LISP - IPv6 over IPv4 core

« LISP - Basic overlay IPv4 to IPv4 (/workbook/view/service-provider-v4/task/lisp-basic-overlay-ipv4-to-ipv4-Mjk1MQ%3D%3D) | Multicast VPNs - PIM Core Tree (/workbook/view/service-provider-v4/task/multicast-vpns-pim-core-tree-Mjg5MA%3D%3D) »

Last updated: April 22, 2016

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **LISP**, which can be found in CCIE SPv4 Topology Diagrams & Initial Configurations (http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs). Refer to the **LISP Diagram** in order to complete this task.

Task

› CONTENTS

- Configure a LISP overlay to provide IPv6 connectivity between Site A and Site B over the provider IP core.
- Routing throughout the core network and sites has been pre-configured. There is no need to change routing protocols or add static routes throughout this lab.
- Configure R9 as the LISP Map-Server and Map-Resolver. R9 should allow IPv6 EID registrations from the LISP sites as follows:
  - LISP Site A:
    - EID 2001:10:1::/48 - allow more specifics
    - EID 2001:1:1:1::1/128
    - Use authentication key "SITEA_xTR-R2"
  - LISP Site B:
    - EID 2001:10:19::/48 - allow more specifics
    - EID 2001:20:20:20::20/128
    - Use authentication key "SITEB_xTR-XR1"
- Configure R2 as the LISP xTR router for Site A as follows:
  - Use the Loopback0 of R9, 9.9.9.9, as the map-server and map-resolver. Authenticate with key "SITEA_xTR-R2", as previously configured on R9.
  - Register the following EIDs:
    - 2001:10:1:2::/64
    - 2001:1:1:1::1/128
  - R2 should use both of its links into the SP Core as active RLOCs.
    - Gig1.23 - 20.2.3.2
      - priority 1
      - weight 50
    - Gig1.24 - 20.2.4.2
      - priority 1
      - weight 50

    - Ensure R2 registers the site's EIDs using both RLOCS.
- Configure XR1 as the LISP xTR router for Site B as follows:
    - Use the Loopback0 of R9, 9.9.9.9, as the map-server and map-resolver. Authenticate with key "SITEB_xTR-XR1", as previously configured on R9.
    - Register the following EIDs:
        - 2001:10:19:20::/64
        - 2001:20:20:20::20/128
    - XR1 should use both of its links into the SP Core as active RLOCs:
        - Gig0/0/0/0.519 - 20.5.19.19.
            - priority 1
            - weight 50
        - Gig0/0/0/0.619 - 20.6.19.19
            - priority 1
            - weight 50
    - Ensure XR1 registers the site's EIDs using both RLOCS.
- R1 and XR2 should have IPv6 reachability between their Loopback0 and connected interfaces by the end of this task.

Configuration **Click to *collapse***

CONTENTS

```
!
! LISP SiteA xTR config
!

R2:
router lisp
 locator-set SITEA_RLOC_SET
  20.2.3.2 priority 1 weight 50
  20.2.4.2 priority 1 weight 50
  exit
 !
 database-mapping 2001::1:1:1:1/128 locator-set SITEA_RLOC_SET
 database-mapping 2001:10:1:2::/64 locator-set SITEA_RLOC_SET
 ipv6 itr map-resolver 9.9.9.9
 ipv6 itr
 ipv6 etr map-server 9.9.9.9 key SITEA_xTR-R2
 ipv6 etr
 exit



!
! Map Server config
!

R9:
router lisp
 site SITE_A
  authentication-key SITEA_xTR-R2
  eid-prefix 2001::1:1:1:1/128
  eid-prefix 2001:10:1::/48 accept-more-specifics
  exit
 !
 site SITE_B
  authentication-key SITEB_xTR-XR1
  eid-prefix 2001:10:19::/48 accept-more-specifics
  eid-prefix 2001::20:20:20:20/128
  exit
 !
 ipv6 map-server
 ipv6 map-resolver
 exit



!
! LISP SiteB xTR config
!

XR1:
router lisp
 address-family ipv6 unicast
 !
 locator-set SITEB_RLOC_SET
  20.5.19.19 priority 1 weight 50
```


CONTENTS

```
 20.6.19.19 priority 1 weight 50
!
eid-table default instance-id 0
 address-family ipv6 unicast
  etr map-server 9.9.9.9 key encrypted 122A2C233729331C1E1969100164
  etr
  itr map-resolver 9.9.9.9
  itr
  database-mapping 2001:10:19:20::/64 locator-set SITEB_RLOC_SET
  database-mapping 2001::20:20:20:20/128 locator-set SITEB_RLOC_SET
 !
!
locator-table default
!
```

Verification

One of the benefits of the layer of indirection LISP uses between EID and RLOC spaces is that address families can be mixed - an xTR can use IPv4 RLOCs to reach IPv6 EIDs, and vice-verse. In this example, xTR routers R2 and XR1 only have IPv4 connectivity into the core network. Notice that only the IPv4 stack has been enabled on R2's Gig1.23 and Gig1.24, and XR1's Gig0/0/0/0.519 and Gig0/0/0/0.619. However, both xTRs are running IPv6 within the site.

CONTENTS ›

```
R2#show ipv6 interface brief
GigabitEthernet1       [up/up]
    unassigned
GigabitEthernet1.12    [up/up]
    FE80::250:56FF:FE9E:35D1
     2002:10:1:2::2
GigabitEthernet1.23    [up/up]
    unassigned
GigabitEthernet1.24    [up/up]
    unassigned
GigabitEthernet1.112   [administratively down/down]
    unassigned
GigabitEthernet1.212   [administratively down/down]
    unassigned
GigabitEthernet2       [up/up]
    unassigned
GigabitEthernet3       [up/up]
    unassigned
GigabitEthernet4       [up/up]
    unassigned
LISP0                  [up/up]
     FE80::21E:7AFF:FE2A:C900
    unnumbered (Loopback0)
Loopback0              [up/up]
    FE80::21E:7AFF:FE2A:C900
     2002::2:2:2:2
```

```
RP/0/0/CPU0:XR1#show ipv6 interface brief
Sat Jul 25 19:49:00.529 UTC
Loopback0              [Up/Up]
    fe80::945:13ff:feb3:804a
     2001::19:19:19:19
MgmtEth0/0/CPU0/0      [Shutdown/Down]
    unassigned
GigabitEthernet0/0/0/0 [Up/Up]
    unassigned
GigabitEthernet0/0/0/0.519 [Up/Up]
    unassigned
GigabitEthernet0/0/0/0.619 [Up/Up]
    unassigned
GigabitEthernet0/0/0/0.1920 [Up/Up]
    fe80::250:56ff:fe9e:59fe
     2001:10:19:20::19
GigabitEthernet0/0/0/0.1921 [Up/Up]
    unassigned
GigabitEthernet0/0/0/0.1922 [Up/Up]
    unassigned
GigabitEthernet0/0/0/1 [Shutdown/Down]
    unassigned
```

R2 and XR1 are running EIGRPv6 with R1 and XR2 respectively, and are each redistributing a default route into EIGRP. Like in our IPv4 example, this is a static default route to null0. The purpose of this route is to give R1 and XR2 enough routing information to exit the site.

```
R2#show ipv6 eigrp neighbors
EIGRP-IPv6 VR(LISP_SITE_A) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   Link-local address:      Gi1.12                   10 02:35:38 1279  5000  0  7
    FE80::250:56FF:FE9E:5913

RP/0/0/CPU0:XR1#show eigrp ipv6 neighbors
Sat Jul 25 19:52:20.846 UTC

IPv6-EIGRP VR(LISP_SITE_B) Neighbors for AS(100) VRF default

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   Link Local Address:      Gi0/0/0/0.1920    11 02:32:20   17   200  0  2
    fe80::250:56ff:fe9e:27ac


R1#show ipv6 route eigrp
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       la - LISP alt, lr - LISP site-registrations, ld - LISP dyn-eid
       a - Application
 EX  ::/0 [170/568320]
     via FE80::250:56FF:FE9E:35D1, GigabitEthernet1.12
D   2002::2:2:2:2/128 [90/10880]
     via FE80::250:56FF:FE9E:35D1, GigabitEthernet1.12
D   2002:10:1:2::/64 [90/15360]
     via FE80::250:56FF:FE9E:35D1, GigabitEthernet1.12


RP/0/0/CPU0:XR2#show route ipv6 eigrp
Sat Jul 25 19:53:47.744 UTC

 D*EX ::/0
      [170/568320] via fe80::250:56ff:fe9e:59fe, 02:31:48, GigabitEthernet0/0/0/0.1920
D    2001::19:19:19:19/128
      [90/10752] via fe80::250:56ff:fe9e:59fe, 02:34:28, GigabitEthernet0/0/0/0.1920
```

R2 and XR1 have empty map-caches before R1 and XR2 attempt to reach each other:

```
R2#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache for EID-table default (IID 0), 1 entries

::/0, uptime: 00:00:09, expires: never, via static  send map-request
  Negative cache entry, action: send-map-request
```

CONTENTS ›

```
RP/0/0/CPU0:XR1#show lisp ipv6 map-cache

Sat Jul 25 19:59:58.624 UTC



LISP IPv6 Mapping Cache for EID-table default (IID 0), 1 entries


::/0, uptime: 00:00:50, expires: never, via static  send map-request

  Negative cache entry, action: send-map-request
```

Lets begin by testing bidirectional IPv6 reachability between R1 and XR2 over the LISP overlay. Notice that just like in our previous example, the first two packets are dropped. This is the time it takes for the xTRs to build their cache entries bidirectionally.

```
R1#ping 2001::20:20:20:20 source 2001::1:1:1:1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2001::20:20:20:20, timeout is 2 seconds:

Packet sent with a source address of 2001::1:1:1:1

 ..!!!

Success rate is 60 percent (3/5), round-trip min/avg/max = 12/12/14 ms


R1#ping 2001:10:19:20::20

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2001:10:19:20::20, timeout is 2 seconds:

 ..!!!

Success rate is 60 percent (3/5), round-trip min/avg/max = 11/13/16 ms
```

Notice that the map-caches are now populated as expected, mapping an IPv6 EID to an IPv4 RLOC. The forwarding model is the same as in our previous example. The only difference is that the EID prefixes are now IPv6 instead of IPv4. LISP has great potential for newer forwarding paradigms due to this indirection - for example, using MAC addresses as EIDs.

```
R2#show ipv6 lisp map-cache

LISP IPv6 Mapping Cache for EID-table default (IID 0), 3 entries


::/0, uptime: 00:03:05, expires: never, via static send map-request

  Negative cache entry, action: send-map-request

 2001::20:20:20:20/128 , uptime: 00:00:14, expires: 23:59:45, via map-reply, complete

  Locator     Uptime    State      Pri/Wgt

   20.5.19.19  00:00:14  up          1/50

   20.6.19.19  00:00:14  up          1/50

 2001:10:19:20::/64 , uptime: 00:00:07, expires: 23:59:52, via map-reply, complete

  Locator     Uptime    State      Pri/Wgt

   20.5.19.19  00:00:07  up          1/50

   20.6.19.19  00:00:07  up          1/50


RP/0/0/CPU0:XR1#show lisp ipv6 map-cache

Sat Jul 25 20:06:45.246 UTC




LISP IPv6 Mapping Cache for EID-table default (IID 0), 3 entries


::/0, uptime: 00:07:37, expires: never, via static send map-request

  Negative cache entry, action: send-map-request

 2001::1:1:1:1/128 , uptime: 00:04:48, expires: 23:55:35, via map-reply, complete

  Locator   Uptime    State      Pri/Wgt

   20.2.3.2  00:04:48  up          1/50

   20.2.4.2  00:04:48  up          1/50

 2001:10:1:2::/64 , uptime: 00:04:41, expires: 23:55:42, via map-reply, complete

  Locator   Uptime    State      Pri/Wgt

   20.2.3.2  00:04:41  up          1/50

   20.2.4.2  00:04:41  up          1/50
```

Note that it is also possible to use IPv6 in the RLOC space. All of the combinations bellow are fully supported by LISP. More than one combination can be used at the same time. For example, we can map IPv4 and IPv6 EIDs to IPv4 RLOCs - allowing Site A and Site B to have both IPv4 and IPv6 connectivity simultaneously.

- IPv4 EID - IPv4 RLOC
- IPv4 EID - IPv6 RLOC
- IPv6 EID - IPv4 RLOC
- IPv6 EID - IPv6 RLOC

CONTENTS

The Map Server sees the registrations from each site as it did previously, the only difference in the address family of the EID.

```
R9#show lisp site
LISP Site Registration Information

Site Name      Last      Up   Who Last            Inst     EID Prefix
               Register       Registered          ID
SITE_A         00:00:35   yes  20.2.3.2                     2001::1:1:1:1/128
               never     no   --                           2001:10:1::/48
               00:00:35   yes  20.2.3.2                     2001:10:1:2::/64
SITE_B         00:00:24   yes  20.5.19.19                   2001::20:20:20:20/128
               never     no   --                           2001:10:19::/48
               00:00:24   yes  20.5.19.19                   2001:10:19:20::/64
```

Each site owns a /48, but is allowed to register more specific subnets within the aggregate, such as the /64 being registered in this example.

```
R9#show lisp site name SITE_B 2001:10:19:20::/64
LISP Site Registration Information

Site name: SITE_B
Allowed configured locators: any
Requested EID-prefix:
  EID-prefix: 2001:10:19:20::/64
    First registered:     04:25:32
    Routing table tag:    0
     Origin:               Dynamic, more specific of 2001:10:19::/48
    Merge active:         No
    Proxy reply:          No
    TTL:                  1d00h
    State:                complete
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
    ETR 20.5.19.19, last registered 00:00:14, no proxy-reply, map-notify
                TTL 1d00h , no merge, hash-function sha1, nonce 0x9896B945-0x400F7C9B
                state complete, no security-capability
                xTR-ID 0x12FEE45B-0xEB6F520B-0xC2FADD50-0x0A415064
                site-ID unspecified
      Locator     Local  State      Pri/Wgt  Scope
       20.5.19.19  yes    up          1/50   IPv4 none
       0.6.19.19  yes    up          1/50   IPv4 none
```

Just like in the IPv4 example, each site builds a local database which it then advertises to the mapping system.

CONTENTS ›

```
R2#show ipv6 lisp database
LISP ETR IPv6 Mapping Database for EID-table default (IID 0), LSBs: 0x3, 2 entries


 2001::1:1:1:1/128, locator-set SITEA_RLOC_SET
  Locator   Pri/Wgt  Source     State
  20.2.3.2    1/50   cfg-addr   site-self, reachable
  20.2.4.2    1/50   cfg-addr   site-self, reachable
2001:10:1:2::/64, locator-set SITEA_RLOC_SET
  Locator   Pri/Wgt  Source     State
  20.2.3.2    1/50   cfg-addr   site-self, reachable
  20.2.4.2    1/50   cfg-addr   site-self, reachable



RP/0/0/CPU0:XR1#show lisp ipv6 database
Sun Jul 26 00:08:31.533 UTC



LISP ETR IPv6 Mapping Database for EID-table default (IID 0), LSBs: 0x3, 2 entries


 2001::20:20:20:20/128, locator-set SITEB_RLOC_SET
  Locator     Pri/Wgt  Source     State
  20.5.19.19    1/50   cfg-addr   site-self, reachable
  20.6.19.19    1/50   cfg-addr   site-self, reachable
2001:10:19:20::/64, locator-set SITEB_RLOC_SET
  Locator     Pri/Wgt  Source     State
  20.5.19.19    1/50   cfg-addr   site-self, reachable
  20.6.19.19    1/50   cfg-addr   site-self, reachable
```

Even though the xTRs don't have native IPv6 reachability towards the provider, they can encapsulate the IPv6 user payload using an outer IPv4 header destined to the RLOC of the remote LISP site.

Notice that R2 has its forwarding tables programmed to send traffic towards 2001::20:20:20:20/128 using RLOCS 20.5.19.19 and 20.6.19.19 in a 1:1 ratio.

CONTENTS >

```
R2#show ipv6 lisp forwarding eid remote 2001::20:20:20:20
Prefix                 Fwd action  Locator status bits
 2001::20:20:20:20/128  encap       0x00000003
  packets/bytes       4/472
  path list 7F2529E34000, 4 locks, per-destination, flags 0x49 [shble, rif, hwcn]
    ifnums:
      LISP0(21): 20.5.19.19, 20.6.19.19
    2 paths
      path 7F2529EE60C0, share 50/50, type attached nexthop, for IPv6
        nexthop 20.5.19.19 LISP0, IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
      path 7F2529EE5E20, share 50/50, type attached nexthop, for IPv6
        nexthop 20.6.19.19 LISP0, IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
    1 output chain
      chain[0]: loadinfo 7F25226450F8, per-session, 2 choices, flags 0005, 6 locks
                  flags [Per-session, for-rx-IPv6]
                  16 hash buckets
                    < 0 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 1 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 2 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 3 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 4 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 5 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 6 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 7 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 8 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    < 9 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <10 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <11 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <12 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <13 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <14 > IPV6 midchain out of LISP0, addr 20.5.19.19 7F2529FF44E8
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                    <15 > IPV6 midchain out of LISP0, addr 20.6.19.19 7F2529FF4308
                          IP adj out of GigabitEthernet1.23, addr 20.2.3.3 7F2529E354E0
                  Subblocks:
                    None
```

You may have noticed that although R2 is sending to both of XR1's RLOCs, all egress traffic is using Gig1.23. LISP optimizes how traffic ingresses a site - egress traffic routing is still based on normal forwarding. Due to our BGP config, R2 uses Gig1.23 to reach 20.5.19.19 and 20.6.19.19.

```
R2#show ip route 20.5.19.19
Routing entry for 20.5.19.0/24
  Known via "bgp 200" , distance 20, metric 0
  Tag 100, type external
  Last update from 20.2.3.3 04:38:32 ago
  Routing Descriptor Blocks:
  * 20.2.3.3, from 20.2.3.3 , 04:38:32 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none

R2#show ip cef 20.5.19.19
20.5.19.19/32
  nexthop 20.2.3.3 GigabitEthernet1.23

R2#show ip route 20.6.19.19
Routing entry for 20.6.19.0/24
  Known via "bgp 200" , distance 20, metric 0
  Tag 100, type external
  Last update from 20.2.3.3 04:38:36 ago
  Routing Descriptor Blocks:
  * 20.2.3.3, from 20.2.3.3 , 04:38:36 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none

R2#show ip cef 20.6.19.19
20.6.19.19/32
  nexthop 20.2.3.3 GigabitEthernet1.23
```

XR1 displays the IPv4 RLOCs as the IPv4 compatible IPV6 addresses, using the ::ffff prefix. You will also notice that although XR1 is encapsulating traffic towards Site-A using both of R2's RLOCs, all traffic exiting Site-B is using XR1's Gig0/0/0/0.519.

```
RP/0/0/CPU0:XR1#show cef ipv6 2001::1:1:1:1/128 detail

Sun Jul 26 00:18:38.341 UTC

 2001::1:1:1:1/128 , version 0, internal 0x4000001 0x0 (ptr 0xa0ef5f74) [1], 0x0 (0x0), 0x0 (0x0)

  Updated Jul 25 20:01:56.476

  Prefix Len 128, traffic index 0, precedence n/a, priority 0

   gateway array (0xa0d49ed8) reference count 2, flags 0x4010, source lisp lo (7), 0 backups

                 [1 type 3 flags 0x10048289 (0xa0e0b758) ext 0x0 (0x0)]

   LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]

    via ::ffff:20.2.3.2, 4 dependencies, recursive [flags 0x1000]

     path-idx 0 NHID 0x0 [0xa1704050 0x0]

      next hop ::ffff:20.2.3.2 via ::ffff:20.2.3.2

     IP ENCAP (3 headers), locks: 4

        Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)

       Parent: DATA_TYPE_ECD_TRACKER, 0xa0b8f17c

        Payload AF/MTU: ipv6/1456

       HDR: IP4

         src: 20.5.19.19

         dst: 20.2.3.2

         prot: 17, ttl: 255, tos: dyn, df: set

       HDR: UDP

         src: dyn, dst: 0x10f5, chksum: disabled (0x0000)

       HDR: LISP

         instance ID: not set, mapv: not set

         lsb: 0x00000003, nonce: not set, echo nonce req: not set

    via ::ffff:20.2.4.2, 4 dependencies, recursive [flags 0x1000]

     path-idx 1 NHID 0x0 [0xa17040c4 0x0]

      next hop ::ffff:20.2.4.2 via ::ffff:20.2.4.2

     IP ENCAP (3 headers), locks: 4

        Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)

       Parent: DATA_TYPE_ECD_TRACKER, 0xa0b8f1e0

        Payload AF/MTU: ipv6/1456

       HDR: IP4

         src: 20.5.19.19

         dst: 20.2.4.2

         prot: 17, ttl: 255, tos: dyn, df: set

       HDR: UDP

         src: dyn, dst: 0x10f5, chksum: disabled (0x0000)

       HDR: LISP

         instance ID: not set, mapv: not set

         lsb: 0x00000003, nonce: not set, echo nonce req: not set


     Weight distribution:

      slot 0, weight 50, normalized_weight 1, class 0

      slot 1, weight 50, normalized_weight 1, class 0


     Load distribution: 0 1 (refcount 1)


     Hash  OK  Interface                 Address

      0     Y   recursive                 ::ffff:20.2.3.2

      1     Y   recursive                 ::ffff:20.2.4.2



RP/0/0/CPU0:XR1#show route 20.2.4.2
```


CONTENTS

```
Sun Jul 26 00:20:50.132 UTC


Routing entry for 20.2.4.0/24
  Known via "bgp 1900", distance 20, metric 0
  Tag 100, type external
  Installed Jul 25 17:09:42.564 for 07:11:07
  Routing Descriptor Blocks
    20.5.19.5, from 20.5.19.5, BGP external
      Route metric is 0
  No advertising protos.



RP/0/0/CPU0:XR1#show route 20.2.3.2
Sun Jul 26 00:20:55.382 UTC


Routing entry for 20.2.3.0/24
  Known via "bgp 1900", distance 20, metric 0
  Tag 100, type external
  Installed Jul 21 00:46:39.070 for 4d23h
  Routing Descriptor Blocks
    20.5.19.5, from 20.5.19.5, BGP external
      Route metric is 0
  No advertising protos.
```

There are multiple ways to solve the egress traffic issue, such as enabling BGP Multipath on the xTRs, however this is outside the scope of LISP.

Lets take a look at a LISP encapsulated packet as it traverses the core network. Shutdown Gig1.24 between R2 and R4 in order to force LISP traffic between R2 and XR1 to traverse R3.

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int g1.24
R4(config-if)#shutdown



R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#int g1.24
R2(config-if)#shutdown
```

Clear the caches on R2 and XR1.

CONTENTS >

```
R2#clear ipv6 lisp map-cache



RP/0/0/CPU0:XR1#clear lis ipv6 map-cache

Sun Jul 26 01:09:16.893 UTC



R3:

ip access-list extended LISP

 permit udp any any eq 4341

 permit udp any any eq 4342

 permit udp any eq 4342 any

 permit udp any eq 4341 any


R3#monitor capture LISP_CAP access-list LISP interface g1.23 both

R3#monitor capture LISP_CAP limit packet-len 1500 packets 100

R3#monitor capture LISP_CAP start



R1#ping 2001::20:20:20:20 source 2001::1:1:1:1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2001::20:20:20:20, timeout is 2 seconds:

Packet sent with a source address of 2001::1:1:1:1

 ..!!!

Success rate is 60 percent (3/5), round-trip min/avg/max = 12/14/18 ms



R3#monitor capture LISP_CAP export ftp://cisco:cisco@169.254.254.1/lisp.v6.in.v4.001.pcap

Writing lisp.v6.in.v4.001.pcap

Exported Successfully

Notice the sequence of events detailed in the packet capture.
```

Notice that the control plane messaging exchange that takes place before the ICMPs are able to flow are very similar to the previous example with IPv4 over IPv4.

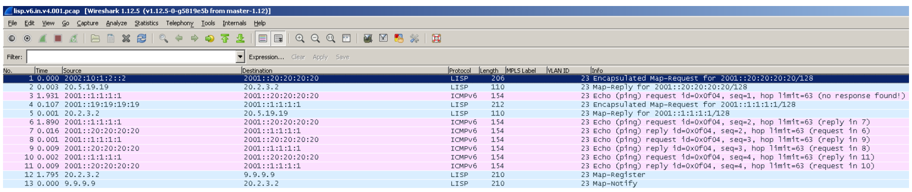

Zooming in on packet #10, we can see the encapsulated ICMP/IPv6 headers from the user payload.

```
10 0.002 2001::1:1:1:1 2001::20:20:20:20 ICMPv6 154   Echo (ping) request id=0x0f04, seq=4, hop limit=63 (reply in 11)
+ Frame 10: 154 bytes on wire (1232 bits), 154 bytes captured (1232 bits)
+ Ethernet II, Src: Vmware_9e:35:d1 (00:50:56:9e:35:d1), Dst: Vmware_9e:6e:6a (00:50:56:9e:6e:6a)
+ 802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 23
- Internet Protocol Version 4, Src: 20.2.3.2 (20.2.3.2), Dst: 20.5.19.19 (20.5.19.19)
    Version: 4
    Header Length: 20 bytes
  + Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00: Not-ECT (Not ECN-Capable Transport))
    Total Length: 136
    Identification: 0x009f (159)
  + Flags: 0x02 (Don't Fragment)
    Fragment offset: 0
    Time to live: 255
    Protocol: UDP (17)
  + Header checksum: 0x3caa [validation disabled]
    Source: 20.2.3.2 (20.2.3.2)
    Destination: 20.5.19.19 (20.5.19.19)
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
- User Datagram Protocol, Src Port: 1090 (1090), Dst Port: 4341 (4341)
    Source Port: 1090 (1090)
    Destination Port: 4341 (4341)
    Length: 116
  + Checksum: 0x0000 (none)
    [Stream index: 3]
- Locator/ID Separation Protocol (Data)
  + Flags: 0xc0
    Nonce: 4964621 (0x4bc10d)
    0000 0000 0000 0000 0000 0000 0000 0001 = Locator-Status-Bits: 0x00000001
- Internet Protocol Version 6, Src: 2001::1:1:1:1 (2001::1:1:1:1), Dst: 2001::20:20:20:20 (2001::20:20:20:20)
  + 0110 .... = Version: 6
  + .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00000000
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 60
    Next header: ICMPv6 (58)
    Hop limit: 63
    Source: 2001::1:1:1:1 (2001::1:1:1:1)
    [Source Teredo Server IPv4: 0.0.0.0 (0.0.0.0)]
    [Source Teredo Port: 65534]
    [Source Teredo Client IPv4: 255.254.255.254 (255.254.255.254)]
    Destination: 2001::20:20:20:20 (2001::20:20:20:20)
    [Destination Teredo Server IPv4: 0.0.0.0 (0.0.0.0)]
    [Destination Teredo Port: 65503]
    [Destination Teredo Client IPv4: 255.223.255.223 (255.223.255.223)]
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
+ Internet Control Message Protocol v6
```

« LISP - Basic overlay IPv4 to IPv4 (/workbook/view/service-provider-v4/task/lisp-basic-overlay-ipv4-to-ipv4-Mjk1MQ%3D%3D) | Multicast VPNs - PIM Core Tree (/workbook/view/service-provider-v4/task/multicast-vpns-pim-core-tree-Mjg5MA%3D%3D) »

CONTENTS ›