configured with the same IP address (Anycast IP address), so that Map-Requests originated from LISP ITR devices can be received on the MR device that is “closer” from a routing table point of view. A unique IP address is instead leveraged for the Map-Server, because the LISP ETRs must register their EID subnets with both standalone Map-Servers.

Step 3 Configure the remote branch site.**NX-OS**

```
lisp site BRANCH
  eid-prefix 10.10.10.0/24
  authentication-key 0 abcd
```

IOS

```
router lisp
  site BRANCH
    authentication-key abcd
    eid-prefix 10.10.10.0/24
```

Step 4 Configure the West and East Data Center sites.**NX-OS**

```
lisp site WEST_DATA_CENTER
  eid-prefix 10.17.0.0/16 accept-more-specifics
  authentication-key 0 abcd
!
lisp site EAST_DATA_CENTER
  eid-prefix 10.18.0.0/16 accept-more-specifics
  authentication-key 0 abcd
```

IOS

```
site WEST_DATA_CENTER
  authentication-key abcd
  eid-prefix 10.17.0.0/16 accept-more-specifics
!
site EAST_DATA_CENTER
  authentication-key abcd
  eid-prefix 10.18.0.0/16 accept-more-specifics
```

It is important to notice the “accept-more-specifics” keyword associated to the DC EID prefix. This needs to be configured for the sites where LISP Host Mobility is enabled, since specific /32 prefixes part of the larger aggregate prefix will be registered by the DC xTRs. The reasoning behind this behavior will be clarified in detail in the following section.

Remote Clients Communicating to EIDs before a Mobility Event


Assuming the LISP configuration previously described has been applied to the various devices, let's now clarify how traffic flows to and from the mobile workloads can be established. Please refer to the previous chapter for more details on the show commands that can be used to verify that the mapping database has successfully registered the LISP EID subnets.

[Figure 5-6](#) highlights the content of the routing tables on the DC xTRs. It is worth noting that the routing table on the “home site” has no special entries installed by LISP. However, note that on the East DC xTRs, a /24 Null0 router (owned by the LISP process) is installed (for reasons described in detail below).




Note In all the following network diagrams, the MS/MR is generically represented as a standalone device connected to the core, to ease the explanation of the discussed functionalities.

Figure 5-6 Routing Table Content in DC xTRs Before Workload Migration



The West DC represents the home site for the mobile IP subnet 10.17.1.0/24. As a consequence, the only route present in the routing table of the local xTRs is the directly connected /24 prefix. This has two immediate consequences:

1. The LISP xTRs does not dynamically discover workloads that are connected to the home site and source IP traffic in the northbound direction (as instead happens in Extended Subnet Mode). As we'll clarify later in this section, the only exception to this behavior is when a workload is initially moved back to the home site from a remote data center. In this latter case, the workload is going to be dynamically discovered and temporarily registered in the local dynamic-eid table and with the Map-Server.
2. Traffic originated from the Layer 3 domain (for example from a remote client) and destined to a host belonging to the 10.17.1.0/24 subnet, will be directed by LISP toward the West DC, decapsulated and routed directly to the locally attached EID subnet. In the current LISP Host Mobility Across Subnets implementation this is a basic principle to keep in mind: communication to a workload belonging to a LISP mobile subnet and connected to the home site can be established without the need to dynamically discover that EID prefix ([Figure 5-7](#)).


Figure 5-7 Establishing Client Initiated Traffic Flows

The example above focuses only on the client-to-server traffic flows. For what concerns the return traffic, two different scenarios are possible:

1. The branch subnets where the clients are deployed are injected in the core of the network. In this case, the DC xTRs will receive routing information about the branch subnets and as a consequence traffic will be natively routed.
2. The branch subnets are EID as well (hence not injected in the routing protocol running in the core of the network). In this case, communication between the DC EIDs and the clients must happen through LISP. The mechanism is similar to the one discussed above, with the only difference that now the remote LISP devices become ETRs and the DC devices play the role of ITRs. In this case, the DC xTRs have to populate their map-cache tables to be able to encapsulate traffic to the remote locations ([Figure 5-8](#)).

■ Remote Clients Communicating to EIDs after a Mobility Event

Figure 5-8 Establishing of Server to Clients Communication




Note Server-client traffic flows will also be balanced across both of DC xTRs, assuming they are connected to the edge of the network leveraging an active/active mechanism like vPC. In a STP based POD topology, all the traffic flows will be instead handled by the same xTR (the HSRP active one).

Remote Clients Communicating to EIDs after a Mobility Event

Figure 5-9 highlights the sequence of events to move a workload from the West DC to the East DC and the consequent updates of the LISP data structures on the MS and on the DC xTRs.

Figure 5-9 Update of LISP Data Structures after Workload Migration



1. The workload is migrated from the West DC (with subnet 10.17.1.0/24) to the East DC (with subnet 10.18.1.0/24). The workload VM retains its IP address and MAC address, and sources an IP packet that reaches one of the two DC xTR devices. This triggers an “URPF-like” failure event, since the packet is received on a local SVI belonging to the 10.18.1.0/24 subnet and the source of the packet is from the 10.17.1.0/24 subnet. The result of the check failure is that the packet is punted to the CPU, causing the dynamic discovery of the EID.



Note IP packets sourced by the migrated workload are steered toward the local gateway because of the proxy-ARP function (if the ARP cache of the workload is empty after the move is completed) or because a consistent vMAC/VIP pair is locally provided (in case the assumption is that the ARP cache of the workload after the migration retains the information populated while in the original home site).

2. The xTR that discovered the EID installs in routing table a local /32 route associated to the EID. This is important to allow exchange of traffic (like ARP) to happen with the workload, since that traffic would normally be prevented by the /24 Null0 entry.
3. The discovering xTR sends out a Map-Notify-Group message on the interface (SVI 1301) where the discovery of the EID happened.
4. The multicast message reaches the peer xTR in the East DC side, which also install a valid /32 route for the EID.
5. The discovering xTR sends a Map-Register messages for the /32 EID address to the Map-Server.
6. The Map-Server adds to the database the entry for the specific EID, associated to the RLOCs (12.2.2.1 and 12.2.2.2) assigned to the xTRs in the East DC.
7. The Map-Server sends a Map-Notify message to the last xTR in the West DC that registered the 10.17.0.0/16 prefix. This message is to notify the xTR that a specific workload belonging to that subnet has just been discovered in a remote DC site.

8. The xTR in the West DC receives the Map-Notify message from the Map-Server and adds to the routing table a /32 Null0 route associated to the 10.17.1.10 EID.
9. The West DC xTR also notifies its peer by leveraging a site local Map-notify message.
10. The West DC peer xTR also installs in the routing table the /32 Null0 route associated to the 10.17.1.10 EID.


**Note**

Once an EID is discovered in the remote data center site, a periodic liveness check is performed by the xTR that discovered it. The check is performed by pinging the EID every 60 seconds. If three consecutive PINGs are not responded to, the LISP xTR removes the EID from its dynamic EID table and stops registering it with the Map-Server. This is done to handle scenarios where EIDs are silently disconnected from the network.

The process above allows updating the information in the DC xTR devices and in the mapping database. To establish successful communication between the remote client and the migrated workload it is necessary to complete a further step: updating the map-cache information in the map-cache of the remote xTR devices. Even after the move of the workload, the map-cache of the remote xTR may still have old mapping information, associating EIDs part of the global /16 prefix with the RLOCs of the xTRs in the West site. This is the case for example if communication between the client and the workload was started before the move.


The consequence is that data packet destined to the migrated EID, will be sent toward the West DC. Once one of the xTRs in that site receives it and decapsulates the first packet, it will perform a routing lookup and find that the destination address is associated to the Null0 route installed at steps 8 and 10 above. Because of this, the packet will be dropped and punted to the CPU to be handled by the LISP process. This will allow the steps shown in [Figure 5-10](#) to happen:

Figure 5-10 Updating the Remote xTR Map-Cache Entry



1. The LISP process on the xTR receiving the first data packet creates a control plane message (called Solicit-Map-Request – SMR) and sends it to the remote ITR that generated the packet. This is to inform the remote ITR that there is a need to refresh its map-cache information because the destination workload has been moved to a different location. It is important to notice that for the SMR message to be created it is mandatory to have a /32 Null0 entry associated to the EID, since without that traffic will be natively routed to the 10.17.1.0/24 home subnet and subsequently dropped. Hence, it is critical to verify that the xTRs in the original site have the entry populated once the EID is discovered in the East DC.
2. The remote ITR receives the SMR and send a new Map-Request for the desired destination (10.17.1.10) to the Map-Server. The Map-Request is forwarded by the Map-Server to the DC xTR in the East site that registered last the /32 EID address.
3. The DC xTR in the East DC replies with updated mapping information to the remote ITR.
4. The remote ITR updates the information in its map-cache, adding the specific /32 EID address associated to the xTRs deployed in the East site (12.2.2.1 and 12.2.2.2).
5. Traffic is now optimally steered toward the East DC site.

Per Figure 5-9, once a workload is discovered in the remote DC site, a corresponding /32 Null0 route is added to both LISP xTRs in the home site (West DC is the home site for 10.17.1.0/24 in our example). The presence of this route forces the dynamic discovery of the EID once it returns to the home site. The sequence of events happening in this case is described in Figure 5-11.

Figure 5-11 Workload Moving back to the Home Site

1. The workload is migrated from the East DC back to the home subnet (10.17.1.0/24) in the West DC. The workload VM retains its IP address and MAC address, and sources an IP packet that reaches one of the two DC xTR devices. This triggers an “URPF-like” failure event, because of the existence of the /32 Null0 entries associated to the IP address of the workload. As a result the packet is punted to the CPU, causing the dynamic discovery of the EID.
2. The xTR that discovered the EID installs in routing table a local /32 route associated to the EID. This is important to allow exchange of traffic (like ARP) to happen with the workload, since that traffic would normally be prevented by the /24 Null0 entry.
3. The discovering xTR sends out a Map-Notify-Group message on the interface (SVI 1301) where the discovery of the EID happened.
4. The multicast message reaches the peer xTR in the East DC side, which also install a valid /32 route for the EID.
5. The discovering xTR sends a Map-Register messages for the /32 EID address to the Map-Server.
6. The Map-Server adds to the database the entry for the specific EID, associated to the RLOCs (12.1.1.1 and 12.1.1.2) assigned to the xTRs in the West DC.
7. The Map-Server sends a Map-Notify message to the xTR in the East DC that last registered the 10.17.1.10/32 prefix. This message is to notify the xTR that workload has just been discovered in a different DC site.
8. The xTR in the East DC receives the Map-Notify message from the Map-Server and remote the valid /32 route associated to 10.17.1.10 from its routing table.
9. The East DC xTR also notifies its peer by leveraging a site local Map-notify message.
10. The East DC peer xTR also removes the /32 valid route associated to the 10.17.1.10 EID.

It is worth noticing that the discovery of the EID in the home site is mostly needed to ensure a proper update of the remote xTRs routing tables, mapping database, remote ITR map-caches, etc. However, as previously mentioned, communication to EIDs attached to the home subnet is normally allowed without

requiring installing host routes in the routing table of the local xTRs. As a consequence, the existence of the host routes highlighted in [Figure 5-11](#) is only temporary: a timer (3 minutes) is started after the dynamic discovery of the EID and at its expiration, the EID 10.17.1.10 is removed from the dynamic-table of the local xTRs, the corresponding host route from their routing tables and the IP address is also removed from the mapping database. New ITRs needing to send traffic to that EID would be allowed to do so by following the global 10.17.0.0/16 prefix (as initially highlighted in [Figure 5-7](#)).


East-West Traffic Flows Considerations

The discussion in the previous sections focused on the establishment of traffic flows between clients situated in remote locations (behind an ITR) and a workload in the DC (before and after a workload migration). The next step is clarifying how communication can instead happen between a workload migrated to the East DC and resources still deployed in the original West site. From that point of view, there are two types of traffic flows to consider: intra-subnet and inter-subnets.

Intra-Subnet Traffic Flows

The establishment of intra-subnet communication between two workloads part of the same IP subnet but connected in separate DC sites is shown in [Figure 5-12](#), respectively for the West-to-East and East-to-West directions.

Figure 5-12 Intra-Subnet Communication between DC Sites



- **West to East:** the initial assumption is that the host X does not have any information for Y in its ARP table, so that it will send out a new ARP request for it. This will allow proxy-ARP enabled on the SVI 1301 to provide the vMAC of the default gateway as response. Notice that for this to happen, it is critical that the xTRs in the West DC have the 10.17.1.10/32 route associated to Null0 (otherwise proxy-ARP would not reply because the 10.17.10/24 subnet is locally connected to the SVI).

The above assumption is surely valid if no communication was established between the two machines before the migration of Y to the East site. In cases where X and Y started communicating in the West DC site, the specific ARP entry associated to Y will be removed from X's ARP cache by a Gratuitous ARP generated by the West DC xTR when it receives the Map-Notify message from the Map-Server when Y is discovered in the East site (step 7 of the process shown in [Figure 5-9](#)).

■ East-West Traffic Flows Considerations

In both cases, traffic originated by X and destined for Y would be steered toward the LISP xTRs in the West DC site. At that point, the lookup for the 10.17.1.10 destination will hit the specific /32 Null0 entry as result, causing the packet to be punted to the CPU and triggering the LISP control plane. A Map-request is then sent to retrieve mapping information for Y's address and traffic will start to be encapsulated to one of the RLOCs deployed in the East DC site.

- **East to West:** also in this case the assumption is that Y does not have any information for X in its ARP table. Differently from what has been discussed above, in the current implementation there is no dynamic mechanism to clear that specific entry if it is already available in Y's ARP cache. As a consequence, it is mandatory that the ARP cache of Y be empty once the migration to the East site is completed. **This is one of the main reasons why the deployment of LISP Across Subnets Mode is currently positioned for cold migration scenarios.**


Under the above assumption, Y would ARP to discover X MAC address. The proxy-ARP function enabled on the SVI 1301 will provide the vMAC of the default gateway as response. Notice that for this to happen, it is critical that the xTRs in the East DC site have the 10.17.0.0/24 route associated to Null0 (this route is installed by LISP because of the dynamic-EID mapping configuration). The traffic is then steered to the default gateway on the East DC (one of the two LISP xTR devices), LISP control plane is triggered and traffic is encapsulated and sent to the West DC xTRs.

Inter-Subnets Traffic Flows

Let's assume the desire is to establish inter-subnet communication between a host Z in subnet 10.17.2.0/24 in the West site and the usual workload Y (10.17.1.10) that was migrated to the East DC. There are two scenarios to consider:

1. 10.17.2.0 is also a mobile subnet, which means there is a dynamic-eid mapping for the subnet defined on all the LISP xTRs. Also, in the home location (West DC) there is also a global mapping covering the subnet. This scenario is highlighted in [Figure 5-13](#).

Figure 5-13 Traffic Flows between Workloads Belonging to Mobile Subnets



In this case, when Z wants to send a packet to Y, it first sends it to its local default gateway, positioned on one of the DC xTR in the West site. The xTR tries to route the packet to Y, but hits the Null0 route installed when Y was discovered in the East site. This punts the packet to the CPU

and allows triggering a LISP Map-Request to the Map-Server. Once the xTR receives valid information for Y, it will start encapsulating traffic to it. This means that traffic in the Z-to-Y direction will flow using LISP encapsulation across the L3 infrastructure.

Similarly, when Y tries to send packets back to Z, the LISP control plan is triggered (a /24 Null0 is installed for the subnet of Z in the East DC xTRs, since it is mobile) and data packet are then LISP encapsulated and sent across the L3 infrastructure.


2. 10.17.2.0 is not a mobile subnet, which means no LISP mobility commands are configured under the SVI associated to it.

In this case, the communication in the Z-to-Y direction happens identically to the scenario above. However, in the current implementation, an xTR performs a check on the source IP address before triggering the LISP control plane. This means that Z must be an EID to be able to communicate to Y via LISP. Since the subnet to which Z belongs is not a mobile subnet, this essentially means that Z needs to be part of the global mapping defined with the “ip lisp database-mapping” command. This is indeed an additional reason why the global mapping should always be defined, covering all the non mobile subnets deployed in a specific DC site.


For what concerns the communication in the Y-to-Z direction, two scenarios are possible:

1. Z subnet is advertised in the protocol used to exchange routing information between the West and East DC sites. In this case, the xTR in the East site will have valid routing information for the subnet Z belongs to and traffic will hence be natively routed back (Figure 5-14).

Figure 5-14 Natively Routing Traffic from East to West DC



2. Z subnet is not injected in the routing protocol, because the goal is to make it only reachable via LISP encapsulation. In this case, when the packet from Y reaches the East xTR, a Map-Request for Z will be sent out and the reply from the Map-Server will be the global prefix covering Z subnet (10.17.0.0/16 in our configuration example), associated to the RLOCs of the DC xTRs belonging to the West DC site. Traffic will then be LISP encapsulated and sent toward that site across the L3 network infrastructure and eventually routed to Z (Figure 5-15).

Figure 5-15 LISP Encapsulating Traffic from East to West DC

Summary

In summary, the LISP Host Mobility Across Subnets solution provides the following key benefits:

- Provides automated move detection, map-cache update and provide a direct data Path to the mobile workloads current Location.
- Off-Subnet Connections are maintained after the move.
- Preserving on-Subnet connections across move would require refreshing the ARP cache on the moving workload. Because of that, the across subnet deployment model is currently targeted to "cold" workload migration use cases, like for example Disaster Recovery scenarios.
- No routing re-convergence required during move.
- No DNS updates required.
- Absolutely no host protocol stack changes, OS changes, or configuration changes required.