



External Connectivity—LISP

- [External Connectivity, on page 1](#)

External Connectivity

Layer-3 DCI for VXLAN BGP EVPN fabrics—LISP

VXLAN BGP EVPN data center fabrics can be connected across Layer-3 boundaries using MPLS L3VPN, VRF IP Routing (VRF lite), or LISP as the mechanism of transport outside the VXLAN fabric. The LISP scenario is explained in this chapter.

Some pointers about LISP as the data center interconnect (DCI) technology is given below:

- A Cisco Nexus 7000 (or 7700) Series switch (with F3/M3 line cards) is considered as the border leaf switch in the LISP scenario explained below.
- F3/M3 line cards provide multiple data plane encapsulation in hardware and control plane protocols. VXLAN encapsulation is implemented in hardware on the southbound side, and LISP is implemented in hardware on the northbound side of the F3/M3 cards.



Note For the F3 line card, support was introduced in the Cisco NX-OS Software Release 7.2. For the M3 line card, support was introduced in the 8.2(1) release.

- LISP provides an optimal way to resolve ingress route optimization challenges that result from workload mobility across data centers.
- Unlike other traditional routing protocols, LISP is a PULL based model and only interested entities (switches or routers) ask for information. The PULL based model results in a smaller state in the switch tables and conserves hardware space.
- LISP has other advantages, as noted below:
 - Mobility—IP prefix address family portability.
 - Scalability—On demand routing.
 - Security—Tenant ID based segmentation.

VM mobility across datacenters

VM mobility across VXLAN BGP EVPN datacenter fabrics works the same way as it does within the datacenter fabric. When VM mobility takes place, the VM generates RARP and GARP messages. You should enable a Layer-2 DCI such as OTV, Classical Ethernet or VPLS to transport broadcast RARP and GARP packets generated due to the VM movement.



Note Additional configuration is not required to support VM movement across fabrics.

VM mobility across fabrics cannot take place in these scenarios:

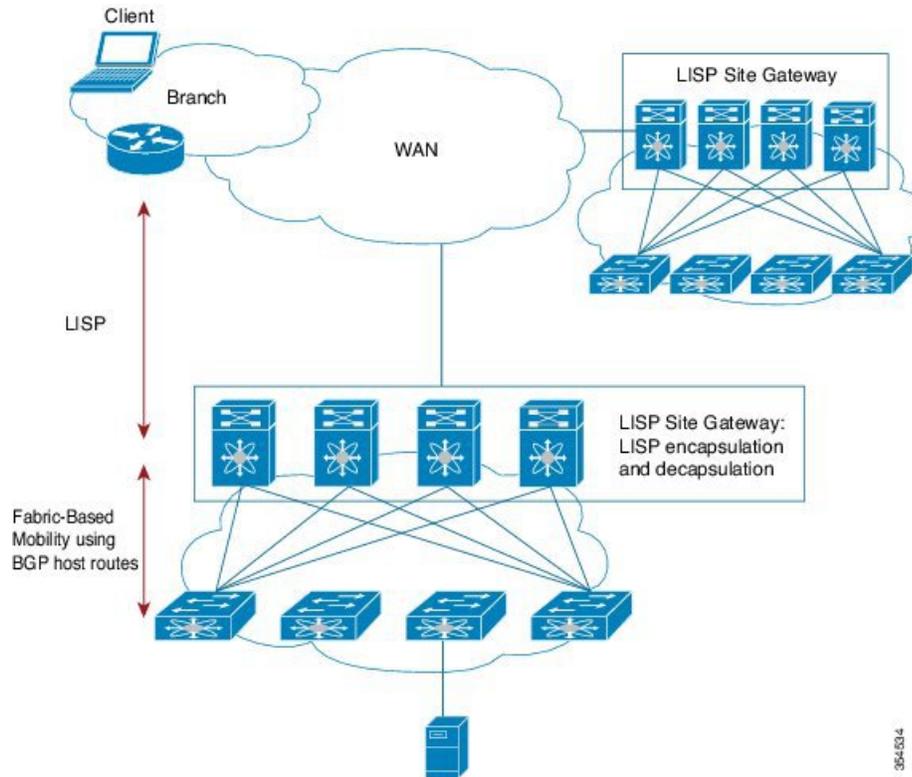
1. VM movement takes place when MAC chaining (multiple IP addresses mapped to the same MAC address) is in effect.
2. When an end host sends a non broadcast packet such as ARP on VM move.

Host Mobility across VXLAN BGP EVPN Fabrics using LISP

Host move detection within the fabric

In the VXLAN BGP EVPN fabric, the host routes and MAC address information are distributed in the MP-BGP EVPN control plane, which means that the fabric itself performs the host detection. The LISP site gateways use these host routes for triggering the LISP mobility encapsulation and decapsulation. LISP, when integrated with the VXLAN BGP EVPN fabric, provides ingress route optimization for traffic from the clients to the data center (Refer figure below).

Figure 1: LISP functional roles in a VXLAN BGP EVPN fabric



Host Mobility

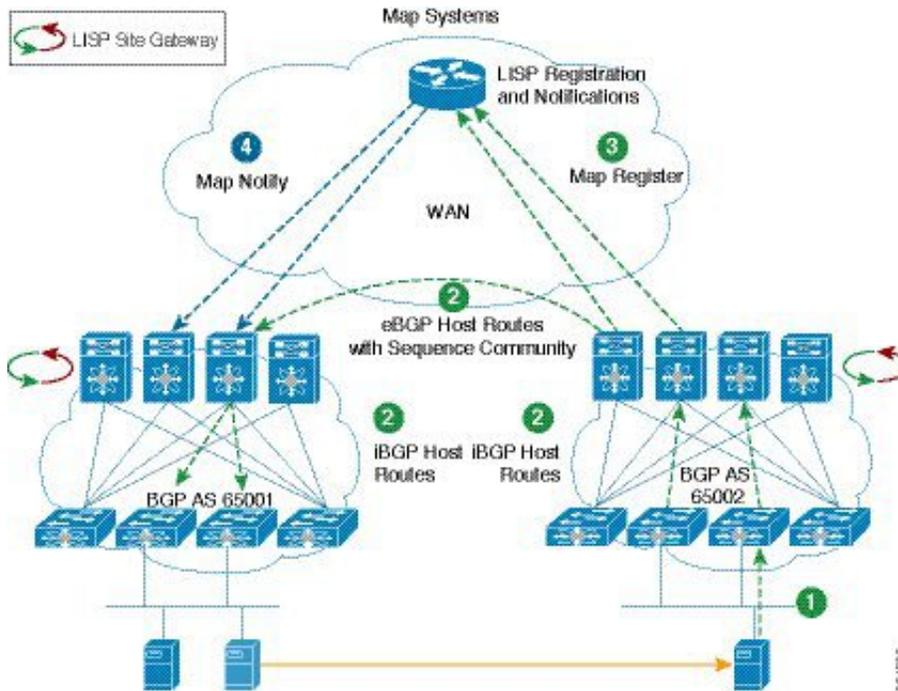
When the leaf or ToR switch detects a host movement across data centers, it injects that host route into the MP-BGP EVPN control plane with an updated sequence number. The sequence number is a mobility community attribute that represents the state of mobility. It increments every time the server moves from one location to another. This sequence number attribute has to be carried to the original leaf or ToR switch from which the host moved, because it needs to withdraw that particular host route from BGP. The host route withdrawal happens only when the leaf or ToR switch receives a route with an updated sequence number. LISP currently cannot carry the mobility community attribute across the data center through the WAN.

To help LISP achieve mobility semantics across VXLAN BGP EVPN fabrics, you need to establish an Exterior BGP (eBGP) relationship between the data centers. This eBGP relationship is used to carry the mobility community attribute in BGP EVPN across the data center sites for the state reachability information. Details are given below.



Note For traffic between hosts that are in different datacenters but within the same subnet, a Layer-2 tunnel (using a Layer-2 DCI like OTV) is required between the two datacenter fabric pods.

Figure 2: Host Mobility across VXLAN BGP EVPN fabrics



The numbered sequence in the above image is elaborated below:

Step 1—The end system or server, after moving to a new location, sends a DHCP and ARP packet to join the new network.

Step 2—The leaf or ToR switch detects the new host and redistributes the IP address and MAC reachability information in the MP-BGP EVPN control plane with an updated sequence number. This sequence number attribute is carried across the data centers using an eBGP relationship between AS 65001 and 65002. When the original leaf or ToR switch receives the route information with an updated sequence number, it withdraws its original route from BGP.

When the host first comes online (before moving across data centers), the sequence number attribute will be 0. This value indicates that this was the first time that the host is coming online in any data center (refer *Host mobility with sequence number 0* image below).

Table 1: Host mobility with sequence number 0

MAC	IP	VNI	Next-Hop	Encap	Seq
MAC1	IPHOST1	5000	VTEP L1	VXLAN	0

After the host moves from one location to another, the sequence number is updated to 1, which triggers the route update through the eBGP connection and the route withdrawal from the original leaf or ToR switch (refer *Host mobility with sequence number 1* image below).

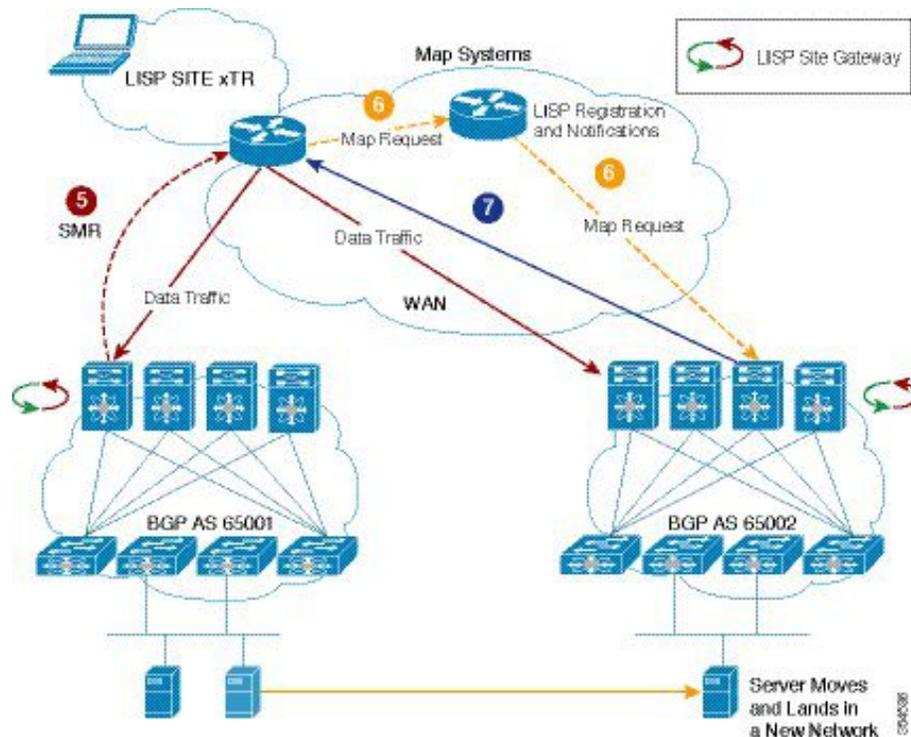
Table 2: Host mobility with sequence number 1

MAC	IP	VNI	Next-Hop	Encap	Seq
MAC1	IPHOST1	5000	VTEP L3	VXLAN	1

Step 3—When the LISP site gateway (also running MP-BGP EVPN in the fabric) detects this new host, it sends a map-register message to the map-system database to register the new IP address in its own data center (BGP AS 65002).

Step 4—When the map system receives the map-register message from BGP, AS 65002 sends a map-notify message to the old LISP site gateways, notifying them that the host has moved from their data center. This message helps ensure that the LISP site gateways install a Null 0 route for that prefix in their routing tables. This Null 0 prefix indicates that the host is in a location remote to that data center.

Figure 3: LISP Map System updates



Step 5—When the clients in the remote branch sites try to send traffic to the LISP site gateways at which the host was present (BGP AS 65001) before the mobility event, the site gateways see that the host is reachable through a Null 0 route. This event triggers a solicit-map request (SMR) from the site gateways to the LISP-enabled router in the branch site asking it to update its database.

Step 6—The branch router then sends a map request to the mapping system asking for the new location of the host. This request is relayed to the LISP site gateways to which the host has moved (BGP AS 65002).

Step 7—The LISP site gateways in BGP AS 65002 unicast a map reply to the LISP-enabled branch router asking it to update its database with the new location.

Now data traffic starts to flow to the correct data center (BGP AS 65002).

Configuration for the VXLAN BGP EVPN – LISP scenario

Hardware and software versions used in the configuration example are given below:

Table 3: Hardware and software versions

Functional Role	Hardware Platform	Software Version
Border spine and border leaf switch	Cisco Nexus 7000 and 7700 Series switches with F3/M3 line card	For the F3 line card, support was introduced in the Cisco NX-OS Software Release 7.2. For the M3 line card, support was introduced in the 8.2(1) release.
Map server and map resolver	Cisco ASR 1000 Series Aggregation Services Routers	Cisco IOS XE Software Release 3.13.2

Border spine switch configuration in Data Center 1 (BGP AS 65001)

This section summarizes the steps for configuring LISP for a hand-off from VXLAN on the border spine or border leaf switch.



Note **Important**—This document contains LISP related configurations on the border spine/leaf switch. You should also enable VXLAN BGP EVPN configurations on this switch. For VXLAN BGP EVPN configuration details, refer the *Forwarding Configurations* chapter, *Cisco Nexus 7000 Series switch configuration* section.

Enable the LISP feature and LISP encapsulation/decapsulation functionality

(config) #

```
feature lisp
ip lisp itr
ip lisp etr
```

Configure LISP map server and map resolver reachability

(config) #

```
ip lisp itr map-resolver 10.8.12.9
ip lisp etr map-server 10.8.12.9 key 0 123456789
```

The above commands configure the LISP map resolver address for the LISP ITR functionality, and the LISP map server address (along with a shared key) for the LISP ETR functionality.

Configure the LISP hand-off for tenant VRF instances—The following example shows a two-tenant VRF instance configuration.

(config) #

```
vrf context evpn-tenant-1
  lisp instance-id 10
  ip lisp locator-vrf default
  lisp dynamic-eid lisp-subnets
    ip lisp database-mapping 10.0.0.0/8 10.8.12.10 priority 1 weight 1
```

```
register-route-notifications tag 65001
```

- The above commands creates a LISP instance ID (which needs to be the same across data centers) and defines the RLOC VRF to use (The mapping database is reachable through that RLOC VRF).
- The tag in the route notification (65001) has to match the BGP Autonomous Systems Number that datacenter 1 and the spine switch belong to.
- The **database-mapping** command configures an EID-to-RLOC mapping relationship and its associated traffic policy for the LISP site.
- The **register-route-notifications** command triggers mobility registration on reception of host routes with the tag configured.

A sample configuration of the second tenant VRF is given below

(config) #

```
vrf context evpn-tenant-2
  lisp instance-id 20
  ip lisp locator-vrf default
  lisp dynamic-eid lisp-subnets
    ip lisp database-mapping 10.21.0.0/8 10.8.12.10 priority 1 weight 1
  register-route-notifications tag 65001
```

If you need to configure additional EID (IP address) subnets to map to the VRF instance, then you will have to create another dynamic EID subnet name. A sample configuration is given below.

(config) #

```
vrf context evpn-tenant-2
  lisp dynamic-eid lisp-subnets-1
    ip lisp database-mapping 209.165.200.225/24 10.0.0.2 priority 1 weight 1
  register-route-notifications tag 650001
```

The LISP instance ID provides a means of maintaining unique address spaces in the control and data plane. Instance IDs are numerical tags defined in the LISP canonical address format (LCAF). The instance ID has been added to LISP to support virtualization.

When multiple organizations within a LISP site are using private addresses as EID prefixes, their address spaces must remain segregated to prevent address duplication. An instance ID in the address encoding can be used to create multiple segmented VPNs within a LISP site at which you want to keep using EID-prefix based subnets. The LISP instance ID is currently supported in LISP ingress tunnel routers and egress tunnel routers (ITRs and ETRs), map server (MS), and map resolver (MR).

The LISP locator VRF is used to associate a VRF table through which the routing locator address space is reachable with a router LISP instantiation.

Border Leaf Configuration in Data Center 2 (BGP AS 65002)

Configuration of the border leaf switch is similar to that of the border spine switch. A consolidated configuration is given below.

(config) #

```

feature lisp
ip lisp itr
ip lisp etr
ip lisp itr map-resolver 10.8.2.45
ip lisp etr map-server 10.8.2.45 key 0 123456789
vrf context evpn-tenant-1
  lisp instance-id 10
  ip lisp locator-vrf default
  lisp dynamic-eid lisp-subnets
    ip lisp database-mapping 10.0.0.0/8 10.8.2.46 priority 1 weight 50
    register-route-notifications tag 65002
vrf context evpn-tenant-2
  lisp instance-id 20
  ip lisp locator-vrf default
  lisp dynamic-eid lisp-subnets
    ip lisp database-mapping 10.21.0.0/8 10.8.2.46 priority 1 weight 50
    register-route-notifications tag 65002

```

The tag in the route notification (65002) has to match the BGP Autonomous Systems Number that datacenter 2 and the border leaf switch belong to.

LISP Map-System Database Configuration

Configure the map server and map resolver on the switch. The map server and map resolver can be either on the same device or multiple devices. The scenario here uses an ASR 1000 Series router as the map server and map resolver.

(config) #

```

router lisp
lisp site DC
  authentication-key shared-key
  eid-prefix instance-id 10 10.0.0.0/8 accept-more-specifics
  eid-prefix instance-id 10 10.21.0.0/8 accept-more-specifics

```

The above commands defines the data center administrative scope, and maps the EID prefixes from the data center sites to the LISP mapping system.

(config) #

```

lisp site branch
  authentication-key shared-key
  eid-prefix instance-id 10 10.12.0.0/8 accept-more-specifics

```

The above commands define the branch location administrative scope.

Configure the branch site

(config) #

```

router lisp
  eid-table default instance-id 10
  database-mapping 10.1.0.0/24 10.100.0.1 priority 1 weight 50

router lisp
  ipv4 itr map-resolver 10.9.9.9

```

```
ipv4 etr map-server 10.9.9.9 key s3cr3t-k3y
```

The above commands configure the LISP Map resolver and Map server addresses.

(config) #

```
router lisp
  ipv4 itr
  ipv4 etr
```

The above commands configure the device as a LISP ITR and ETR.

Verification

To check for the EID (host IP address) learnt on the LISP site gateway on a Cisco Nexus 7000 Series or 7700 platform, use the command shown here.

```
N7700-Border-Spine# show lisp dynamic-eid summary vrf evpn-tenant-1
```

```
LISP Dynamic EID Summary for VRF " evpn-tenant-1"
* = Dyn-EID learned by site-based Map-Notify
! = Dyn-EID learned by routing protocol
^ = Dyn-EID learned by EID-Notify
```

Dyn-EID-Name	Dynamic-EID	Interface	Uptime	Last-Packet	Pending-Ping-Count
lisp-subnets	!10.1.1.12	Eth4/1	06:50:21	00:12:12	0
lisp-subnets	!10.1.1.13	Eth4/2	03:20:01	00:10:12	0

In the above sample output, **lisp-subnets** refers to the EID subnet name, **Dynamic-EID** column refers to the End host IP addresses (EID), and **Interface** refers to the interface connecting to the leaf switches.

To check for LISP map-cache entries on the map server, use the command shown below:

```
Map-Server# show lisp site
```

```
LISP Site Registration Information
* = Some locators are down or unreachable
```

Site-Name	Last-Register	Up	Who-Last-Registered	Inst-ID	EID-Prefix
DC	Never	No	---	10	10.0.0.0/8
	00:00:50	Yes	10.8.2.46	10	10.1.1.12/32
	00:00:50	Yes	10.8.2.46	10	10.1.1.13/32
	00:00:40	Yes	10.8.12.10	10	10.1.1.15/32
	00:00:40	Yes	10.8.12.10	10	10.1.1.16/32
	Never	No	---	10	10.21.0.0/8

In the above example, the **Who Last Registered** column refers to the LISP site gateway/RLOC.

